

CONTRIBUTIONS TO STATISTICS

Alessandra Giovagnoli

Anthony C. Atkinson

Bernard Torsney

Editors

Caterina May

Co-Editor

mODa 9 – Advances in Model-Oriented Design and Analysis



Physica-Verlag

A Springer Company

Contributions to Statistics

For further volumes:

<http://www.springer.com/series/2912>

Alessandra Giovagnoli · Anthony C. Atkinson ·
Bernard Torsney
Editors

Caterina May
Co-Editor

mODa 9 – Advances in Model-Oriented Design and Analysis

Proceedings of the 9th International
Workshop in Model-Oriented Design
and Analysis held in Bertinoro, Italy,
June 14–18, 2010



Physica-Verlag

Editors

Prof. Alessandra Giovagnoli
University of Bologna
Dipt. di Scienze Statistiche
Via Belle Arti 41
40126 Bologna
Italy
alessandra.giovagnoli@unibo.it

Prof. Anthony C. Atkinson
London School Economics
and Political Science
Department of Statistics
Houghton Street
WC2A 2AE London
United Kingdom
a.c.atkinson@lse.ac.uk

Dr. Bernard Torsney
University of Glasgow
Department of Statistics
University Gardens 15
G12 8QW Glasgow
United Kingdom
b.torsney@stats.gla.ac.uk

Co-Editor

Caterina May
Università del Piemonte Orientale
SEMEQ
via Perrone 18
28100 Novara
Italy
caterina.may@eco.unipmn.it

ISSN 1431-1968

ISBN 978-3-7908-2409-4

e-ISBN 978-3-7908-2410-0

DOI 10.1007/978-3-7908-2410-0

Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2010929257

© Springer-Verlag Berlin Heidelberg 2010

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, reuse of illustrations, recitation, broadcasting, reproduction on microfilm 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.

Cover design: Integra Software Services Pvt. Ltd, Puducherry

Printed on acid-free paper

Physica-Verlag is a brand of Springer-Verlag Berlin Heidelberg
Springer-Verlag is part of Springer Science+Business Media (www.springer.com)

*This volume is dedicated to Valeri Fedorov,
Ivan Vuchkov and Henry Wynn, Men of
Algorithms, on the occasion of their
birthdays (70, 70 and 65).*

Preface

This volume contains a substantial number of the papers presented at the mODa 9 conference in Bertinoro, Forlì, Italy, in June 2010; mODa stands for Model Oriented Data Analysis and Optimal Design. Design of experiments (DOE) is that part of statistics which provides tools for gathering data from experimentation in order to be able to draw conclusions in an efficient way. This subject began in an agricultural context, but nowadays is applied in many areas, both in science and industry, and a principal field of application is pharmacological research. Due to increasing competition, DOE has become crucial in drug development and clinical trials. Currently an important field of application is genomic, with the need to design and analyse microarray experiments. This increased competition requires ever increasing efficiency in experimentation, thus necessitating new statistical designs.

The theory for the design of experiments has accordingly developed a variety of approaches. A model-oriented view, where some knowledge of the form of the data-generating process is assumed, naturally leads to the so-called optimum design of experiments. Standard methods of DOE are no longer adequate in drug testing and biomedical statistics and research into new ways of planning clinical and non-clinical trials for dose-finding is receiving keen attention. Furthermore, in recent years the use of experimentation in engineering design has found renewed impetus through the practice of computer experiments, which has been steadily growing over the last two decades. These experiments are run on a computer code implementing a simulation model of a physical system of interest. This enables one to explore complex relationships between input and output variables. The main advantage should be that the system becomes more “observable”, since computer runs might be expected to be easier and cheaper than measurements taken in a physical set-up. However, with very complicated models, only a relatively few simulation runs are possible and good interpolators have to be found. The need to find optimal or sub-optimal ways of integrating simulated experiments and physical ones is paramount.

Leading experts on DOE have come together in the mODa group to promote new research topics, joint studies and financial support for research in DOE and related areas. In order to stimulate the necessary exchange of ideas, the mODa group

organises workshops. Previous conferences have been held on the Wartburg, then in the German Democratic Republic (1987), St Kirik Monastery, Bulgaria (1990), Petrodvorets, St Petersburg, Russia (1992), the Island of Spetses, Greece (1995), the Centre International des Rencontres Mathématiques, Marseille, France (1998), Puchberg / Schneeberg, Austria (2001), Kappellerput, Heeze, Holland (2004), and Almagro, Spain, (2007). The purpose of these workshops has traditionally been to bring together two pairs of groups: firstly scientists from the East and West of Europe with an interest in optimal design of experiments and related topics; and secondly younger and senior researchers. Thus an implicit aim of the mODa meetings has always been to give young researchers in DOE the opportunity to establish personal contacts with leading scholars in the field. These traditions remain vital to the health of the series. In recent years Europe has seen increasing unity and the scope of mODa has expanded to countries beyond Europe, including the USA, South Africa and India. Presentation of the work done by young researchers is very much encouraged in these workshops, either orally or by poster. The poster sessions have been developed according to a new format of one-minute introductory presentations by all, which ensures attention by the entire audience.

The 2010 edition of the conference is organized by the University of Bologna. Bologna University began to take shape at the end of the eleventh century and is probably the oldest university in the western world. Its history is one of great thinkers in science and the humanities, making this university an indispensable reference point in the panorama of European culture. Unfortunately, the workshop happens to take place in the middle of a world-wide economic crisis that has affected research opportunities in many countries, especially Italy, so that we are particularly grateful to our sponsors for making it possible, with their support, nevertheless to hold the workshop. GlaxoSmithKline have very kindly continued their support of the series of conferences. New sources have been: JMP, UK, who have generously funded the publication of these proceedings; the University of Bologna; the Department of Statistics at Bologna University; and CEUB itself, namely the Centre where the conference is hosted. We are very grateful for these contributions.

The major optimal design topics featuring in these proceedings include models with covariance structures, generalized linear models, sequential designs, applications in clinical trials, computer/screening experiments and designs for model discrimination; also new models appear, and classical design topics feature too. A breakdown is as follows:

1. The most common theme is that of covariance structures with the papers by Ginsbourger and Le Riche, by Pázman and W. Müller, by Pepelyshev, by Biswas and Mandal, by Rodríguez-Díaz, Santos-Martín, Stehlík and Waldl, and by Vazquez and Bect.
2. Non-linear models feature in the contributions of C. Müller and Schäfer, of Manukyan and Rosenberger and of Torsney. Optimal designs for linear logistic test models are investigated by Graßhoff, Holling and Schwabe
3. The topic of clinical trials arises both in the papers by Anisimov and by Fedorov, Leonov and Vasiliev, and in the form of dose finding studies in Roth and in Fedorov, Wu and Zhang.

4. Screening experiments appear in the papers by Jones and Nachtsheim, and by Peterson, whereas the paper by Roustant, Franco, Carraro, and Jourdan deals with computer experiments.
5. The topic of both the papers by Atkinson and by Tommasi, Santos-Martín and Rodríguez-Díaz is discrimination between models.
6. Sequential design has been investigated by several authors: by Yao and Flournoy, by Maruri-Aguilar and Trandafir, by Baldi Antognini and Zagoraiou, by Flournoy, May, Moler and Plo, and by Pronzato.
7. The papers by Bischoff and by Mielke and Schwabe deal with optimality criteria for experimental design; Bonnini, Corain, and Salmaso's paper is about sample size determination. Coetzer and Haines write about optimal design for compositional data.
8. Finally, topics covered by just one paper are microarray experiments and split-plot and robust designs. The authors thereof are Schiffel and Hilgers on the one hand, and Berni on the other.

Bologna,
March 2010

Alessandra Giovagnoli
Anthony C. Atkinson
Ben Torsney
with the collaboration of *Caterina May*

Acknowledgements

We thank JMP, a division of SAS Institute, UK



who have generously funded the publication of these Proceedings.
<http://www.jmp.com/uk/>

We also thank the following organizations for their financial support:

- GlaxoSmithKline GSK
<http://www.gsk.com/>
- Department of Statistical Sciences, University of Bologna
<http://www.stat.unibo.it/ScienzeStatistiche/default.htm>
- Alma Mater Studiorum, University of Bologna
<http://www.eng.unibo.it/PortaleEn/default.htm>
- CEUB - University Residential Centre of Bertinoro
http://www.ceub.it/index_en.cfm

Contents

Impact of Stratified Randomization in Clinical Trials	1
Vladimir V. Anisimov	
1 Introduction	1
2 Recruitment in Different Strata	2
3 Randomization Effects	3
3.1 Impact of Randomization on the Power and Sample Size ..	4
4 Conclusions	8
References	8
The Non-Uniqueness of Some Designs for Discriminating Between Two Polynomial Models in One Variable	9
Anthony C. Atkinson	
1 Introduction	9
2 Background	10
3 Examples	10
4 Designs for Higher-Order Models	13
5 Design and the Analysis of Data	15
References	16
Covariate Adjusted Designs for Combining Efficiency, Ethics and Randomness in Normal Response Trials	17
Alessandro Baldi Antognini and Maroussa Zagoraiou	
1 Introduction	17
2 Optimal Allocations for Inference	18
3 Optimal Allocations for Ethics	20
4 Compound Optimal Designs	20
5 Doubly-Adaptive Biased Coin Designs with Covariates	22
References	24

Split-Plot and Robust Designs: Weighting and Optimization in the Multiple Response Case 25
 Rossella Berni
 1 Introduction 25
 2 Split-Plot Theory 26
 3 The Optimization Procedure 27
 3.1 Desirability Function and Weighting 28
 4 An Application to a Case Study 28
 5 Optimization Results 30
 6 Concluding Remarks 32
 References 32

An Improvement in the Lack-of-Fit Optimality of the (Absolutely) Continuous Uniform Design in Respect of Exact Designs 33
 Wolfgang Bischoff
 1 Introduction 33
 2 Preliminaries 34
 3 LOF Optimality 35
 4 Appendix 38
 References 39

Optimal Allocation Proportion for a Two-Treatment Clinical Trial Having Correlated Binomial Responses 41
 Atanu Biswas and Saumen Mandal
 1 Introduction 41
 2 Optimal Allocation Proportions in the Presence of Correlation 43
 3 Numerical Computations 45
 4 Concluding Remarks 47
 References 48

Sample Size Determination for Multivariate Performance Analysis with Complex Designs 49
 Stefano Bonnini, Livio Corain and Luigi Salmaso
 1 Introduction 49
 2 Global Ranking Methods 50
 3 Simulation Study and Sample Size Determination 52
 4 Conclusions 54
 References 55

Optimal Design for Compositional Data 57
 Roelof L. J. Coetzer and Linda M. Haines
 1 Introduction 57
 2 Additive Logistic Normal 58
 2.1 Model 58
 2.2 Design 59
 3 Dirichlet Model 59

3.1	Model	59
3.2	Design	60
4	Example	61
4.1	Additive Logistic Normal	62
4.2	Dirichlet	63
5	Conclusions	63
	References	64
Dose Finding Experiments: Responses of Mixed Type		65
Valerii V. Fedorov, Yuehui Wu and Rongmei Zhang		
1	Introduction	65
2	Model	66
2.1	Information Matrix for a Single Observation	66
2.2	Utility and Penalty Functions	67
3	Optimal Designs	68
3.1	Adaptive Designs	69
4	Examples	70
5	Conclusions	71
	References	72
Pharmacokinetic Studies Described by Stochastic Differential Equations: Optimal Design for Systems with Positive Trajectories		73
Valerii V. Fedorov, Sergei L. Leonov and Vyacheslav A. Vasiliev		
1	Introduction	73
2	Response Models	74
2.1	Stochastic Systems with Positive Trajectories	75
3	Optimal Designs	76
3.1	Sampling Times and Examples of Optimal Design	76
4	Discussion	78
	Appendix: Proof of Lemma 1	79
	References	80
On Testing Hypotheses in Response-Adaptive Designs Targeting the Best Treatment		81
Nancy Flournoy, Caterina May, Jose A. Moler and Fernando Plo		
1	Introduction	81
2	RRU-Designs and Test of Hypotheses	82
2.1	Test Based on the t-Statistic	83
2.2	Test Based on the 'Proportion of Black Balls' Statistic ...	84
3	Numerical Results	85
4	Discussion and Further Developments	87
	References	88

Towards Gaussian Process-based Optimization with Finite Time Horizon 89
David Ginsbourger and Rodolphe Le Riche

- 1 Introduction 89
- 2 What is a Strategy and How to Measure its Performance? 91
- 3 Towards Deriving the Optimal Finite Time Strategy 92
- References 96

Optimal Designs for Linear Logistic Test Models 97
Ulrike Graßhoff, Heinz Holling and Rainer Schwabe

- 1 Introduction 97
- 2 Optimal Design 100
- 3 Discussion 102
- References 103

A Class of Screening Designs Robust to Active Second-Order Effects 105
Bradley Jones and Christopher J. Nachtsheim

- 1 Introduction 105
- 2 Design Structure: An Example 106
- 3 Algorithm 107
- 4 Design Diagnostic Comparisons 108
- 5 Suggestions for Analysis 111
- 6 Summary 111
- References 112

D-Optimal Design for a Five-Parameter Logistic Model 113
Zorayr Manukyan and William F. Rosenberger

- 1 Introduction 113
- 2 Methods 114
- 3 Results 116
- 4 Discussion 117
- 5 Appendix: Information Matrix 118
- References 119

Sequential Barycentric Interpolation 121
Hugo Maruri-Aguilar and Paula Camelia Trandafir

- 1 Introduction 121
- 2 Barycentric Lagrange Interpolation 122
- 3 Sequential Interpolation 122
 - 3.1 Response-based Update 123
 - 3.2 Sequential Design Algorithm 123
- 4 Performance and Large Sample Properties 124
 - 4.1 Interpolating Performance 125
 - 4.2 Large Sample Properties 125
- 5 Discussion and Future Work 127
- Appendix A: Proof of Theorem 1 127
- References 128

Some Considerations on the Fisher Information in Nonlinear Mixed Effects Models 129
 Tobias Mielke and Rainer Schwabe

- 1 Introduction 129
- 2 Non-linear Models 130
- 3 Mixed-Effects Models 132
- 4 Example 134
- 5 Discussion 136
- References 136

Designs with High Breakdown Point in Nonlinear Models 137
 Christine H. Müller and Christina Schäfer

- 1 Introduction 137
- 2 Identifiability and d Fullness 140
- 3 Nonlinear Models with Unrestricted Parameter Space 141
- 4 Nonlinear Models with Restricted Parameter Space 143
- 5 Discussion 143
- References 144

A Note on the Relationship between Two Approaches to Optimal Design under Correlation 145
 Andrej Pázman and Werner G. Müller

- 1 Introduction 145
- 2 Information Matrices 146
- 3 Conclusions 147
- References 148

The Role of the Nugget Term in the Gaussian Process Method 149
 Andrey Pepelyshev

- 1 Introduction 149
- 2 The Likelihood for a Gaussian Process Without the Nugget Term 150
- 3 The Likelihood for a Gaussian Process With a Nugget Term 153
 - 3.1 MLE for a Gaussian Process 153
 - 3.2 MLE for Stationary Processes 154
- 4 Conclusions 155
- References 156

A Bonferroni-Adjusted Trend Testing Method for Excess over Highest Single Agent 157
 John J. Peterson

- 1 Testing for Excess Over Highest Single Agent (EOHSA) 157
 - 1.1 Model and Testing for EOHSA 157
 - 1.2 Approaches Based Upon Trend Tests 159
 - 1.3 Multiplicity Adjusted p -Values 160
- 2 An Example 162
- References 164

Asymptotic Properties of Adaptive Penalized Optimal Designs over a Finite Space 165

Luc Pronzato

- 1 Introduction 165
- 2 Asymptotic Properties of Estimators when \mathcal{X} is Finite 166
- 3 Adaptive Penalized D-optimal Design 167
 - 3.1 A bound on the sampling rate of nonsingular designs 168
 - 3.2 λ_n is bounded in (6) 169
 - 3.3 λ_n tends to infinity in (6) 170
- References 171

Filling and D-optimal Designs for the Correlated Generalized Exponential Model 173

Juan M. Rodríguez-Díaz, Teresa Santos-Martín, Milan Stehlík and Helmut Waldl

- 1 Introduction 173
- 2 Assuming m Known 175
- 3 Case of Unknown m 176
- 4 Parabolic Designs 176
- 5 Illustrative Example 178
- 6 Conclusions and Discussion 179
- References 179

Designs for Dose Finding Studies on Safety and Efficacy 181

Katrin Roth

- 1 Introduction 181
- 2 Optimal Design in a Bivariate Model 182
 - 2.1 Definition of the Model 182
 - 2.2 Optimal Designs for This Model 183
- 3 Sequential Approach and Simulation Study 186
- 4 Discussion 188
- References 188

A Radial Scanning Statistic for Selecting Space-filling Designs in Computer Experiments 189

Olivier Roustant, Jessica Franco, Laurent Carraro and Astrid Jourdan

- 1 Introduction 189
- 2 The Radial Scanning Statistic 190
 - 2.1 Selecting a Goodness-of-fit Test for the Uniform Distribution 192
 - 2.2 Graphical Properties 193
 - 2.3 Decisional Issues 193
- 3 Usage and Applications 194
- 4 Conclusion and Further Research 195
- References 196

Optimal Designs for Two-Colour Microarray Experiments for Estimating Interactions 197
 Katharina Schiffl and Ralf-Dieter Hilgers

- 1 Introduction 197
- 2 Preliminaries 198
- 3 Optimal Designs 200
 - 3.1 Interactions in Multi-factor Settings for the Estimation of All Pairwise Comparisons 200
 - 3.2 Interactions in Two-factor Settings 201
- 4 Discussion 202
- Appendix: Proof of Theorem 2 202
- References 204

Discrimination Between Random and Fixed Effect Logistic Regression Models 205
 Chiara Tommasi, Maria Teresa Santos-Martín and Juan Manuel Rodríguez-Díaz

- 1 Introduction 205
- 2 Logistic Regression Model 206
- 3 D_S -Optimality Criterion 207
- 4 KL-Optimality Criterion 209
- 5 Some Results 209
- 6 Appendix 210
- References 212

Estimation and Optimal Designing under Latent Variable Models for Paired Comparisons Studies via a Multiplicative Algorithm 213
 Bernard Torsney

- 1 Paired Comparisons 213
 - 1.1 Introduction 213
 - 1.2 The Data 214
 - 1.3 Models 214
- 2 Parameter Estimation 215
- 3 Optimality Conditions 216
- 4 Algorithms 217
 - 4.1 Multiplicative Algorithm 217
 - 4.2 Properties of the Algorithm 217
- 5 Fitting Bradley-Terry Models 217
- 6 Local Optimal Designing 218
- 7 Discussion 220
- References 220

Pointwise Consistency of the Kriging Predictor with Known Mean and Covariance Functions 221
 Emmanuel Vazquez and Julien Bect

- 1 Introduction 221
- 2 Several Formulations of Pointwise Consistency 222

- 3 Pointwise Consistency in L^2 -Norm and the No-Empty-Ball Property 224
- 4 Pointwise Consistency for Continuous Sample Paths 225
- 5 Proof of Proposition 1 227
- References 228

Information in a Two-stage Adaptive Optimal Design for Normal Random Variables having a One Parameter Exponential Mean Function

..... 229

Ping Yao and Nancy Flournoy

- 1 Introduction 229
- 2 A Two-stage Design for Normal Random Variables having a One Parameter Exponential Mean Function 230
- 3 Properties of the Stage 2 Design Point 231
- 4 Information Measures 232
 - 4.1 A Simulated Illustration 234
- 5 Discussion 236
- References 236

Index 237

List of Contributors

Vladimir V. Anisimov

Research Statistics Unit, GlaxoSmithKline, New Frontiers Science Park
(South), Third Avenue, Harlow, Essex, CM19 5AW, United Kingdom, e-mail:
Vladimir.V.Anisimov@gsk.com

Anthony C. Atkinson

London School of Economics, London WC2A 2AE, UK, e-mail:
a.c.atkinson@lse.ac.uk

Alessandro Baldi Antognini

Department of Statistical Sciences, University of Bologna, Via Belle Arti 41, 40126
Bologna, Italy, e-mail: a.baldi@unibo.it

Julien Bect

SUPELEC, 3 rue Joliot-Curie, 91192 Gif-sur-Yvette, France, e-mail:
julien.bect@supelec.fr

Rossella Berni

Department of Statistics “G.Parenti”, Viale Morgagni 59, 50134 Florence, Italy,
e-mail: berni@ds.unifi.it

Wolfgang Bischoff

Catholic University Eichstätt-Ingolstadt, D-85071 Eichstätt, Germany, e-mail:
wolfgang.bischoff@ku-eichstaett.de

Atanu Biswas

Applied Statistics Unit, Indian Statistical Institute, 203 B. T. Road, Kolkata – 700
108, India, e-mail: atanu@isical.ac.in

Stefano Bonnini

Department of Mathematics, University of Ferrara, Via Savonarola, 9 - 44121
Ferrara, Italy, e-mail: stefano.bonnini@unife.it

Laurent Carraro

Telecom Saint-Etienne, France, e-mail: laurent.carraro@telecom-st-etienne.fr

Roelof L. J. Coetzer

Sasol Technology, PO Box 1, Sasolburg 1947, South Africa, e-mail:
roelof.coetzer@sasol.co.za

Livio Corain

Department of Management and Engineering, University of Padova, Italy, e-mail:
livio.corain@unipd.it

Valerii V. Fedorov

GlaxoSmithKline, Collegeville, PA, USA, e-mail: Valeri.V.Fedorov@gsk.com

Nancy Flournoy

Department of Statistics, 146 Middlebush Hall, University of Missouri, Columbia,
MO 65203, USA, e-mail: flournoyn@missouri.edu

Jessica Franco

Total, Pau, France, e-mail: jessica.franco@total.com

David Ginsbourger

CHYN, Rue Emile-Argand 11, CP 158 CH-2009 Neuchâtel, Switzerland, e-mail:
david.ginsbourger@unine.ch

Ulrike Graßhoff

Institut für Mathematische Stochastik, Otto-von-Guericke-Universität Magdeburg,
PF 4120, 39016 Magdeburg, Germany, e-mail: ulrike.grasshoff@mathematik.uni-
magdeburg.de

Linda M. Haines

Department of Statistical Sciences, University of Cape Town, Rondebosch 7700,
South Africa, e-mail: linda.haines@uct.ac.za

Ralf-Dieter Hilgers

RWTH Aachen University, Department of Medical Statistics, 52074 Aachen,
Germany, e-mail: rhilgers@ukaachen.de

Heinz Holling

Psychologisches Institut IV, Westfälische Wilhelms-Universität Münster, Fliedner-
str. 21, 48149 Münster, Germany, e-mail: holling@uni-muenster.de

Bradley Jones

SAS Institute, Cary, NC 27513, USA, e-mail: bradley.jones@sas.com

Astrid Jourdan

Ecole Internationale De Sciences Et Traitement Informatique, Pau, France, e-mail:
astrid.jourdan@eisti.fr

Rodolphe Le Riche

CNRS and CROCUS, Ecole Nationale Supérieure des Mines, 158 cours Fauriel,
Saint-Etienne, France, e-mail: leriche@emse.fr

Sergei L. Leonov

GlaxoSmithKline, 1250 So. Collegeville Rd, PO Box 5089, Collegeville, PA, 19426, USA, e-mail: Sergei.2.Leonov@gsk.com

Saumen Mandal

Department of Statistics, University of Manitoba, Winnipeg, MB, R3T 2N2, Canada, e-mail: saumen_mandal@umanitoba.ca

Zorayr Manukyan

The EMMES Corporation, 401 N. Washington St., Rockville, MD 20850, USA, e-mail: zmanukyan@emmes.com

Hugo Maruri-Aguilar

School of Mathematical Sciences, Queen Mary, University of London, Mile End Road, London E1 4NS, UK, e-mail: H.Maruri-Aguilar@qmul.ac.uk

Caterina May

Università degli Studi del Piemonte Orientale, Via Perrone 18, 28100 Novara, Italy, e-mail: caterina.may@eco.unipmn.it

Tobias Mielke

Otto-von-Guericke University, Magdeburg, Germany, e-mail: tobias.mielke@ovgu.de

Jose A. Moler

Universidad Publica de Navarra, Campus de Arrosadia s/n. 31006-Pamplona, Spain, e-mail: jmoler@unavarra.es

Christine H. Müller

Fakultät Statistik, Technische Universität Dortmund, D-44221 Dortmund, Germany, e-mail: cmueller@statistik.tu-dortmund.de

Werner G. Müller

Department of Applied Statistics, Johannes Kepler University, Freistädterstr. 315, 4040 Linz, Austria, e-mail: werner.muller@ku.at

Christopher J. Nachtsheim

Carlson School of Management, University of Minnesota, Minneapolis, MN 55455, USA, e-mail: nacht001@umn.edu

Andrej Pázman

Comenius University, Bratislava, Slovakia, e-mail: pazman@fmph.uniba.sk

Andrey Pepelyshev

University of Sheffield, Sheffield S3 7RH, UK, e-mail: a.pepelyshev@sheffield.ac.uk

John J. Peterson

GlaxoSmithKline Pharmaceuticals, 1250 So. Collegeville Rd, PO Box 5089, Collegeville, PA, 19426, USA, e-mail: john.peterson@gsk.com

Fernando Plo

Universidad de Zaragoza, Pedro Cerbuna, 12 50009-Zaragoza, Spain, e-mail: fplo@unizar.es

Luc Pronzato

Laboratoire I3S, CNRS/Université de Nice–Sophia Antipolis, France, e-mail: pronzato@i3s.unice.fr

Juan M. Rodríguez-Díaz

Department of Statistics, University of Salamanca, Pl. de los Caídos, Salamanca, Spain, e-mail: juanmrod@usal.es

William F. Rosenberger

Department of Statistics, George Mason University, 4400 University Dr., MS 4A7, Fairfax, VA 22030-4444, USA, e-mail: wrosenbe@gmu.edu

Katrin Roth

Bayer Schering Pharma AG, Müllerstr. 178, 13353 Berlin, and Otto-von-Guericke-Universität, Magdeburg, Germany, e-mail: katrin.roth@bayerhealthcare.com

Oliver Roustant

Ecole Nationale Supérieure des Mines, 158 cours Fauriel, Saint-Etienne, France, e-mail: roustant@emse.fr

Luigi Salmaso

Department of Management and Engineering, University of Padova, 35121 Padua, Italy, e-mail: salmaso@gest.unipd.it

Teresa Santos-Martín

Department of Statistics, University of Salamanca, Pl. de los Caídos, Salamanca, Spain, e-mail: maysam@usal.es

Christina Schäfer

Institut für Mathematik, Carl von Ossietzky Universität Oldenburg, D-26111 Oldenburg i.O., Germany, e-mail: christina-schaefer@hotmail.de

Katharina Schiff

RWTH Aachen University, Department of Medical Statistics, 52074 Aachen, Germany, e-mail: kschiff@ukaachen.de

Rainer Schwabe

Institut für Mathematische Stochastik, Otto-von-Guericke-Universität Magdeburg, PF 4120, 39016 Magdeburg, Germany, e-mail: Rainer.Schwabe@mathematik.uni-magdeburg.de

Milan Stehlík

Department of Applied Statistics, Johannes Kepler University, Freistädterstr. 315, 4040 Linz, Austria, e-mail: Milan.Stehlik@jku.at

Paula Camelia Trandafir

Departamento de Estadística e Investigación Operativa, Universidad de Valladolid, Prado de la Magdalena s/n, 47005, Valladolid, Spain, e-mail: camelia@eio.uva.es

Chiara Tommasi

Department of Economics, Business and Statistics, University of Milan, Via Conservatorio, 7 - 20122 Milano, Italy, e-mail: chiara.tommasi@unimi.it

Ben Torsney

Department of Statistics, University of Glasgow, 15 University Gardens, Glasgow G12 8QW, UK, e-mail: bent@stats.gla.ac.uk

Vyacheslav A. Vasiliev

Tomsk State University, Department of Applied Mathematics and Cybernetics, Tomsk, Russia, e-mail: vas@mail.tsu.ru

Emmanuel Vazquez

SUPELEC, 3 rue Joliot-Curie, 91192 Gif-sur-Yvette, France e-mail: emmanuel.vazquez@supelec.fr

Helmut Waldl

Department of Applied Statistics, Johannes Kepler University, Freistädterstr. 315, 4040 Linz, Austria, e-mail: Helmut.Waldl@jku.at

Yuehui Wu

GlaxoSmithKline, Collegeville, PA, USA, e-mail: Yuehui.2.Wu@gsk.com

Ping Yao

Public Health and Health Education Program, 209K Wirtz Hall, Northern Illinois University, Dekalb, IL 60115, USA, e-mail: pyao@niu.edu

Maroussa Zagoraiou

Department of Statistical Sciences, University of Bologna, Via Belle Arti 41, 40126 Bologna, Italy, e-mail: maroussa.zagoraiou@unibo.it

Rongmei Zhang

CCEB, University of Pennsylvania, Philadelphia, PA, USA, e-mail: rongmei@mail.med.upenn.edu

Impact of Stratified Randomization in Clinical Trials

Vladimir V. Anisimov

Abstract This paper deals with the analysis of randomization effects in clinical trials. The two randomization schemes most often used are considered: unstratified and stratified block-permuted randomization. A new analytic approach using a Poisson-gamma patient recruitment model and its further extensions is proposed. The prediction of the number of patients randomized in different strata to different treatment arms is considered. In the case of two treatments, the properties of the total imbalance in the number of patients on treatment arms caused by using stratified randomization are investigated and for a large number of strata a normal approximation of imbalance is proved. The impact of imbalance on the power of the trial is considered. It is shown that the loss of statistical power is practically negligible and can be compensated by a minor increase in sample size. The influence of patient dropout is also investigated.

1 Introduction

The properties of various types of randomization schemes are studied in the papers Hallstrom and Davis (1988), Lachin (1988), Matts and Lachin (1988), and books by Pocock (1983), Rosenberger and Lachin (2002). However, the impact of randomness in patient recruitment and the prediction of the number of randomized patients in the case of multiple centres have not been fully investigated.

To investigate these phenomena, a new analytic approach using a Poisson-gamma patient recruitment model developed in Anisimov and Fedorov (2006, 2007) is proposed. The model accounts for the variation in recruitment over time and in recruitment rates between strata. The prediction of the number of patients randomized in different strata to different treatment arms is considered. In the case of two treat-

Prof. Vladimir V. Anisimov

Research Statistics Unit, GlaxoSmithKline, New Frontiers Science Park (South), Third Avenue, Harlow, Essex, CM19 5AW, United Kingdom, e-mail: Vladimir.V.Anisimov@gsk.com

ments, the properties of the total imbalance in the number of patients randomized to different treatment arms caused by using stratified randomization are investigated as well. For a large number of strata a normal approximation of imbalance is proved. These results are used for investigating the impact of randomization on the power and sample size of the trial. Note that in a special case of a centre-stratified randomization some results in these directions are obtained in Anisimov (2007). The effect of patient dropout is also considered. These results form the basis for comparing randomization schemes using combined criteria including statistical power, study costs, drug supply costs, etc.

2 Recruitment in Different Strata

Consider a multicentre clinical trial carried out with the aim to recruit in total n patients. Suppose that the patient population is divided into S strata. Strata can stand for different countries, centres or regions, groups of population specified by some covariates, etc. Upon registration, patients are randomized to one of the treatment arms according to some randomization scheme. The recruitment is stopped when the total number of recruited patients reaches n . Assume that the patients in different strata are recruited independently. Accounting for a natural variation in recruitment between strata, we can consider the following model: the recruitment in s -th stratum is described by a Poisson process with rate μ_s , where μ_s is viewed as a realization of a gamma distributed variable with parameters $(\alpha N_s, \beta)$ (shape and rate parameters), and the values N_s reflect the sizes of strata. Denote $N = \sum_s N_s$.

As a natural illustration of this model, assume that there are N clinical centres divided among S regions, where a region s has N_s centres. Let us associate the region s with s -th stratum. Suppose that the recruitment in centres is described by a Poisson-gamma model (Anisimov and Fedorov, 2006,2007): in centre i the patients are recruited according to a Poisson process with rate λ_i , where $\{\lambda_i\}$ are viewed as a sample from a gamma distributed population with parameters (α, β) . Then the recruitment in s -th region is described by a Poisson process with rate μ_s which is gamma distributed with parameters $(\alpha N_s, \beta)$. For this case, in Anisimov and Fedorov (2007) a ML-procedure for estimating parameters is proposed.

Consider now the prediction of the total number of patients n_s recruited in a particular strata s . The variable n_s has a mixed binomial distribution with parameters (n, g_s) where $g_s = \mu_s / \mu$, $\mu = \sum_{s=1}^S \mu_s$. Thus, μ has a gamma distribution with parameters $(\alpha N, \beta)$ and g_s has a beta distribution with parameters $(\alpha N_s, \alpha(N - N_s))$. Denote by $\mathcal{B}(a, b)$ a beta function. Then n_s has a beta-binomial distribution and $\mathbf{P}(n_s = k) = P(n, N, N_s, \alpha, k)$, where

$$P(n, N, N_s, \alpha, k) = \binom{n}{k} \frac{\mathcal{B}(\alpha N_s + k, \alpha(N - N_s) + n - k)}{\mathcal{B}(\alpha N_s, \alpha(N - N_s))}, \quad k = 0, \dots, n. \quad (1)$$

3 Randomization Effects

Description of randomization schemes can be found in the books by Pocock (1983), Rosenberger and Lachin (2002). Consider the two often used in clinical trials randomization schemes: unstratified and stratified block-permuted randomization. Unstratified randomization means that the patients registered for the study are randomized to treatment arms according to the independent randomly permuted blocks of a fixed size without regard to stratum. Stratified randomization means that the patients are randomized according to randomly permuted blocks separately in each stratum. Clearly, unstratified randomization minimizes the imbalance in the number of patients on treatment arms for the whole study, but in general is likely to increase the imbalance within each stratum compared to stratified randomization.

Assume that there are K treatments with the allocations (k_1, \dots, k_K) within a randomly permuted block of a size $K_1 = \sum_{j=1}^K k_j$. Denote by $n_s(j)$ the number of patients randomized to treatment j in s -th stratum.

Consider first an unstratified randomization. Assume that the value $M = n/K_1$ is integer. Then there are Mk_j patients on treatment j and all patients can be divided into K groups with Mk_j patients in group j , $j = 1, \dots, K$. Within each group the patients are distributed among strata independently of other groups according to a beta-binomial distribution as described in section 2. Thus, for any stratum s ,

$$\mathbf{P}(n_s(j) = i_j, j = 1, \dots, K) = \prod_{j=1}^K P(Mk_j, N, N_s, \alpha, i_j). \quad (2)$$

Consider now a stratified randomization. In this case in each stratum randomization is carried out independently of other strata according to block-permuted randomization. If in some stratum s , n_s is not a multiple of K_1 , then the last block is incomplete. The incomplete block may contain an unequal number of patients on treatment arms and cause an imbalance in this stratum. Many incomplete blocks in different strata may cause an imbalance between the total number of patients on treatment arms and this may lead to power loss in the study.

Assume that s -th stratum contains an incomplete block of size m , $m = 1, \dots, K_1 - 1$, and denote by $\xi_j(m)$ the number of instances of treatment j in this block. Then $\xi_j(m)$ has a hypergeometric distribution and $\mathbf{P}(\xi_j(m) = l) = \binom{k_j}{l} \binom{K_1 - k_j}{m - l} \binom{K_1}{m}^{-1}$, $l = 0, 1, \dots, \min(k_j, m)$. Therefore, $\mathbf{E}[\xi_j(m)] = k_j m / K_1$, $\mathbf{Var}[\xi_j(m)] = k_j m (K_1 - k_j) \times (K_1 - m) / (K_1^2 (K_1 - 1))$. Let $\text{int}(a)$ be the integer part of a , and $\text{mod}(a, k) = a - \text{int}(a/k)k$. Then

$$n_s(j) = \text{int}(n_s / K_1) k_j + \xi_j(\text{mod}(n_s, K_1)). \quad (3)$$

As the distribution of n_s is given by (1), the characteristics of $n_s(j)$ can be numerically calculated. Closed-form expressions for the mean and the variance of $n_s(j)$ also can be derived. In the case when strata are associated with different geographical regions, these results allow prediction of supply needed to cover patient demand in regions, number of places in hospitals, etc.

3.1 Impact of Randomization on the Power and Sample Size

Let us consider the impact of randomization scheme on the sample size and the power of a statistical test. If one might expect a statistically significant stratum-by-treatment interaction, then stratified randomization should be preferable from a statistical point of view as it provides better balance within each stratum. Therefore, let us assume that there is no stratum-by-treatment interaction. As stratified randomization in general causes the random imbalance between treatment arms, one would expect that unstratified randomization should be preferable. However, we prove that in general the size of imbalance is rather small compared to the total sample size and its impact on the power and sample size is practically negligible.

3.1.1 Properties of Imbalance in Stratified Randomization

Assume for simplicity that there are only two treatments, a and b with equal treatment allocations. Denote by $\eta_s = n_s(a) - n_s(b)$ an imbalance in stratum s . Let n_j^* be the total number of patients on treatment j , $j = a, b$, and $\Delta = n_a^* - n_b^*$ be the total imbalance in the number of patients on both treatments. Then $\Delta = \sum_{s=1}^S \eta_s$.

Theorem 1. *For large enough n and S such that $n \min(N_s)/N \geq K_1$, the imbalance Δ is well approximated by a normal distribution with mean zero and variance $s_0^2 S$, where $s_0^2 = (K_1 + 1)/6$.*

Proof. For equal treatment proportions $k_j = K_1/2$ and $\mathbf{E}[\xi_j(m)] = m/2$, $\mathbf{Var}[\xi_j(m)] = m(K_1 - m)/(4(K_1 - 1))$, $j = 1, 2$. Thus, if in s -th stratum the incomplete block has a size m , then the imbalance in this stratum is $\eta_s(m) = \xi_1(m) - (m - \xi_1(m)) = 2\xi_1(m) - m$, and $\mathbf{E}[\eta_s(m)] = 0$, $\mathbf{Var}[\eta_s(m)] = 4\mathbf{Var}[\xi_1(m)] = m(K_1 - m)/(K_1 - 1)$. In general, in stratum s the imbalance η_s is a random variable: $\eta_s = \eta_s(m)$ with probability $q_m(n, N_s, K_1)$, $m = 0, \dots, K_1 - 1$, where $\eta_s(0) = 0$, and $q_m(n, N_s, K_1) = \mathbf{P}(\text{mod}(n_s, K_1) = m)$. Thus, $\mathbf{E}[\eta_s] = 0$ and from (1) it follows

$$q_m(n, N_s, K_1) = \sum_{l=0}^{n/K_1-1} P(n, N, N_s, \alpha, m + lK_1), \quad m = 0, 1, \dots, K_1 - 1. \quad (4)$$

Furthermore, if on average the number of patients in a stratum is not less than $2K_1$, one can use the approximation $q_m(\cdot) \approx 1/K_1$ (compare with Hallstrom and Davis (1988)). This is also supported by numerical calculations and Monte Carlo simulations (Anisimov 2007). For example, for $n = 60, S = 6, N_s = 1$ (on average 10 patients in a stratum), $K_1 = 4$ and $\alpha = 1.2$, numerical calculations give $(q_0, q_1, q_2, q_3) = (0.269, 0.259, 0.244, 0.228)$ and simulated values for 10^6 runs coincide with these values up to 3 digits.

Thus, using the approximation $q_m(n, N_s, K_1) = 1/K_1, m = 0, \dots, K_1 - 1$, we have $\mathbf{Var}[\eta_s] \approx s_0^2 = (K_1 + 1)/6$. The variables η_s and η_p are not correlated as $s \neq p$ and conditionally independent. Thus, $\mathbf{E}[\eta_s \eta_p] = 0$, $\mathbf{Var}[\Delta] \approx s_0^2 S$, and at large S , Δ is

approximated by a normal distribution with parameters $(0, s_0^2 S)$. This is supported by Monte Carlo simulations (Anisimov 2007). \square

Remark 1. As shown above, for large enough numbers of patients the imbalance η_s in each stratum can be approximated by a mixed hypergeometric distribution $\tilde{\eta}_s = 2\xi(U) - U$, where $\mathbf{P}(U = m) = 1/K_1, m = 0, \dots, K_1 - 1$, $\mathbf{E}\tilde{\eta}_s = 0$, $\mathbf{Var}\tilde{\eta}_s = s_0^2$, and the variables $\tilde{\eta}_s$ are independent. Thus, for a few strata ($S < 10$), the imbalance Δ can be approximated by the variable $\tilde{\Delta} = \sum_{s=1}^S \tilde{\eta}_s$, where $\mathbf{E}\tilde{\Delta} = 0$, $\mathbf{Var}\tilde{\Delta} = s_0^2 S$.

3.1.2 Impact of Imbalance on the Power and Sample Size

In general imbalance is rather small compared to the sample size. Theorem 1 implies that with probability $1 - \varepsilon$, for large S ($S \geq 10$), $|\Delta| \leq s_0 \sqrt{S} z_{1-\varepsilon/2}$. If $S < 10$, then $|\Delta| \leq s_0 \sqrt{S/\varepsilon}$ (basing on Remark 1 and Chebyshev inequality). In particular, for $n \geq 100$, $K_1 \leq 4$ with probability 0.95, $|\Delta| \leq 8$ as $S = 20$, and $|\Delta| \leq 6$ as $S = 6$.

Let us evaluate the increase in sample size required to maintain the same power as for the balanced study accounting for possible imbalance. Consider as an example a standard test that compares means in two patient populations.

Assume that n patients are randomized to two treatments, a and b , in S strata. If one can expect a stratum-by-treatment interaction, then the stratified randomization should be more preferable from a statistical point of view. Consider the case where there is no stratum-by-treatment interaction. Then general guidelines indicate that unstratified randomization should be more preferable from a statistical point of view. However, we prove that stratified randomization leads practically to the same results.

Consider a stratified randomization by blocks of size K_1 and equal treatment allocations. Let n_j^* be the total number of patients randomized to treatment j , $j = a, b$, and $\{x_1, x_2, \dots, x_{n_a^*}\}$ and $\{y_1, y_2, \dots, y_{n_b^*}\}$ be the patient responses on each treatment. Suppose that the observations are independent with unknown means m_a and m_b and the known variance σ^2 . It is known that for testing the hypothesis: $H_0 : m_a - m_b = 0$ against $H_1 : m_a - m_b \geq h$ with probabilities γ and δ of type I and type II errors, the values n_a^* and n_b^* should satisfy the relation

$$h \left(\sigma \sqrt{1/n_a^* + 1/n_b^*} \right)^{-1} = z_{1-\gamma/2} + z_{1-\delta}. \quad (5)$$

For a balanced study $n_a^* = n_b^* = n/2$ (assuming that n is even). Thus, in the balanced case a sample size is $n_{bal} = 4\sigma^2(z_{1-\gamma/2} + z_{1-\delta})^2/h^2$. Denote by $\Delta = n_b^* - n_a^*$ the imbalance between treatment arms. Let us evaluate a sample size increase $n_+ = n - n_{bal}$ required to achieve the same power as for a balanced trial.

Theorem 2. At small S/n_{bal}^2 , $n_+ \approx s_0^2 S(1 + \sqrt{2}z_{1-\delta})(1 + \zeta)/n_{bal}$, where ζ is the error term of approximation, $\zeta = O(s_0^2 S/n_{bal}^2)$.

Proof. Consider a standard test statistic

$$T^* = \frac{\bar{x}_a - \bar{y}_b}{\sigma \sqrt{1/n_a^* + 1/n_b^*}}, \quad (6)$$

where \bar{x}_a and \bar{y}_b are sample means. Under the hypothesis H_0 for large enough n_a^* and n_b^* , $T^* \approx \mathcal{N}(0, 1)$, where $\mathcal{N}(0, 1)$ has a standard normal distribution. Thus, for testing H_0 with error probabilities γ and δ , the acceptance region is the interval $(-z_{1-\gamma/2}, z_{1-\gamma/2})$, and under the hypothesis H_1 it should be

$$\mathbf{P}_{H_1}(T^* \leq z_{1-\gamma/2}) = \delta. \quad (7)$$

Accounting for random imbalance, let us find n satisfying (7). Let ζ_i be the values of the magnitude $O(s_0^2 S/n_{bal}^2)$. Then, under the hypothesis H_1 , given the imbalance Δ and assuming that $m_a - m_b = h$ and Δ/n is small, one can use the approximation: $T^* \approx \frac{h}{2\sigma} \sqrt{n}(1 - \Delta^2(1 + \zeta_1)/(2n^2)) + \mathcal{N}(0, 1)$. As $z_{1-\gamma/2} + z_{1-\delta} = \sqrt{n_{bal}} \frac{h}{2\sigma}$, relation (7) is asymptotically equivalent to a quadratic equation $n_+^2 + n_{bal}n_+ - Q(1 + \zeta_2) = 0$, where $Q = s_0^2 S(1 + \sqrt{2}z_{1-\delta})$. Thus, $n_+ = \frac{n_{bal}}{2} (\sqrt{1 + 4Q(1 + \zeta_2)}/n_{bal}^2 - 1) = Q(1 + \zeta_3)/n_{bal}$. Results of Monte Carlo simulation support this statement for rather wide range of parameters and even for not so large n , e.g. $n = 30$. \square

As usually $S < n_{bal}/2$ and for two treatments $K_1 = 4$, this implies that in general $n_+ \leq 2$. Thus, both randomization schemes lead practically to the same sample size.

Note that the impact of imbalance is concentrated in the term $\Delta^2/2n^2 = O(S/n^2)$ and is negligible at large n . This is in agreement with Lachin (1988).

3.1.3 Impact of patient dropout

Consider the impact of a random patient dropout on a sample size for both randomization schemes on the example of the test that compares means (see Section 3.1.2). Assume that each patient randomized to treatment j will stay till the end of the trial with probability p_j , $j = a, b$. Only these patients will be included into the analysis. The values $q_j = 1 - p_j$, $j = a, b$, define the probabilities of dropout. Let v_j be the number of patients initially randomized to treatment j . Assume that $v_a - v_b = G$, where G is a random variable with mean zero and variance D^2 . As $v_a + v_b = n$, then $v_a = n/2 + G/2$, $v_b = n/2 - G/2$. In this general setting we can combine together the cases of unstratified and stratified randomization, as in the first case $G = 0$, and in the second case $G = \Delta$ and according to Theorem 1, $D^2 \approx s_0^2 S$.

Let n_j^* be the remaining number of patients on treatment j after dropout. Then $n_a^* = \text{Bin}(n/2 + G/2, p_a)$, $n_b^* = \text{Bin}(n/2 - G/2, p_b)$, where $\text{Bin}(k, p)$ is a binomial variable with parameters (k, p) . If G is random, n_a^* and n_b^* are dependent and $\mathbf{E}[n_j^*] = np_j/2$, $\mathbf{Var}[n_j^*] = np_j q_j/2 + D^2 p_j^2/4$, $\mathbf{E}[n_a^* n_b^*] = p_a p_b (n^2 - D^2)/4$. Thus, at large n

$$(n_a^*, n_b^*) \approx \left((n/2)p_a(1 + \psi_a \xi_a / \sqrt{n}), (n/2)p_b(1 + \psi_b \xi_b / \sqrt{n}) \right), \quad (8)$$

where $\psi_j = \sqrt{2q_j/p_j + D^2/n}$, $j = a, b$, and vector (ξ_a, ξ_b) has a bivariate normal distribution, $\mathbf{E}\xi_j = 0$, $\mathbf{Var}\xi_j = 1$, $\mathbf{E}[\xi_a \xi_b] = -D^2/(n\psi_1\psi_2)$. Denote

$$M = \frac{h}{\sigma^2} \sqrt{\frac{p_a p_b}{2(p_a + p_b)}}, \quad R = \frac{D^2 p_a p_b (p_a - p_b)^2}{2n}, \quad B^2 = \frac{h^2}{\sigma^2} \frac{q_a p_b^3 + q_b p_a^3 + R}{4(p_a + p_b)^3}.$$

Under the hypothesis H_1 , after some algebra one can get an approximation for statistic (6) in the form $T^* \approx \sqrt{n}M + \sqrt{1+B^2}\mathcal{N}(0, 1)$. This relation together with (7) implies the relation for the required sample size:

$$n \approx \frac{2\sigma^2(p_a + p_b)}{h^2 p_a p_b} (z_{1-\gamma/2} + \sqrt{1+B^2}z_{1-\delta})^2. \quad (9)$$

Consider now the averaged design (the number of patients on treatments a and b are fixed and equal to $(n/2)p_a$ and $(n/2)p_b$, respectively). Using (5) one can easily establish that the sample size for the averaged design is

$$n_{aver} \approx \frac{2\sigma^2(p_a + p_b)}{h^2 p_a p_b} (z_{1-\gamma/2} + z_{1-\delta})^2.$$

Thus, the sample size increase compared to the averaged design is concentrated in the term B^2 and is practically negligible. For example, if B^2 is rather small,

$$n - n_{aver} \approx \frac{q_a p_b^3 + q_b p_a^3 + R}{2p_a p_b (p_a + p_b)^2} z_{1-\delta} (z_{1-\gamma/2} + z_{1-\delta}). \quad (10)$$

In particular, for $\gamma = \delta = 0.05$ and $p_a = p_b = p$, in the region $p \geq 0.4$ (dropout less than 60%), $n - n_{aver} \leq 2$ (sample size increases by no more than two patients).

The impact of the randomization scheme is concentrated in the term R . For unstratified randomization $R = 0$, while in the case of stratified randomization $R = s_0^2 p_a p_b (p_a - p_b)^2 / (2n)$ and is also rather small. Calculations show that using stratified randomization practically does not lead to sample size increase.

Table 1: Sample size calculations.

h	0.5	0.6	0.7	0.8	0.9	1.0	1.1	1.2	1.3	1.4	1.5
Averaged design	409	284	209	160	127	103	85	71	61	53	46
Unstratified	411	286	211	162	129	105	87	73	63	55	48
Stratified	411	286	211	162	129	105	87	73	63	55	48

Table 1 shows the calculated values of sample sizes for a particular scenario. Consider a study with $S = 10$ strata of equal sizes ($N_s = 1$). Let $\gamma = 0.05$, $\delta = 0.05$, $p_a = 0.4$, $p_b = 0.7$, $K = 2$, block size $K_1 = 4$. Consider three cases: averaged design (randomness in dropout is not accounted for), unstratified randomization and stratified randomization. We set $\sigma^2 = 1$. The sample size is calculated for different values of h in interval $[0.5, 1.5]$. As one can see, a sample size increase accounting for random patient dropout is only two patients, and using stratified randomization does not

lead to an additional sample size increase compared to unstratified randomization. Similar results are true for other scenarios and large number of strata.

4 Conclusions

Using the advanced patient recruitment model allows prediction at the design stage of the number of patients randomized to different treatment arms in different strata and investigation of the properties of imbalance caused by stratified randomization and its impact on the power and sample size of the trial. For two treatment arms with interest in a statistical test that compares means, it is shown, that the sample size increase required to compensate for random imbalance is practically negligible. Randomness in patient dropout also leads to a negligible sample size increase compared to averaged design (fixed number of randomized patients). These results show that stratified randomization even for a large number of strata does not lead to a visible sample size increase compared to unstratified randomization.

The type of randomization may affect other characteristics of the trial, e.g. centre-stratified randomization in general requires less drug supply compared to unstratified randomization. Thus, in the cases when the choice of randomization is not dictated by the type of data, it is beneficial to use various criteria accounting for sample size, recruitment and supply costs, etc., when choosing a randomization scheme.

References

- Anisimov, V. V. (2007). Effect of imbalance in using stratified block randomization in clinical trials. *Bulletin of the International Statistical Institute - LXII, Proc. of the 56 Annual Session, Lisbon*, 5938–5941.
- Anisimov, V. V. and V. V. Fedorov (2006). Design of multicentre clinical trials with random enrolment. In *Advances in Statistical Methods for the Health Sciences* (N. Balakrishnan, J. L. Auget, M. Mesbah, and G. Molenberghs eds). Berlin: Birkhäuser, 387–400.
- Anisimov, V. V. and V. V. Fedorov (2007). Modeling, prediction and adaptive adjustment of recruitment in multicentre trials. *Statistics in Medicine* 26, 4958–4975.
- Hallstrom, A. and K. Davis (1988). Imbalance in treatment assignments in stratified blocked randomization. *Controlled Clinical Trials* 9, 375–382.
- Lachin, J. M. (1988). Statistical properties of randomization in clinical trials. *Controlled Clinical Trials* 9, 289–311.
- Matts, J. P. and J. M. Lachin (1988). Properties of permuted-block randomization in clinical trials. *Controlled Clinical Trials* 9, 327–345.
- Pocock, S. J. (1983). *Clinical Trials. A Practical Approach*. New York: Wiley.
- Rosenberger, W. F. and J. M. Lachin (2002). *Randomization in Clinical Trials*. New York: Wiley.

The Non-Uniqueness of Some Designs for Discriminating Between Two Polynomial Models in One Variable

Anthony C. Atkinson

Abstract T-optimum designs for discriminating between two nested polynomial regression models in one variable that differ in the presence or absence of the two highest order terms are studied as a function of the values of the parameters of the true model. For the value of the parameters corresponding to the absence of the next-highest order term, the optimum designs are not unique and can contain an additional support point. A numerical exploration of the non-uniqueness reveals a connection with D_S -optimality for models which do contain the next highest term. Brief comments are given on the analysis of data from such designs

1 Introduction

T-optimum designs for discriminating between two regression models were introduced by Atkinson and Fedorov (1975). More recently, Dette and Titoff (2008) explored the structure of T-optimum designs in some detail. One of their examples was of discrimination between linear and cubic models in one variable. For particular parameter values the T-optimum design was not unique, consisting of convex combinations of two extreme designs. This example can be thought of as an extension of Example 1 of Atkinson and Fedorov in which designs were found for discrimination between a constant and a general quadratic. The paper illustrates how the designs depend upon the parameters of the true model and gives a geometric interpretation of the occurrence of non-unique designs as a function of the response.

The non-unique designs occur when the larger model contains a term of order x^k and all lower order terms except that of order x^{k-1} , the smaller model containing terms up to order x^{k-2} . The structure of these non-unique designs is explored numerically for k in the range two to six. A relationship is indicated with D_S -optimum designs for the estimation of the coefficient of x^k in a polynomial model which adds a term in x^{k-1} to those of the larger model.

Prof. Anthony C. Atkinson
London School of Economics, London WC2A 2AE, UK, e-mail: a.c.atkinson@lse.ac.uk

The plan of the paper is as follows. The background theory for T-optimality is in the next section. Two examples are in §3. Breaks in the structure of the designs as functions of the parameters are shown to occur for the two polynomial examples as one parameter goes to zero. Section 4 explores the structure of the designs when the coefficient of x^{k-1} is zero. The paper concludes with brief comments on data analysis and the power of tests as a function of the number of support points of a design.

2 Background

The T-optimum design for discriminating between two models depends upon which model is true and, usually, on the values of some of the parameters of the true model. Without loss of generality let this be the first model and write

$$y = \eta(x) + \varepsilon, = \eta_1(x, \theta_1) + \varepsilon, \quad (1)$$

where the errors ε are i.i.d $\mathcal{N}(0, \sigma^2)$. A good design for discriminating between the models will provide a large lack-of-fit sum of squares for the second model. When the second model is fitted to the data, the least squares estimates of the $p_2 \times 1$ parameter θ_2 depend on the experimental design as well as on the value of θ_1 and the errors. In the absence of error the parameter estimates are

$$\hat{\theta}_2(\xi) = \operatorname{argmin}_{\theta_2} \int_{\mathcal{X}} \{\eta(x) - \eta_2(x, \theta_2)\}^2 \xi(dx), \quad (2)$$

yielding a residual sum of squares

$$\Delta(\xi) = \int_{\mathcal{X}} [\eta(x) - \eta_2\{x, \hat{\theta}_2(\xi)\}]^2 \xi(dx). \quad (3)$$

For linear models $\Delta(\xi)$ is proportional to the non-centrality parameter of the χ^2 distribution of the residual sum of squares for the second model when the design is ξ . T-optimum designs maximise $\Delta(\xi)$ and so provide the most powerful test for lack of fit of the second model when the first is true. In general, T-optimum designs have $p_2 + 1$ points of support.

3 Examples

Example 1. Constant Against Quadratic

Atkinson and Fedorov (1975) exhibit designs for discrimination between the models

$$\eta(x) = \beta_0 + \beta_1 x + \beta_2 x^2 \text{ and } \eta_2 = \beta_0. \tag{4}$$

The T-optimum design depends on the ratio β_1/β_2 , but not on the magnitude of the parameters which will, however, affect the magnitude of the non-centrality parameter. Atkinson and Fedorov (1975) reparameterise by taking $\beta_1 = \cos \phi$ and $\beta_2 = \sin \phi$. Their Figure 1 shows the support points of the design for $0 \leq \phi \leq 90^\circ$.

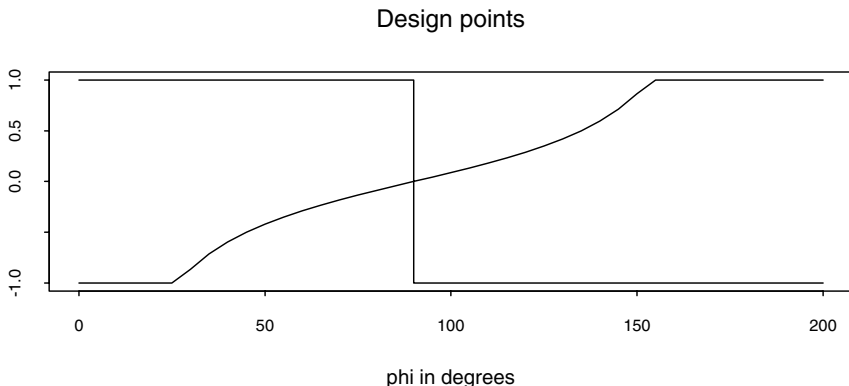


Fig. 1: Example 1: constant against quadratic model. Support points of T-optimum design with $\beta_1 = \cos \phi$ and $\beta_2 = \sin \phi$ when $\mathcal{X} = [-1, 1]$

In the general case the T-optimum design puts equal weight at the two points of support of the design which are at the minimum and maximum of the quadratic over the design region, taken as $\mathcal{X} = [-1, 1]$. Differentiation of $\eta(x)$ shows that the turning point of the quadratic is at $x^* = -0.5 \cot(\phi)$. When $\phi < 26^\circ 54'' = \arctan(0.5)$ this value lies outside the experimental region and, as Figure 1 shows, the support points of the design are at ± 1 . For larger values of ϕ the support points, up to 90° are at x^* and 1. Above 90° the support points are -1 and x^* until $\phi \geq 153^\circ 26''$ when the points again become -1 and 1. The figure repeats for values of $\phi > 180^\circ$.

Three special values are of interest. When $\phi = 0$, $\beta_2 = 0$ and the model is a straight line, when the maximum and minimum of $\eta(x)$ are unambiguous. However, when $\phi = 90^\circ$ the model is a pure quadratic. There are two equal maxima of the function at -1 and $+1$ with a minimum at $x = 0$. Thus one T-optimum design puts half the weight at -1 and 0 and another, equally good, design is its reflection putting half the weight at 0 and half at 1. Any convex linear combination of these designs will also be T-optimum so that the most general T-optimum design is

$$\xi_T^* = \left\{ \begin{matrix} -1 & 0 & 1 \\ 0.5\lambda & 0.5 & 0.5(1-\lambda) \end{matrix} \right\} \quad (0 \leq \lambda \leq 1). \tag{5}$$

Perhaps the most interesting of these designs is that for $\lambda = 0.5$ which is also the D_1 -optimum design for β_2 in $\eta(x)$. We return to this design in §4. For values of ϕ close to 90° this design has good T-efficiency as measured by the value of $\Delta(\xi)$.

The third value of interest in Figure 1 is $\phi = 180^\circ$ when the model is again first-order, although with a negative slope. For values of ϕ around 180° the design puts half the weight at -1 and the other half at 1 . The only break in the smooth evolution of the designs in the figure with ϕ is at 90° , for which value there is the family of designs given by (5). The same design is optimum when $\phi = 270^\circ$; now the minima of the quadratic are at $x = \pm 1$ and the maximum is at 0 .

Example 2. Linear Against Cubic

Dette and Titoff (2008) extend Example 1 to a linear regression against a cubic so that (4) becomes

$$\eta(x) = \beta_0 + \beta_1 x + \beta_2 x^2 + \beta_3 x^3 \quad \text{and} \quad \eta_2(x) = \beta_0 + \beta_1 x. \quad (6)$$

With $\eta_2(x)$ containing two parameters, the unique T-optimum designs have three points of support.

Again consider a trigonometric transformation. We now take $\beta_2 = \cos \phi$ and $\beta_3 = \sin \phi$, again with $\mathcal{X} = [-1, 1]$. The support points of the T-optimum designs are shown in the upper panel of Figure 2 with the design weights in the lower panel.

The general structure of the designs is similar to that shown in Figure 1, with the non-unique design at $\phi = 90^\circ$. When $\phi = 0$, $\eta(x)$ is a pure quadratic and the design is the D_1 -optimum design for β_2 , namely with support points $-1, 0$ and 1 and weights $0.25, 0.5$ and 0.25 . As ϕ increases to 45° the value of the central support point increases as does the weight on $x = 1$. For all designs the weight on the central support point is 0.5 .

When $\phi = 45^\circ$, $\beta_2 = \beta_3$. The design weights are $1/6, 1/2$ and $1/3$, which values are optimum for all designs up to $\phi = 90^\circ$. Above $\phi = 45^\circ$ the lower design point increases away from -1 , so that the designs no longer span the design region. The two lower design points continue to increase until $\phi = 90^\circ$ when $\beta_2 = 0$ and $\eta(x)$ contains a cubic term, but no quadratic. Again at this value of ϕ there are two extreme T-optimum designs; one design has support points $-0.5, 0.5$ and 1 . Another is the reflection of this with support points $-1, -0.5$ and 0.5 . As for Example 1, the convex linear combination of these designs will also be T-optimum so that the most general T-optimum design is

$$\xi_T^* = \left\{ \begin{array}{cccc} -1 & -0.5 & 0.5 & 1 \\ \lambda/3 & (1+2\lambda)/6 & (3-2\lambda)/6 & (1-\lambda)/3 \end{array} \right\} \quad (0 \leq \lambda \leq 1), \quad (7)$$

which is a reparameterisation of Dette and Titoff's (2.14). When $\lambda = 0.5$ we obtain the D_1 -optimum design for β_3 in $\eta_1(x)$, extending the result for the same design criterion when $\phi = 90^\circ$ but for β_2 in Example 1.

For values of $\phi > 90^\circ$ the designs are the reflection in \mathcal{X} of those for $180^\circ - \phi$. As the figure shows, the cycle of designs repeats itself for values of ϕ above 180° .

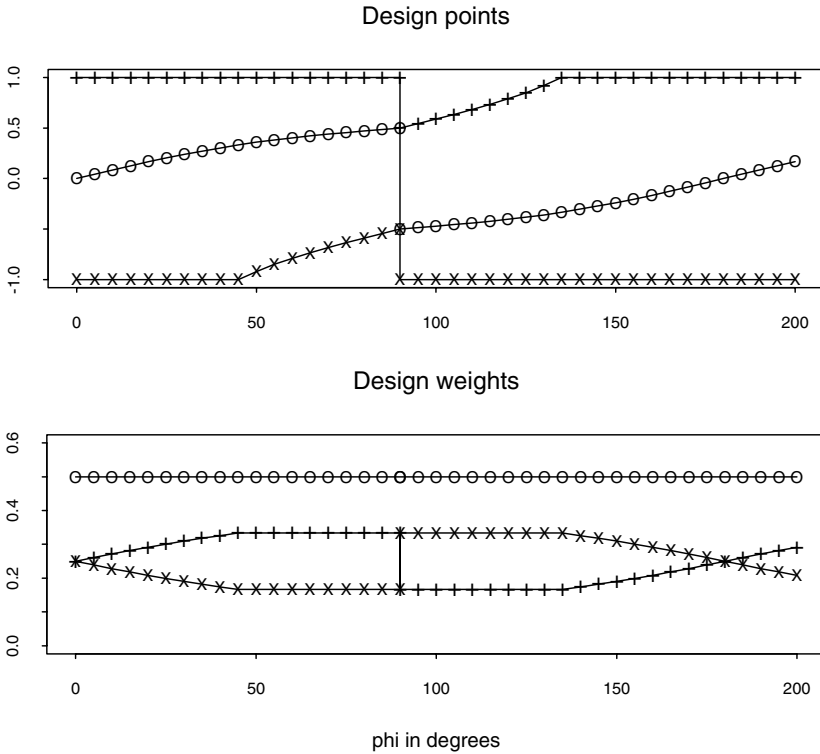


Fig. 2: Example 2: linear against cubic model. Upper panel: support points of T-optimum design with $\beta_2 = \cos \phi$ and $\beta_3 = \sin \phi$ when $\mathcal{X} = [-1, 1]$. Lower panel: design weights. The same symbols are used for the three support points in the two panels

4 Designs for Higher-Order Models

The designs obtained above for $\phi = 90^\circ$ are special cases of a more general discrimination problem in which the models are

$$\eta_2(x) = \sum_{j=0}^{k-2} \beta_j x^j \text{ and } \eta(x) = \eta_2(x) + \beta_k x^k, \tag{8}$$

where, now, β_k is not constrained to equal one. The two models thus differ by a single term, but with the term in x^{k-1} absent from both.

For linear models differing by a single parameter the value of $\Delta(\xi)$ for the T-optimum design depends on the value of the extra parameter, here β_k . However, the T-optimum design does not depend on this value and is identical to the D_1 -optimum design.

Table 1 gives numerically obtained T- and D_1 -optimum designs for k from two to six. The designs shown have one support point at $x = -1$. Otherwise the support

Table 1: Identical T- and D_1 -optimum designs for the term of order k in the pair of polynomial models (8). Reflected designs with the signs of all x values reversed are also optimum

k	Optimum Design
2	$\begin{Bmatrix} -1 & 0 \\ 1/2 & 1/2 \end{Bmatrix}$
3	$\begin{Bmatrix} -1 & -0.5 & 0.5 \\ 1/3 & 1/2 & 1/6 \end{Bmatrix}$
4	$\begin{Bmatrix} -1 & -\sqrt{2}/2 & 0 & \sqrt{2}/2 \\ 1/4 & 0.427 & 1/4 & 0.073 \end{Bmatrix}$
5	$\begin{Bmatrix} -1 & -0.809 & -0.309 & 0.309 & 0.809 \\ 1/5 & 0.362 & 0.262 & 0.138 & 0.038 \end{Bmatrix}$
6	$\begin{Bmatrix} -1 & -\sqrt{3}/2 & -0.5 & 0 & 0.5 & \sqrt{3}/2 \\ 1/6 & 0.311 & 1/4 & 1/6 & 0.083 & 0.022 \end{Bmatrix}$

points, but not the weights, are symmetrical around $x = 0$. There is appreciable structure in the results. For example, the weights at $x = -1$ are $1/k$. These and the other ratios in the table, including $\sqrt{2}/2$ and $\sqrt{3}/2$, are accurate to 5 decimal places in the numerical results.

To demonstrate that these numerically obtained designs are indeed optimum, the derivative function for the appropriate equivalence theorem was used. In general, for D_s -optimum designs, the variance $d_s(x, \xi^*)$ (see, for example, Atkinson, Donev, and Tobias 2007, p. 139), takes its maximum value of s at the points of support of the design. Figure 3 shows the plot of the variance function over the design region for the case of $k = 6$. Indeed the maximum values of the function are one and occur at the points of support of the design.

The main interest in this section is whether the designs are unique for these higher values of k . Figure 3 also provides an answer to this question. The curve of the variance is symmetrical with a value of one at $x = 1$, which is not a support point of the design, a phenomenon indicative of non-uniqueness of the design. Indeed, from the symmetry of the reflected designs, it follows that the mirror image of the design for $k = 6$ in Table 1 will have the same plot of the variance function as that of Figure 3. Thus, as for the examples for $k = 2$ and 3 in the previous section, the design is not unique and any convex linear combination will also be a T- and D_1 -optimum design for $k = 6$. Similar numerical results hold for the other values of k in Table 1.

A last comment is on the designs found by averaging the designs of Table 1 and their reflections, that is the combinations with $\lambda = 0.5$. The numerical results in the table show that such designs have weights $1/2k$ at the ends of the design region and weights $1/k$ at the $k - 1$ remaining points. They are, in fact, the D_1 -optimum designs given by Kiefer and Wolfowitz (1959) for β_k , but not in $\eta(x)$ in (8), but rather for the model also including a term in x^{k-1} . The support points of these designs are

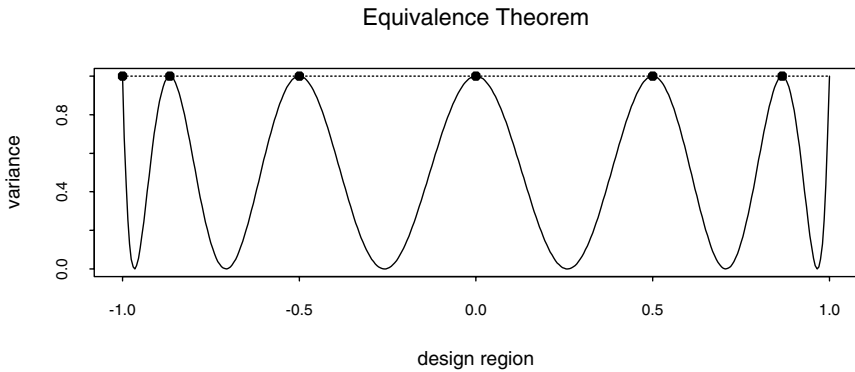


Fig. 3: Fourth-order model against sixth-order: variance $d_S(x, \xi^*)$ for the optimum design of Table 1. The value is one at the points of support of the design, which are marked \bullet . Note that $x = 1$ is not a support point of the design

$$x_j = -\cos\{(j\pi)/k\} \quad (j = 0, \dots, k),$$

so extending the results of §3 for $\lambda = 0$. An advantage of the designs for this value of $\lambda = 0.5$ is that the weights are much more equal than those for the designs of Table 1. They are therefore more accurately approximable to give small integer designs. However they do have one more point of support than the asymmetrical designs.

5 Design and the Analysis of Data

In all examples the designs with minimum support have $p_2 + 1$ points of support. In the analysis of data from such designs it will therefore not be possible to estimate $\eta(x)$. The analysis of variance table will consist of a sum of squares for $\eta_2(x)$, a sum of squares for pure error from replication and a lack of fit sum of squares, with non-centrality parameter a multiple of $\Delta(\xi)$, on one degree of freedom.

The larger model $\eta(x)$ can be estimated from designs with at least $p_2 + 2$ support points. For the special case of §4, designs with $(0 < \lambda < 1)$ have this support. But now the lack of fit sum of squares will have the same non-centrality parameter as the component designs with $p_2 + 1$ support points, but on two degrees of freedom so that the power will be reduced. A test with one degree of freedom in the numerator can be obtained by breaking this sum of squares on two degrees of freedom into individual components. However the result, for a sufficiently large number of observations, will be a significance test for the highest term in $\eta(x)$, with the next highest term not significant. Such models are not usually recommended. Indeed, Nelder (1998) states that they “are of very limited interest”. If the purpose of the experiment is not

only to test the smaller model but to fit the larger if the smaller is rejected, the family of DT-optimum designs (Atkinson 2008) is appropriate.

The discussion in this paper is in terms of least squares and known linear models. For nonlinear models the relationship between T- and D_1 -optimality is more complicated (López-Fidalgo, Tommasi, and Trandafir 2008; Atkinson and Bogacka 2010). López-Fidalgo, Trandafir, and Tommasi 2007 describe designs using Kullback-Leibler distance for discriminating between non-normal models. Wiens (2009) extends this work to designs when the models are only approximately specified and considers designs both for model discrimination and parameter estimation.

Acknowledgements I am most grateful to a referee who suggested exploring the properties of designs for discriminating between the pairs of models (8) and so led me to the results reported in §4.

References

- Atkinson, A. C. (2008). DT-optimum designs for model discrimination and parameter estimation. *Journal of Statistical Planning and Inference* 138, 56–64.
- Atkinson, A. C. and B. Bogacka (2010). Optimum designs for the equality of parameters in enzyme inhibition kinetic models. (Submitted).
- Atkinson, A. C., A. N. Donev, and R. D. Tobias (2007). *Optimum Experimental Designs, with SAS*. Oxford: Oxford University Press.
- Atkinson, A. C. and V. V. Fedorov (1975). The design of experiments for discriminating between two rival models. *Biometrika* 62, 57–70.
- Dette, H. and S. Titoff (2008). Optimal discrimination designs. *Annals of Statistics* 37, 2056–2082.
- Kiefer, J. and J. Wolfowitz (1959). Optimum designs in regression problems. *Annals of Mathematical Statistics* 30, 271–294.
- López-Fidalgo, J., C. Tommasi, and C. Trandafir (2008). Optimal designs for discriminating between some extensions of the Michaelis-Menten model. *Journal of Statistical Planning and Inference* 138, 3797–3804. doi: 10.1016/j.jspi.2008.01.014.
- López-Fidalgo, J., C. Trandafir, and C. Tommasi (2007). An optimal experimental design criterion for discriminating between non-normal models. *Journal of the Royal Statistical Society, Series B* 69, 231–242.
- Nelder, J. A. (1998). The selection of terms in response surface models - how strong is the weak heredity principle? *The American Statistician* 52, 315–318.
- Wiens, D. P. (2009). Robust discrimination designs. *Journal of the Royal Statistical Society, Series B* 71, 805–829.

Covariate Adjusted Designs for Combining Efficiency, Ethics and Randomness in Normal Response Trials

Alessandro Baldi Antognini¹ and Maroussa Zagoraiou¹

Abstract This paper deals with the problem of allocating patients to two competing treatments in the presence of covariates in order to achieve a good trade-off between efficiency, ethical concern and randomization. We propose a compound criterion that combines inferential precision and ethical gain by flexible weights depending on the unknown treatment effects. In the absence of treatment-covariate interactions, this criterion leads to a locally optimal allocation which does not depend on the covariates and can be targeted by a suitable implementation of the doubly-adaptive biased coin design aimed at balancing the roles of randomization, ethics and information. Some properties of the suggested procedure are described.

1 Introduction

Patients arrive sequentially in a clinical trial for comparing two treatments, say A and B , where some concomitant variables, like the subjects' gender, medical history etc., potentially affect the experimental outcome. Suppose that for each incoming subject, before assigning either treatment we observe a vector \mathbf{Z} of categorical covariates (also called block factors), which is assumed to be random, i.e. not under the experimenters' control. Let the treatments be assigned according to a given randomization rule, with $\delta_i = 1$ if the i -th subject is allocated to A and 0 otherwise, and suppose that a normal response Y_i is observed belonging to the linear homoscedastic model in the form

$$E(Y_i) = \delta_i \mu_A + (1 - \delta_i) \mu_B + \mathbf{z}_i' \boldsymbol{\beta}, \quad V(Y_i) = \sigma^2, \quad i = 1, \dots, n, \quad (1)$$

Alessandro Baldi Antognini, Maroussa Zagoraiou
Department of Statistical Sciences, University of Bologna, Via Belle Arti 41, 40126, Bologna, Italy, e-mail: a.baldi@unibo.it; maroussa.zagoraiou@unibo.it

where \mathbf{z}_i represents the covariate profile of the i -th individual and $\boldsymbol{\beta}$ is the vector of covariate effects. To avoid cumbersome notation, from now on we take into account only one categorical covariate Z with levels z_1, \dots, z_J , so that \mathbf{z}_i is the $(J-1)$ -dim vector of indicators for the i -th subject. Also, we denote by $p_j = \Pr(Z = z_j)$ for $j = 1, \dots, J$ the probability distribution of the covariate in the population of interest.

At the end of the trial, suppose that n assignments of either treatment A or B have been made to patients with i.i.d. covariates Z_1, \dots, Z_n where, conditionally on the covariates and the treatment allocations, patients' responses are assumed to be independent. Let $N_{nj} = \sum_{i=1}^n \mathbb{1}_{\{Z_i=z_j\}}$ denote the number of subjects with covariate profile z_j , where $\mathbb{1}_{\{\cdot\}}$ is the indicator function, and $\mathbf{N}_n = (N_{n1}, \dots, N_{nJ})^t$ with $\sum_{j=1}^J N_{nj} = n$. Moreover, $\boldsymbol{\pi}_n = (\pi_{n1}, \dots, \pi_{nJ})^t$ represents the vector of allocation proportions to A for each profile, where $\pi_{nj} = \sum_{i=1}^n \delta_i \mathbb{1}_{\{Z_i=z_j\}} / N_{nj}$, and $D_{nj} = N_{nj}(2\pi_{nj} - 1)$ is the current imbalance between the two groups for the j -th covariate level. Hence $D_n = \sum_{j=1}^J D_{nj}$ is the global imbalance.

2 Optimal Allocations for Inference

Under model (1), the covariates affect the treatment responses in the same way for all patients with the same profile, so that the superiority of A to B or vice-versa, is *uniformly* constant over the covariates and it is customary to regard $\boldsymbol{\beta}$ as a nuisance parameter. Consequently, the inferential interest typically lies in estimating the difference $\alpha = \mu_A - \mu_B$ between the treatment effects as precisely as possible. Following Atkinson (2002), model (1) can be written as

$$E(\mathbf{Y}_n) = (2\boldsymbol{\delta}_n - \mathbf{1}_n)\alpha + \mathbf{F}\boldsymbol{\theta}, \quad (2)$$

where $\mathbf{Y}_n = (Y_1, \dots, Y_n)^t$, $\boldsymbol{\delta}_n = (\delta_1, \dots, \delta_n)^t$ and $\mathbf{1}_n$ is the n -dim vector of 1's. Notice that, under (2) the constant term and covariates are included in the 0-1 matrix \mathbf{F} and the nuisance $\boldsymbol{\theta}$ incorporates $\boldsymbol{\beta}$ and an overall effect.

From an inferential viewpoint the design problem consists in finding the allocation $\boldsymbol{\pi}_n^* = (\pi_{n1}^*, \dots, \pi_{nJ}^*)^t$ which minimizes a suitable measure Φ_I of loss of precision, called the inferential criterion. In this setting it is customary to assume the well-known D_A -optimality, so let $\hat{\alpha}_n$ be the OLS (or ML) estimator of α , given the covariates $\mathbf{Z}_n = (Z_1, \dots, Z_n)^t$ and the design $\boldsymbol{\delta}_n$, we have

$$V(\hat{\alpha}_n | \mathbf{Z}_n, \boldsymbol{\delta}_n) = \sigma^2 \{n - \mathbf{b}'(\mathbf{F}'\mathbf{F})^{-1}\mathbf{b}\}^{-1} = \sigma^2 \{n - L_n\}^{-1},$$

where $\mathbf{b} = \mathbf{F}'(2\boldsymbol{\delta}_n - \mathbf{1}_n) = (D_n, D_{n1}, \dots, D_{n(J-1)})^t$ and $L_n = \mathbf{b}'(\mathbf{F}'\mathbf{F})^{-1}\mathbf{b}$ is the loss after n assignments (see for instance Atkinson 2002). Since

$$\mathbf{F}'\mathbf{F} = \left(\begin{array}{c|c} n & \tilde{\mathbf{N}}_n^t \\ \hline \tilde{\mathbf{N}}_n & \text{diag}(\tilde{\mathbf{N}}_n) \end{array} \right), \quad (3)$$

where $\tilde{\mathbf{N}}_n = (N_{n1}, \dots, N_{n(J-1)})^t$ and $\text{diag}(\tilde{\mathbf{N}}_n)$ is the diagonal matrix with j -th entry N_{nj} (for $j = 1, \dots, J-1$), the loss L_n depends on the covariates only through the number N_{nj} of patients within each profile, which is not under the experimental control, whereas it depends on the design through $\boldsymbol{\pi}_n$, hence from now on we write $L_n = L(\boldsymbol{\pi}_n, \mathbf{N}_n)$ and $V(\hat{\alpha}_n | \boldsymbol{\pi}_n, \mathbf{N}_n)$.

As is well-known, in the case of perfect balance between the treatment groups within each profile the loss is identically zero independently of the covariates.

Lemma 1. *The loss after n steps can be written as*

$$L(\boldsymbol{\pi}_n, \mathbf{N}_n) = \sum_{j=1}^J (2\pi_{nj} - 1)^2 N_{nj}, \quad (4)$$

which is minimized by $\boldsymbol{\pi}_I^* = (\frac{1}{2}, \dots, \frac{1}{2})^t$, independently of \mathbf{N}_n .

Proof. Using some results on the inverses of partitioned matrices, from (3)

$$(\mathbf{F}^t \mathbf{F})^{-1} = \frac{1}{N_{nJ}} \cdot \left(\begin{array}{c|c} 1 & -\mathbf{1}'_{J-1} \\ \hline -\mathbf{1}'_{J-1} & \mathbf{J}_{J-1} \end{array} \right) + \left(\begin{array}{c|c} 0 & \mathbf{0}'_{J-1} \\ \hline \mathbf{0}_{J-1} & (\text{diag}(\tilde{\mathbf{N}}_n))^{-1} \end{array} \right),$$

where \mathbf{J}_k is the $(k \times k)$ matrix of ones and $\mathbf{0}_k$ is the k -dim vector of zeros. Thus, the loss in (4) can be derived by simple algebra since

$$\mathbf{b}^t \left(\begin{array}{c|c} 1 & -\mathbf{1}'_{J-1} \\ \hline -\mathbf{1}'_{J-1} & \mathbf{J}_{J-1} \end{array} \right) \mathbf{b} = D_{nJ}^2$$

and

$$\mathbf{b}^t \left(\begin{array}{c|c} 0 & \mathbf{0}'_{J-1} \\ \hline \mathbf{0}_{J-1} & (\text{diag}(\tilde{\mathbf{N}}_n))^{-1} \end{array} \right) \mathbf{b} = \sum_{j=1}^{J-1} \frac{D_{nj}^2}{N_{nj}}.$$

□

Observe that the loss $L(\boldsymbol{\pi}_n, \mathbf{N}_n)$ in (4) is a r.v. depending on \mathbf{N}_n and therefore, in order to remove the effect due to the random nature of the covariates and to derive a standardized criterion varying in $[0; 1]$, from now on we take into account the loss of design efficiency

$$\Phi_I(\boldsymbol{\pi}_n) = 1 - \text{Eff}(\boldsymbol{\pi}_n) = 1 - [V(\hat{\alpha}_n | \boldsymbol{\pi}_I^*) / V(\hat{\alpha}_n | \boldsymbol{\pi}_n)], \quad (5)$$

where, since $V_{\mathbf{Z}}[E(\hat{\alpha}_n | \boldsymbol{\pi}_n, \mathbf{N}_n)] = 0$, it follows that $V(\hat{\alpha}_n | \boldsymbol{\pi}_n) = E_{\mathbf{Z}}[V(\hat{\alpha}_n | \boldsymbol{\pi}_n, \mathbf{N}_n)] = E_{\mathbf{Z}}[\sigma^2 \{n - L(\boldsymbol{\pi}_n, \mathbf{N}_n)\}^{-1}]$. As is well-known, $V(\hat{\alpha}_n | \boldsymbol{\pi}_n)$ can be approximated by $\sigma^2 \{n - E_{\mathbf{Z}}[L(\boldsymbol{\pi}_n, \mathbf{N}_n)]\}^{-1}$, so that the inferential criterion (5) becomes

$$\Phi_I(\boldsymbol{\pi}_n) = E_{\mathbf{Z}}[n^{-1} L(\boldsymbol{\pi}_n, \mathbf{N}_n)] = \sum_{j=1}^J (2\pi_{nj} - 1)^2 p_j, \quad (6)$$

since \mathbf{N}_n is distributed according to a multinomial r.v. $MN(n; p_1, \dots, p_J)$.

3 Optimal Allocations for Ethics

Assuming that the treatment effects are different and that larger is better, the ethical cost-per-observation can be measured by the percentage of patients who receive the worse treatment, i.e. $\frac{1}{2} + [\frac{1}{2} - \frac{1}{n} \sum_{j=1}^J \pi_{nj} N_{nj}] \text{sgn}(\alpha)$, where $\text{sgn}(x)$ is the sign of x . Note that $\sum_{j=1}^J \pi_{nj} N_{nj}$ is the number of assignments to A in the trial, stressing that, in the absence of treatment-covariate interactions the ethical gain depends on the design only through the total proportion of allocations to A . As previously, from now on we consider the expected percentage of subjects assigned to the worse treatment as our ethical criterion

$$\Phi_E(\boldsymbol{\pi}_n) = \frac{1}{2} + \left(\frac{1}{2} - \sum_{j=1}^J \pi_{nj} p_j \right) \text{sgn}(\alpha). \quad (7)$$

The optimal ethical target minimizing (7) is $\boldsymbol{\pi}_E^* = (\mathbb{1}_{\{\alpha > 0\}}, \dots, \mathbb{1}_{\{\alpha > 0\}})^t$, which is constant over the covariate levels and depends only on $\alpha = \mu_A - \mu_B$.

Note that (7) is well defined if and only if the treatment effects are different. Indeed, $\mu_A = \mu_B$ means that there is no longer a worse treatment, stressing that the comparative experiment degenerates to observing just one treatment. Thus, from now on we exclude the situation $\alpha = 0$, since in this case criterion (7) no longer depends on the design and the need to derive any kind of compromise between ethics and inference vanishes.

4 Compound Optimal Designs

Several approaches have been proposed in the recent literature in order to achieve a compromise between information gain and ethical concern for the subjects involved in the trial. See for instance Bandyopadhyay and Biswas (2001), Rosenberger *et al.* (2001), Atkinson and Biswas (2005), Geraldès *et al.* (2006) and Tymofyeyev *et al.* (2007). In order to obtain a valid trade-off, we suggest a compromise based on a suitable weighted combination of the ethical criterion Φ_E in (7) and the inferential one Φ_I in (6) of the form¹:

$$\Phi_\omega(\boldsymbol{\pi}_n) = \omega \Phi_E(\boldsymbol{\pi}_n) + (1 - \omega) \Phi_I(\boldsymbol{\pi}_n). \quad (8)$$

Clearly, if the treatment effects differ substantially, more care is required for the ethical aspects, whereas when μ_A and μ_B tend to be similar, less attention may be paid to ethics and more to inference, since it is harder to discriminate between A and B . Thus, the relative importance of the two criteria changes on the basis of the treatment effects, so we shall assume the weight ω assigned to ethics to be a function

¹ A thorough discussion of the properties of the suggested compound criterion has been submitted for publication elsewhere.

of the unknown difference $\alpha = \mu_A - \mu_B$. In order to treat A and B symmetrically, we choose $\omega = \omega(|\alpha|) : \mathbb{R}^+ \rightarrow [0; 1)$ to be a continuous and non-decreasing function with $\omega \rightarrow 0$ for $\alpha \rightarrow 0$. It is easy to see that $\Phi_\omega(\cdot)$ is a convex function of $\boldsymbol{\pi}_n$, since it is a weighted combination of a linear and a convex function, and therefore there exists a unique optimal allocation $\boldsymbol{\pi}_\omega^*$ minimizing (8), which depends in general on the unknown parameter α .

Theorem 1. *If the weight function is chosen so that $\omega < \frac{4}{5}$, then the optimal target $\boldsymbol{\pi}_\omega^* = (\pi_{\omega 1}^*, \dots, \pi_{\omega J}^*)^t$ minimizing Φ_ω in (8) is given by*

$$\pi_{\omega j}^* = \frac{\text{sgn}(\alpha)}{8} \frac{\omega}{(1-\omega)} + \frac{1}{2} \in (0; 1) \quad \text{for any } j = 1, \dots, J. \quad (9)$$

Otherwise, if $\omega \geq \frac{4}{5}$ for some values of $|\alpha|$, then the optimal target $\boldsymbol{\pi}_\omega^$ may become trivial and may coincide with the optimal ethical allocation $\boldsymbol{\pi}_E^*$.*

Proof. By putting the first order partial derivatives of the compromise criterion $\Phi_\omega(\cdot)$ in (8) w.r.t. π_{nj} 's equal to zero, we have

$$4[2\pi_{nj} - 1] = \text{sgn}(\alpha)[\omega/(1-\omega)] \quad \text{for any } j = 1, \dots, J, \quad (10)$$

where the left-hand side of the equation varies in $[-4; 4]$ and the right-hand side in \mathbb{R} . Thus, when $\omega(|\alpha|) < \frac{4}{5}$ for any $|\alpha|$, the solution of equation (10) lies in $(0; 1)$; otherwise the compromise criterion $\Phi_\omega(\cdot)$ may become monotonically increasing (if $\alpha < 0$) or decreasing (when $\alpha > 0$), so that the optimal compound target can degenerate to $\boldsymbol{\pi}_E^*$. \square

Assuming model (1), the optimal compound target does not depend on the covariate profiles and is strictly related to the chosen weight function ω , since $\lim_{\omega \rightarrow 0} \boldsymbol{\pi}_\omega^* = \boldsymbol{\pi}_I^*$, $\lim_{\omega \rightarrow 1} \boldsymbol{\pi}_\omega^* = \boldsymbol{\pi}_E^*$ and, from (9), its behaviour depends on

$$\frac{\partial \pi_{\omega j}^*}{\partial \alpha} = \frac{1}{8[1-\omega(|\alpha|)]^2} \left(\frac{\partial \omega(|\alpha|)}{\partial \alpha} \right),$$

which can be regarded as the “ethical improvement rate”. Thus, the weight can be chosen in order to model the optimal compound target in terms of the desired ethical impact as α varies. For instance, a suitable class of weight functions is

$$\omega_s(x) = (4/5) (1+x^{-2})^{-2s} \left[2 - (1+x^{-2})^{-2} \right], \quad \text{with } s \geq 1$$

and the following figure shows the behaviour of ω_s for $s = 2, 3, 5$ and of the corresponding optimal compound target (9).




Fig. 1: Plots of the weight function ω_s (left) and the corresponding optimal compound target π_ω^* (right) for $s = 2, 3, 5$ (solid, dashed and dotted, respectively) as α varies in \mathbb{R}^+ .

5 Doubly-Adaptive Biased Coin Designs with Covariates

Letting $\omega < 4/5$, the compound optimal allocation in (9) is non-trivial and depends on the unknown parameter α , i.e. $\pi_\omega^* = \pi_\omega^*(\alpha)$, so that it can be targeted by applying suitable modified versions of the doubly-adaptive biased coin design (Eisele 1994) adjusted for covariates. A natural extension to the present setting is: i) start with a pilot stage performed to derive initial parameter estimates; ii) at each step k estimate α with all the collected data up to that step by $\hat{\alpha}_k$, so the optimal target (9) can be estimated by $\pi_{\omega_j}^*(\hat{\alpha}_k)$; iii) when the next patient with covariate $Z_{k+1} = z_j$ is ready to be randomized, he/she will be allocated to A with probability $g\left\{\pi_{kj}; \pi_{\omega_j}^*(\hat{\alpha}_k)\right\}$ for $j = 1, \dots, J$, where the function g satisfies the following conditions:

- C1) $g(x; y)$ is continuous on $(0; 1)^2$, with $g(x; x) = x$;
- C2) $g(x; y)$ is strictly decreasing in x and strictly increasing in y .

Assuming this approach, Zhang and Hu (2009) suggest the family

$$g_\gamma(x; y) = \frac{y(y/x)^\gamma}{y(y/x)^\gamma + (1-y)[(1-y)/(1-x)]^\gamma}, \quad (11)$$

where the parameter $\gamma \geq 0$ controls the degree of randomness (as γ grows the allocation tends to be more deterministic). Adopting this procedure (we refer to it as a ZH-design) the randomization function g is the same for each covariate level, so the closeness to the optimal target will be forced in the same way for each pattern. However, the evolution of the procedure (and thus the convergence properties of the allocation proportion in terms of both expectation and variability) depends on the number of subjects belonging to each pattern entering the trial, and therefore it is related to the distribution of the covariates in the population. This may be particularly critical for small samples, where some covariate profiles could be strongly under-represented so that, both from the ethical and inferential viewpoint, the need to force closeness to the target should be greater. Thus, when the covariate distri-

bution is known a-priori we suggest an extension of Zhang and Hu’s design which can change the degree of randomness between the covariate levels in order to force the convergence towards optimality on the basis of the different representativeness of the patterns in the population. Let g_1, \dots, g_J be a set of randomization functions satisfying C1) and C2), we suggest to allocate the $(k + 1)$ -th patient with $Z_{k+1} = z_j$ to A with probability

$$g_j(\pi_{kj}; \pi_{\omega j}^*(\hat{\alpha}_k)), \quad \text{for } j = 1, \dots, J. \tag{12}$$

Theorem 2. *Adopting the compound criterion (8) with a weight function $\omega < 4/5$, under the allocation rule in (12) the following holds:*

$$\lim_{n \rightarrow \infty} \boldsymbol{\pi}_n = \boldsymbol{\pi}_\omega^* \quad a.s.$$

Proof. This result can be easily derived from Zhang and Hu (2009). \square

From (11), a suitable choice consists in setting

$$g_j(x; y) = \frac{y(y/x)^{\gamma_j}}{y(y/x)^{\gamma_j} + (1-y)[(1-y)/(1-x)]^{\gamma_j}} \tag{13}$$

with $\gamma_j \propto p_j^{-1}$ for any $j = 1, \dots, J$, so the allocations for the profiles which may be potentially under-represented can be forced towards the optimal target. This procedure reflects a Covariate-adjusted Randomization and we shall call it a CR-design.

In order to perform some finite sample comparisons between our proposal and the ZH-design, we take into account a binary covariate, say gender with male (M) and female (F), and two different population scenarios where $p_M = 0.4$ and $p_M = 0.1$. The results come from 1000 simulations with $n = 40$, $\mu_B = 0$, $\sigma^2 = 1$, $\beta = 4$ and ω_s with $s = 2$. To allow for homogeneous comparisons, we adopt the ZH-design in (11) with $\gamma = 2$ and the CR-design in (13) with $\gamma_j = k/p_j$ ($j = M, F$) and $\gamma_M + \gamma_F = 4$. As shown in Table 1, the CR-design tends to force the allocation proportion to the

Table 1: Expectation and standard deviation (in brackets) of the proportion of allocations to A for male $\pi_n(M)$ and female $\pi_n(F)$.

$n = 40$	π_ω^*	$p_M = 0.4$				$p_M = 0.1$			
		ZH-design		CR-design		ZH-design		CR-design	
		$\pi_n(M)$	$\pi_n(F)$	$\pi_n(M)$	$\pi_n(F)$	$\pi_n(M)$	$\pi_n(F)$	$\pi_n(M)$	$\pi_n(F)$
$\alpha = 1$	0.512	0.5222 (0.0588)	0.5168 (0.0488)	0.5185 (0.0576)	0.5162 (0.0577)	0.5227 (0.1067)	0.5190 (0.0422)	0.5166 (0.0931)	0.5183 (0.0662)
$\alpha = 3$	0.708	0.6969 (0.0633)	0.7019 (0.0560)	0.6979 (0.0593)	0.7015 (0.0592)	0.6795 (0.1030)	0.7048 (0.0502)	0.6824 (0.0916)	0.7048 (0.0695)
$\alpha = 5$	0.847	0.8363 (0.0461)	0.8382 (0.0391)	0.8368 (0.0458)	0.8375 (0.0449)	0.7672 (0.0875)	0.8400 (0.0341)	0.7827 (0.0718)	0.8320 (0.0514)

target within the under-represented group, and this tendency grows as p_M decreases.

Moreover, the variabilities within the different profiles tend to be similar, since that of the under-represented group decreases and, at the same time, the other grows wrt the ZH-design. We have obtained similar results for $n = 100$.

Acknowledgements This research was partly supported by the Italian Ministry for Research (MIUR) PRIN 2007 “Statistical methods for learning in clinical research”.

References

- Atkinson, A.C. (2002). The comparison of designs for sequential clinical trials with covariate information. *Journal of the Royal Statistical Society, Series A* 165, 349–373.
- Atkinson, A.C. and A. Biswas (2005). Adaptive biased-coin designs for skewing the allocation proportion in clinical trials with normal responses. *Statistics in Medicine* 24, 2477–2492.
- Bandyopadhyay, U. and A. Biswas (2001). Adaptive designs for normal responses with prognostic factors. *Biometrika* 88, 409–419.
- Eisele, J. (1994). The doubly adaptive biased coin design for sequential clinical trials. *Journal of Statistical Planning and Inference* 38, 249–262.
- Geraldes, M., V. Melfi, C. Page, and H. Zhang (2006). The doubly adaptive weighted difference design. *Journal of Statistical Planning and Inference* 136, 1923–1939.
- Rosenberger, W.F., N. Stallard, A. Ivanova, C.N. Harper and M.L. Ricks (2004). Optimal adaptive designs for binary response trials. *Biometrics* 57, 909–913.
- Tymofyeyev, Y., W.F. Rosenberger and F. Hu (2007). Implementing optimal allocation in sequential binary response experiments. *Journal of the American Statistical Association* 102, 224–234.
- Zhang, L. and F. Hu (2009). A new family of covariate-adjusted response adaptive designs and their properties. *Applied Mathematical Journal of Chinese Universities* 24, 1–13.

Split-Plot and Robust Designs: Weighting and Optimization in the Multiple Response Case

Rossella Berni

Abstract This paper deals with experimental planning and optimization in response surface methodology. It aims at addressing two main issues: i) the optimization of a split-plot design in the multiple response case by the use of a robust-design approach and ii) the related problem of weighting the responses according to the actual importance of these variables and the target values when performing simultaneous optimization. An application to the study of a Numerical Control machine in order to improve the accuracy of the measurement process and to reduce the measurement time is presented.

1 Introduction

In the last decade, split-plot design, see Cochran and Cox (1957), has received great attention as a valid plan in the technological field and as a robust-design approach. In this paper, our main aim is to analyze such experimental designs from two points of view: the theoretical basis of a split-plot is first evaluated as a specific and valid experimentation for the robust-design concept, in order to estimate with accuracy the interaction terms related to noise and design (control) factors, as in Box and Jones (1992); the second point relates to the split-plot and optimization in a multiresponse case, within the Response Surface Methodology (RSM) setting, by involving just one objective function and the possibility of weighting the response variables according to their role in achieving the optimal value.

Many authors, building on the theory suggested by Derringer and Suich (1980) and Khuri and Conlon (1981), have proposed methods to synthesize and optimize responses. More recently, the multiple response case has been appreciably extended by consideration of other problems e.g. the correlation between the response vari-

Rossella Berni

Department of Statistics “G.Parenti”, Viale Morgagni 59, 50134 Florence, Italy, e-mail: berni@ds.unifi.it

ables and, above all, the consideration of noise variables. The latter play an obvious and central role in a robust design context (see, for example, Miró-Quesada and Del Castillo 2004), where variability is studied according to the noise effect and the uncertainty of estimated parameters. In Robinson, Brenneman and Myers (2006), attention is focused on robust design and categorical noise variables when the values of the control factors are evaluated jointly with cost and time issues.

We may summarize our theoretical and empirical contributions as follows:

1. In order to solve the problem of optimization in the multiple response case of RSM and in the dual approach (two surfaces which consider location and dispersion effects, we suggest a single measure (a weighted function of several variables of interest) which allows us to fit just one surface in terms of all the dependent variables;
2. We introduce the desirability approach (Derringer and Suich 1980), suitably modified in order to accommodate our weighting;
3. We present a case study in which a robust design approach is studied through a split-plot design, as in Box and Jones (1992), where we distinguish between two types of factors: the first are considered as classification devices and are included in the whole-plots; the second type are the main object of interest and are used as sub-plot factors;
4. Within the case study, we check the quality of the optimization results according to whether the response variables are standardized or not and we carry out a comparison with the standard desirability approach of Derringer and Suich (1980).

The organization of the paper is the following: in Section 2 the theory of split-plot and robust design, in an RSM approach, is briefly explained; in Section 3 the optimization procedure is outlined while Sections 4 and 5 are related to the case study. Section 4 includes the planning of the split-plot experiment while the outcome and results are reported in Section 5. Concluding remarks follow.

2 Split-Plot Theory

In general, let us define the set $C = \{C_1, \dots, C_I\}$ of Whole-Plot (W-P) factors, and the set $X = \{X_1, \dots, X_J\}$ of Sub-Plot (S-P) factors. Then, within each block k , ($k = 1, \dots, K$), for a balanced split-plot design, we have n runs, and, therefore, the total number of trials is $N = nK$. Furthermore, we define $c_i = (c_{i1}, \dots, c_{iu}, \dots, c_{in})$ as the generic vector in $[-1, +1]$, related to the i -th W-P factor, ($i = 1, \dots, I$) Further, we suppose we have a single S-P factor (X_1), $J = 1$, which is a categorical factor at L levels. Let us consider the case of a second-order split-plot model in an RSM setting including a main-effect model for the S-P factor and the interaction terms between the two sets of variables. Therefore, the split-plot model, defined for I W-P variables, one S-P categorical factor, a single block ($K=1$), a single response variable Y for each observation u , ($u = 1, \dots, n$), can be written as:

$$\begin{aligned}
y_u(C, X) = & \beta_0 + \sum_{i=1}^I \gamma_i c_{iu} + \sum_{i=2}^{I-1} \sum_{i'=i+1}^I \gamma_{i'i} c_{iu} c_{i'u} + \sum_{l=1}^{L-1} \beta_l x_{u,l} \\
& + \sum_{i=2}^I \sum_{l=1}^{L-1} \delta_{il} c_{iu} x_{u,l} + \sum_{l=1}^{L-1} \delta_{(I-1)l} c_{(I-1)u} c_{lu} x_{u,l} + \psi_{u(W-P)} + \varepsilon_{u(S-P)}.
\end{aligned} \tag{1}$$

Note that γ_i and $\gamma_{i'i}$ are parameters related to the second-order model for the W-P variables, while β_0 and β_l are intercept and main-effects for the S-P factor, denoted as $x_{u,l}$ for the generic l level. The parameters δ_{il} and $\delta_{(I-1)l}$ are related to the interaction effects between the W-P variables and the S-P factor. These last terms play a relevant role in the robust design approach, since they contain the parameters of the *control * noise* interaction effects. The error terms are represented by $\psi_{u(W-P)}$ (whole-plot error) and $\varepsilon_{u(S-P)}$ (sub-plot error). We suppose that $\psi \sim i.i.d. N(0, \sigma_\psi^2)$ and $\varepsilon \sim i.i.d. N(0, \sigma_\varepsilon^2)$; in addition, we suppose that the error terms are uncorrelated.

3 The Optimization Procedure

Consider a general response surface model, Y_t ($t = 1, \dots, T$), for T dependent variables. The simultaneous optimization is carried out for these T estimated surfaces where adjustment to a target value is to be performed. We define the following distance between the estimated surface \hat{Y}_t (considered as a function of C and X) and the target value τ_t :

$$S_t(C, X) = (\hat{Y}_t(C, X) - \tau_t)^2. \tag{2}$$

S_t can be viewed as a crude measure of variability. The aim is the minimization of the following expression on the coded experimental region jointly with the weights w_t , defined as values in $(0, 1)$:

$$\min_{[C, X]} \left\{ \sum_t w_t S_t(C, X) \right\}, \tag{3}$$

constrained by:

$$\sum_t w_t = 1. \tag{4}$$

Two kinds of weighting choices may be available: ‘‘a-priori’’ fixed weights (as in Berni 2009) or weights as parameters within the constrained optimization procedure, as in (4). Note that our aim is to find the best solution for the set of factors (C, X) ; the role of each weight in the minimization procedure is related to the term $S_t(C, X)$ in (2), i.e. the weight is defined at improving the adjustment of \hat{Y}_t to the target τ_t during the minimization procedure, carried out through the objective function of (3).

3.1 Desirability Function and Weighting

The standard desirability approach suggested by Derringer and Suich (1980) is modified, as shown below, in order to include weighting and information from the experimental data through the polynomial models estimated on the transformed values. In the literature many authors have modified this optimization method; in our case, we start by computing the transformed response values for all the N observations, $y_{u't}^*$; $u' = 1, \dots, N$; $\forall t$ as defined by Derringer and Suich (1980), then we fit the transformed surface Y_t^* . The $y_{u't}^*$ are obtained through a standardization which considers the tolerance interval of the response variable Y_t , and the corresponding target value. In addition, the $y_{u't}^* \in (0, 1) \forall u'$, i.e. the transformed value is null when the experimental point is not desirable because it exceeds the limits of the two-sided interval. The following weighted objective functions are suggested:

$$D_{M1}(C, X) = \left(\prod_t w_t \hat{Y}_t^*(C, X) \right)^{1/T} \quad \sum_t w_t = 1; t = 1, \dots, T \quad (5)$$

$$D_{M2}(C, X) = \left(\prod_t \hat{Y}_t^*(C, X)^{w_t} \right)^{1/T} \quad \sum_t w_t = T; t = 1, \dots, T. \quad (6)$$

The maximization of D_{M1} and D_{M2} is carried out by considering the estimated surfaces and the weighting. Therefore, the main differences in comparison with the standard method are:

- We optimize the values given in (5) and (6) over the experimental coded region, taking care of the relevance of the individual polynomial model for each Y_t^* ;
- Expression (5): we weight each Y_t^* surface through a weight;
- Expression (6): this is a geometric weighted mean;
- In both cases the weights may be determined a-priori or as parameters in $(0, 1)$, with the defined constraints;
- Note that in both (5) and (6), we obtain the best-fitted T surfaces through the desirability values $y_{u't}^*$ ($u' = 1, \dots, N$; $t = 1, \dots, T$) and then we carry out the maximization of formulas (5) or (6).

We point out that both functions (5) and (6) are applied for improving the optimization step and for comparing the suggested function (??).

4 An Application to a Case Study

The aim of the experiment is the improvement of the accuracy in measurements of a Numerical Control (N/C) machine and the reduction of the measuring time. The machine works by a feeler pin and it has a movable bridge framework to facilitate the positioning of the piece which must be checked. For practical purposes, reference will be made to a dental implant as the piece to be measured, but the specific nature of the piece is irrelevant. Note also that the reduction of measurement time

is implicitly the only possibility in order to reduce costs, which are a secondary problem in this case, where the risk for a patient due to measurement accuracy is the most pressing problem. The machine needs specific environment conditions for proper functioning: it has an integrated thermal compensation system which ensures proper measuring conditions and the setting of the external temperature has been solved previously (see Berni and Gonnelli 2006). The steps of the experimental planning can be outlined as:

1. The response variables are five quantitative variables Y_t , $t = 1, \dots, 5$ related to the different positioning of the feeler pin on the dental implant during the process measurement steps (targets in brackets): maximum circle diameter ($\tau_1 : 3.000$ mm), minimum circle diameter ($\tau_2 : 2.790$ mm), circle diameter measured at -3.3 mm ($\tau_3 : 2.827$ mm), outside neck circle diameter ($\tau_4 : 4.100$ mm) and eccentricity ($\tau_5 : 0.000$ mm). The Y_t are expressed in the same units of measurement, and, therefore, the choice of standardization is not compulsory. Note that each type of measurement is carried out as a distinct step; in addition, each response variable is an independent dimension of the measured piece.
2. The full measurement process includes six phases. In order to reduce the measuring time, the only step where we may intervene is the location of the frustum of a cone by 3 circles; i.e. the frustum of the cone is located by 3 circles at 3 different distances.
3. In order to locate each circle, the N/C machine software identifies a circumference by selecting several points by the feeler pin. In the initial setting, the numbers of points are set at (7,7,7) and a measuring time improvement may be achieved by reducing the number of points. We therefore introduce a factor "circle-point" (X_1) with four levels, each represented by a different combination of points in decreasing order of measuring time: (1)7, 7, 7; (2)7, 5, 7; (3)5, 7, 5; (4)5, 5, 5.
4. Two other sources of variability are included in our planning: C_2 is the measurement speed (mm/sec) for each point; C_3 (mm/sec) is the speed of the feeler pin when it is drawn onto the piece or it turns around the piece. Both factors are considered as fixed levels; their setting is chosen before beginning the measurement process.
5. Therefore, a split-plot design with 3 factors is planned; the two whole-plot factors, both at two levels, are C_2 and C_3 , while the single sub-plot factor (of greater interest) is X_1 .
6. It should be noted that the noise related to the positioning of the piece on the clamp is a source of variability out of our control. Each trial is thus performed with replicates also to evaluate this experimental error. Note that, in whole-plots, C_1 corresponds to the block effect. The final split-plot has three W-P factors ($I = 3$), a single S-P factor ($J = 1$) and 112 runs given by seven replicates (or blocks, $K = 7$), see Table 1.
7. Note that in this design, noise and control factors are studied according to a robust process in just one design matrix, where noise factors are assigned to the W-P and control factors to the S-P (as suggested in Box and Jones 1992). In this case, however, the specific block structure at the split-plot level does not

correspond to a combined array *stricto sensu*, i.e. a single array for noise and design factors without replicates.

The design ensures the equivalence of Ordinary Least Squares (OLS) and Generalized Least Squares (GLS) estimates because the conditions are satisfied (for details see Vining, Kowalski and Montgomery 2005). Observe that, in the measurement process, the noise levels, evaluated through C_2 and C_3 , are fixed before the measuring phase. At the same time, the response variables are related to the specific measurement step and each type of measurement is independent of each other.

Table 1: Split-plot design: $2^2 \times 4^1$; description of k -th block; $k = 1, \dots, K$, $K=7$ blocks

.	.			
.	.			
Block	b_k			
W-P: C_2, C_3	1,1	1,2	2,1	2,2
S-P: X_1	1,2,3,4	2,4,3,1	4,2,3,1	1,2,3,4
.	.			
.	.			

5 Optimization Results

The application, (computed with the Statistical Analysis System -SAS- software), starts by considering the results of the analysis of variance and the specific estimated surface for each response variable, as in (8). The estimates are not displayed here; in general, the significant main effects are the “circle point” (factor X_1) and the measurement speed (factor C_2); drift speed (factor C_3) shows significant p -values, except for Y_2 . The significant interaction effects are mainly the interaction $C_2 \star C_3$, except for Y_2 and Y_5 , and the interactions between X_1 and the whole-plot factors, C_2 and C_3 . The results of the statistical models obtained from standardized data are very close to the ones obtained on the basis of the original data. Note also that the first error term is formed by the interactions between the replicate effect and the W-P variables. The second error term is the residual error of the model.

The estimates of the surfaces for each Y_i are used to optimize (??); the results are displayed in Tables 2 and 3. Note that the simultaneous optimization was reached by conditioning on the specific setting of points corresponding to the levels of the “circle-point” factor X_1 . Optimization results are evaluated through the objective function value, the gradient estimates (maximum absolute gradient value) ($\|x\|_\infty$), the determinant of the Hessian matrix ($|H|$). The optimal experimental trial is also identified according to the nearness to the target values jointly with the reduction of measuring time (high level of factor C_2) and the level of factor C_3 .

In Table 2 we show the best results obtained from (2) and (3) applied to original and standardized data. We must point out that, especially for the original data, we gain good and stable results for each level of X_1 . Weights, included in the mini-

mization procedure, are well specified for the original data in comparison with the standardized ones, but however, through the transformed data the ideal target values are almost achieved. Diagnostic results are good and very steady.

Table 3 reports the optimal solution achieved through the modified desirability approach, (5) and (6). As regards (5), by considering the factors C_2 and C_3 , the optimal levels are very close to the solutions shown before, above all to the solutions of Table 2 and the standardized data. The level of factor X_1 is the same level obtained through standardized data (Table 2). The results from application of (6), shown in Table 3, are similar to previous results when considering the gained level for X_1 , level 2, with C_2 and C_3 equal to the maximum value. However, for this approach, some remarks on the diagnostic results and on the weights are highlighted. In general, the only level of X_1 which has never come out as an optimal solution is the level $X_1 = 1$, which is the initial situation. The best solution for the “circle point” factor is the level $X_1 = 4$ achieved by the original data (Table 2); undoubtedly, the level $X_1 = 2$ is also a good result, and it is achieved from applying (3) to standardized data, (5) and (6). Furthermore, a good combination for the whole-plot factors is obtained from (3) with standardized data, while only (6), related to the modified desirability approach, gains the maximum level (+1).

Table 2: Simultaneous optimization results: formula (3) on original and standardized data

Results	Formula (3); original data	Formula (3); standardized data
Optimal Soln	$X_1 = 4; C_2 = 0.710; C_3 = 0.362$	$X_1 = 2; C_2 = 0.863; C_3 = 0.812$
Weights	$w_1 = 0.285; w_2 = 1.5e - 5;$ $w_3 = 0.008 \quad w_4 = 0.651;$ $w_5 = 0.055$	$w_1 = 0.000; w_2 = 0.509;$ $w_3 = 0.490 \quad w_4 = 0.000;$ $w_5 = 0.000$
Estimated Resp. Surf. $\hat{\tau}_i$	$\hat{\tau}_1 = 3.003; \hat{\tau}_2 = 2.785;$ $\hat{\tau}_3 = 2.822; \hat{\tau}_4 = 4.110;$ $\hat{\tau}_5 = 0.001$	$\hat{\tau}_1 = 3.000; \hat{\tau}_2 = 2.788;$ $\hat{\tau}_3 = 2.822; \hat{\tau}_4 = 4.109;$ $\hat{\tau}_5 = 0.000$
Diagnostics of; $\ x\ _\infty; H $	$3.1e - 6; 1.3e-4; < 10e - 8$	$3.4e - 7; 7.5e - 4; < 10e - 8$

Table 3: The modified desirability approach; formulae (10) and (11)

Results	Formula (10)	Formula (11)
Optimal Soln	$X_1 = 2; C_2 = 0.785; C_3 = 0.986$	$X_1 = 2; C_2 = 1; C_3 = 1$
Weights	$w_1 = 0.141; w_2 = 0.334;$ $w_3 = 0.189; w_4 = 0.178;$ $w_5 = 0.158$	$w_1 = 0.001; w_2 = 0.000;$ $w_3 = 0.000; w_4 = 4.999;$ $w_5 = 0.000$
Diagnostics of; $\ x\ _\infty; H $	$4.0e - 7; 9.9e-7; < 10e - 8$	$0.1992; 5.5e-17; < 1.4e - 15$

6 Concluding Remarks

In this paper we have dealt with optimization in the multiresponse case. Weighting is a specific problem and is included in the optimization procedure through the minimization of the objective function with respect to the weights as well as the factors. The optimization methodology related to the split-plot experimental data in the multiple response case of RSM is applied in a case study. We obtain satisfactory results related to the optimizations, according to each level of the sub-plot factor “circle-point”. However, the definition of a specific factor that identifies the frustum of the cone is still critical and will require further analysis.

References

- Berni, R. and C. Gonnelli. (2006). Planning and optimization of a numerical control machine in a multiple response case. *Quality Reliability Engineering International* 22, 517–526.
- Berni, R. (2009). Response Surface Methodology and multiple response case: optimization measures, developments and comparisons. in G.I. Hayworth (ed.) *Reliability Engineering Advances*. New York: Nova Science Publishers, 287–304.
- Box, G.E.P. and S. Jones (1992). Split-plot designs for robust product experimentation. *Journal of Applied Statistics* 19, 3–26.
- Cochran, W. G. and G. M. Cox (1957). *Experimental Designs, 2nd edn*. New York, Wiley.
- Derringer, G. and R. Suich (1980). Simultaneous optimization of several response variables. *Journal of Quality Technology* 12, 214–219.
- Khuri, A.I. and M. Conlon (1981). Simultaneous optimization of multiple responses represented by polynomial regression functions. *Technometrics* 23, 363–375.
- Miró-Quesada, G. and E. Del Castillo (2004). Two approaches for improving the dual response method in robust parameter design. *Journal of Quality Technology* 36, 155–168.
- Robinson, T.J., W.A. Brennenman and W.R. Myers (2006). Process optimization via robust parameter design when categorical noise factors are present. *Quality Reliability Engineering International* 22, 307–320.
- Vining, G.G., S. Kowalski and D.C. Montgomery (2005). Response surface designs within a split-plot structure. *Journal of Quality Technology* 37, 115–129.

An Improvement in the Lack-of-Fit Optimality of the (Absolutely) Continuous Uniform Design in Respect of Exact Designs

Wolfgang Bischoff

Abstract Designs which have equal numbers of observations taken over a finite set of uniformly spaced points are most popular in practice. For optimum properties of such exact uniform designs, however, reference is often made to the lack of fit (LOF)-optimality of the corresponding absolutely continuous uniform design which is considered as the asymptotic design for the exact uniform designs. It is shown that LOF-optimality in its original form has no relation to exact uniform designs. Subsequently we give conditions under which an interpretation for exact uniform designs is possible.

1 Introduction

Many researchers in practice use designs which have equal numbers of observations taken over a finite set of uniformly spaced points. These designs are called exact uniform designs and a main reason for their popularity is their intuitive appeal and ease in implementing them in practice. For optimum properties of exact uniform designs, however, reference is made in a huge number of papers to an optimum property of the corresponding absolutely continuous uniform design. The absolutely continuous uniform design is the uniform distribution on the experimental region \mathcal{E} where \mathcal{E} is a compact interval of \mathbb{R} . It seems to be intuitive to approximate an exact uniform design by the absolutely continuous design when, as is usual, an exact design is considered as a probability measure. That optimum property of the absolutely continuous uniform design is a maxmin property which maximizes (over a class of designs) the minimal power (over a large class of alternatives) of a model check for linear regression. The maxmin result for the absolutely continuous uni-

Prof. Dr. Wolfgang Bischoff
Catholic University Eichstätt-Ingolstadt, D-85071 Eichstätt, Germany,
e-mail: wolfgang.bischoff@ku-eichstaett.de

form design was established by Wiens (1991). He investigated the F-test to check for linear regression. Biedermann and Dette (2001) additionally considered three non-parametric tests to check for linear regression. It is a surprising fact that the power of all these tests is an increasing function, at least asymptotically, of the same expression. This expression depends only on the (asymptotic) design, the regression model and the true regression function. Furthermore, Biedermann and Dette (2001) have generalized the maxmin property to any design being absolutely continuous with positive density on \mathcal{E} .

After introducing the models and notation in the next section we show that the maxmin results of Wiens (1991) and Biedermann and Dette (2001) have no relevance for exact designs. Finally, we give conditions under which an interpretation for exact designs is possible.

2 Preliminaries

We consider a non-parametric regression model

$$\mathbf{Y} = G + \varepsilon \quad (1)$$

where $\mathbf{Y} = (Y_1, \dots, Y_n)^\top$ is the vector of observations, $G = (g(t_1), \dots, g(t_n))^\top$ is the vector of evaluations of the true but unknown regression function $g : \mathcal{E} \rightarrow \mathbb{R}$ at the design points $t_1, \dots, t_n \in \mathcal{E}$ and $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)^\top$ is the vector of errors. We assume that the experimental region \mathcal{E} is a compact subset of \mathbb{R} and that $\varepsilon_1, \dots, \varepsilon_n$ are uncorrelated real random variables with expectation 0 and unknown constant variance $\sigma^2 \in (0, \infty)$. To simplify notation we put $\mathcal{E} = [0, 1]$.

We are interested in a model check for linear regression. For this purpose let $f_1, \dots, f_k : \mathcal{E} \rightarrow \mathbb{R}$ be known regression functions. Let $f = (f_1, \dots, f_k)^\top : \mathcal{E} \rightarrow \mathbb{R}^k$ be the vector of these known regression functions. Under the null-hypothesis we assume that model (1) can be written as linear regression model

$$\mathbf{Y} = X\theta + \varepsilon \quad (2)$$

where $X = X_n = (f(t_1), \dots, f(t_n))^\top$ is the design (model) matrix and $\theta \in \mathbb{R}^k$ is an unknown parameter vector.

An exact design for n observations determines n not necessarily distinct design points t_1, \dots, t_n of the experimental region \mathcal{E} . Each n -tuple $\tau_n = (t_1, \dots, t_n) \in \mathcal{E}^n$ is called an exact design for n observations. Note that an exact design $\tau_n = (t_1, \dots, t_n)$ can be considered as a probability measure ξ_n , say, by giving each design point t_i the mass $1/n$, i.e. $\xi_n = \xi_{\tau_n} = \frac{1}{n} \sum_{i=1}^n \delta_{t_i}$, where δ_t is the Dirac-measure (one-point measure) in t . In the following we do not distinguish between an exact design and its representation as a probability measure. In generalizing the above idea we regard each probability measure on \mathcal{E} (furnished with its Borel- σ -Algebra) as a (continuous) design. Let Ξ be the set of all these continuous designs (probability measures) and let Ξ_λ be the set of all absolutely continuous (with respect to Lebesgue mea-

sure) designs in Ξ . Conversely, given a design $\xi \in \Xi$ how can we realise a suitable sequence of exact designs which converges to ξ in some useful sense? For this purpose we consider the cumulative distribution function F_ξ of ξ and its quantile function Q_ξ . Then we define the exact design ξ_{n+1} for $n+1$ design points by $\xi_{n+1} = (0, Q_\xi(1/n), \dots, Q_\xi((n-1)/n), 1)$. By this construction the cumulative distribution function F_{ξ_n} converges uniformly to F_ξ . For more information, see Bischoff (1998).

The absolutely continuous uniform design λ , say, is the Lebesgue measure restricted to $\mathcal{E} = [0, 1]$. Intuitively the absolutely continuous uniform design λ is considered as the asymptotic design of the sequence of exact uniform designs $\lambda_n = (0, 1/(n-1), \dots, (n-2)/(n-1), 1)$. For our purpose this intuition is made mathematically rigorous through the result that F_{λ_n} converges uniformly to F_λ , see the discussion above.

3 LOF Optimality

Wiens (1991) proved a maxmin property for the absolutely continuous uniform design λ . To explain this property in detail let us consider the linear space $[f_1, \dots, f_k]$ spanned by the known regression functions f_1, \dots, f_k and let $L^2(\xi)$, $\xi \in \Xi$, be the set of square integrable functions with respect to ξ furnished with its canonical norm $\|\cdot\|_{L^2(\xi)}$. The null-hypothesis

$$H_0: \exists \theta = (\theta_1, \dots, \theta_k)^\top \in \mathbb{R}^k \text{ with } g = f^\top \theta, \quad (3)$$

is of primary practical interest when researchers hope to interpret their data by the linear regression model (2). A test for this problem is called a lack-of-fit-test (LOF-test). For several test statistics (Wiens (1991) for the usual F-test, Biedermann and Dette (2001) for three further tests based on nonparametric estimation of the unknown regression function g) the following fact has been proved asymptotically ($n \rightarrow \infty$):

Let $\xi \in \Xi_\lambda$ be an absolutely continuous design and let $\xi_n, n \in \mathbb{N}$, be its corresponding sequence of exact designs defined at the end of Section 2. Then the asymptotic power (the limit of the sequence of powers corresponding to $\xi_n, n \in \mathbb{N}$,) is an increasing function of

$$\|g - pr_{[f_1, \dots, f_k]}^{L^2(\xi)} g\|_{L^2(\xi)}^2 = \int_{\mathcal{E}} \left(g - pr_{[f_1, \dots, f_k]}^{L^2(\xi)} g \right) (x)^2 \xi(dx), \quad (4)$$

where $pr_{[f_1, \dots, f_k]}^{L^2(\xi)}$ is the orthogonal projector onto $[f_1, \dots, f_k]$ in $L^2(\xi)$ (see also Bischoff and Müller (2000) for a general approach).

Note that, given an exact design ξ the expression given in (4) is, up to a constant, the non-centrality parameter of the F-test for the test of LOF. Hence, given an exact design ξ the power of the F-test is an increasing function of (4). For shortness we

call the expression in (4) asymptotic power. The bigger $\|g - pr_{[f_1, \dots, f_k]}^{L^2(\xi)} g\|_{L^2(\xi)}^2$ is, the bigger is the asymptotic power. It is worth mentioning that assumptions are needed to prove that the limit of the sequence of powers (corresponding to $\xi_n, n \in \mathbb{N}$), converges to (4). For instance, Biedermann and Dette (2001) used among others the following assumption for their result:

$$g \in C^{(r)}([0, 1]), \text{ where } r \geq 2. \quad (5)$$

Since the true regression function g is unknown Wiens (1991) suggested a maxmin approach for a specific set of designs Ξ_0 and of alternatives \mathcal{F} .

Definition 1. Let $\Xi_0 \subseteq \Xi$. A design $\xi_0 \in \Xi_0$ is called lack-of-fit (LOF)-optimal in Ξ_0 for (the class of alternatives) \mathcal{F} , if

$$\max_{\xi \in \Xi_0} \inf_{h \in \mathcal{F}} \|h - pr_{[f_1, \dots, f_k]}^{L^2(\xi)} h\|_{L^2(\xi)}^2 = \inf_{h \in \mathcal{F}} \|h - pr_{[f_1, \dots, f_k]}^{L^2(\xi_0)} h\|_{L^2(\xi_0)}^2.$$

Obviously, a design solving the above problem maximizes, in the class Ξ_0 , the minimal asymptotic power over the class \mathcal{F} of alternatives. It is clear that the subset of alternatives \mathcal{F} must be separated from the hypothesis. Wiens (1991) investigated the sets

$$\mathcal{F} = \mathcal{F}_{\lambda, c} = \{h \in L^2(\lambda) \mid \|h - pr_{[f_1, \dots, f_k]}^{L^2(\lambda)} h\|_{L^2(\lambda)}^2 \geq c\}, \quad c > 0,$$

as alternatives. More generally, Biedermann and Dette (2001) considered the sets

$$\mathcal{F} = \mathcal{F}_{\xi, c} = \{h \in L^2(\xi) \mid \|h - pr_{[f_1, \dots, f_k]}^{L^2(\xi)} h\|_{L^2(\xi)}^2 \geq c\}, \quad c > 0,$$

as alternatives where $\xi \in \Xi_\lambda$ with Lebesgue-density $\frac{d\xi}{d\lambda}(t) > 0$ for all $t \in \mathcal{E}$. Here ξ gives more weight to those design points of \mathcal{E} for which a deviation is more serious. Since

$$\inf_{h \in \mathcal{F}_{\xi, c}} \|h - pr_{[f_1, \dots, f_k]}^{L^2(\xi)} h\|_{L^2(\xi)}^2 = c \inf_{h \in \mathcal{F}_{\xi, 1}} \|h - pr_{[f_1, \dots, f_k]}^{L^2(\xi)} h\|_{L^2(\xi)}^2, \quad c > 0,$$

a LOF-optimal design for $\mathcal{F}_{\xi, 1}$ is also LOF-optimal for $\mathcal{F}_{\xi, c}$ for each $c > 0$. Therefore, we can fix $c = 1$ and put

$$\mathcal{F}_\xi := \mathcal{F}_{\xi, 1} = \left\{ h \in L^2(\xi) \mid \|h - pr_{[f_1, \dots, f_k]}^{L^2(\xi)} h\|_{L^2(\xi)}^2 \geq 1 \right\}.$$

Theorem 1. a) (Wiens (1991)) The absolutely continuous uniform design λ is LOF-optimal in $\Xi_0 = \Xi_\lambda$ for $\mathcal{F} = \mathcal{F}_\lambda$.

b) (Biedermann and Dette (2001)) $\xi \in \Xi_\lambda$ with $\frac{d\xi}{d\lambda}(t) > 0, t \in \mathcal{E}$, is LOF-optimal in $\Xi_0 = \Xi_\lambda$ for $\mathcal{F} = \mathcal{F}_\xi$.

Remark 1. It is obvious that in the above theorem the class Ξ_0 of competing designs can be enlarged. For instance, all discrete designs can be added. This was already

pointed out by Wiens (1991). One can even show more: in the above theorem the class of competing designs can be enlarged to $\Xi_0 = \Xi$ and the LOF-optimal designs are unique.

Now the question arises what significance do the results of Theorem 1 have for exact uniform (Wiens) or, more generally, for exact designs corresponding to $\xi \in \Xi_\lambda$ (Biedermann and Dette)?

Wiens' result,

$$\max_{\xi \in \Xi_\lambda} \inf_{h \in \mathcal{F}_\lambda} \|h - pr_{[f_1, \dots, f_k]}^{L^2(\xi)} h\|_{L^2(\xi)}^2 = \inf_{h \in \mathcal{F}_\lambda} \|h - pr_{[f_1, \dots, f_k]}^{L^2(\lambda)} h\|_{L^2(\lambda)}^2, \quad (6)$$

is cited in many papers as an optimality property for exact uniform designs λ_n . If the result has any relevance for exact designs, then there must be a useful interpretation when we replace λ in (6) by an exact uniform design λ_n , at least if n is large enough. In the first instance let us fix the class of alternatives $\mathcal{F} = \mathcal{F}_\lambda$ as in Theorem 1. Then for an arbitrary exact uniform design λ_n with $n > k + 1$ it is the case that,

$$\inf_{h \in \mathcal{F}_\lambda} \|h - pr_{[f_1, \dots, f_k]}^{L^2(\lambda_n)} h\|_{L^2(\lambda_n)}^2 = 0. \quad (7)$$

However by the definition of \mathcal{F}_λ ,

$$\inf_{h \in \mathcal{F}_\lambda} \|h - pr_{[f_1, \dots, f_k]}^{L^2(\lambda)} h\|_{L^2(\lambda)}^2 = 1. \quad (8)$$

Equation (7) is even true if the set of alternatives is restricted to $\mathcal{F} = \mathcal{F}_\lambda \cap C[0, 1] \cap BV[0, 1]$, where $BV[0, 1]$ is the class of functions having bounded variation, being right continuous and having left hand limits. In the Appendix we construct a function $h_0 \in \mathcal{F}_\lambda \cap C[0, 1] \cap BV[0, 1]$ with

$$\|h_0 - pr_{[f_1, \dots, f_k]}^{L^2(\lambda_n)} h_0\|_{L^2(\lambda_n)}^2 = 0. \quad (9)$$

Both papers (Wiens (1991), Biedermann and Dette (2001)) consider too large a class of alternatives for LOF-optimality when we take into account the restrictions on the alternatives imposed by convergence to (4); see (5). Indeed, in both papers the problem of LOF-optimality is treated separately from convergence to the limit (4). An interpretation is possible if the class of alternatives is restricted to the alternatives imposed by convergence to (4).

Next, we consider only the F-test; the additional statistics considered by Biedermann and Dette can be similarly treated. By considering Bischoff (1998) or Bischoff and Miller (2000) and Bischoff (2002) one can discover that the class $\mathcal{F} = BV[0, 1]$ can be chosen when f_1, \dots, f_k are continuous and have bounded variation. By these papers or by direct inspection we have for $\xi \in \Xi_\lambda$ and $h \in BV[0, 1]$

$$\lim_{n \rightarrow \infty} \|h - pr_{[f_1, \dots, f_k]}^{L^2(\xi_n)} h\|_{L^2(\xi_n)}^2 = \|h - pr_{[f_1, \dots, f_k]}^{L^2(\xi)} h\|_{L^2(\xi)}^2, \quad (10)$$

where (ξ_n) is the sequence of exact designs corresponding to ξ ; see the end of Section 2. For the next result we have to restrict the class of designs. Let $\Xi_\lambda^{B[0,1]}$ be the set of all absolutely continuous designs with densities in $B[0, 1]$.

Theorem 2. *Let $f_1, \dots, f_k \in C[0, 1] \cap B[0, 1]$. Then for each $h \in BV[0, 1]$ and $\xi \in \Xi_\lambda$, the convergence given in (10) holds true. Moreover, let $\xi, \eta \in \Xi$. If*

$$\exists k+1 \text{ disjoint intervals } A_1, \dots, A_{k+1} \text{ with } \xi - \eta(A_i) > 0, \quad i = 1, \dots, k+1, \quad (11)$$

then there exists a function $h_\eta \in BV[0, 1] \cap \mathcal{F}_\xi$ with

$$\|h_\eta - pr_{[f_1, \dots, f_k]}^{L^2(\eta)} h_\eta\|_{L^2(\eta)}^2 < \|h_\eta - pr_{[f_1, \dots, f_k]}^{L^2(\xi)} h_\eta\|_{L^2(\xi)}^2 = 1. \quad (12)$$

Hence, $\xi \in \Xi_\lambda^{B[0,1]}$ is the unique LOF-optimal design in $\Xi_0 = \Xi_\lambda^{B[0,1]}$ for the class of alternatives $\mathcal{F} = BV[0, 1] \cap \mathcal{F}_\xi$.

Proof. *Proof of (12).* For designs ξ, η Wiens (see also the corresponding proof of Biedermann and Dette) constructed a measurable function $h_\eta \in \mathcal{F}_\xi$ fulfilling

$$\|h_\eta - pr_{[f_1, \dots, f_k]}^{L^2(\eta)} h_\eta\|_{L^2(\eta)}^2 \leq \|h_\eta - pr_{[f_1, \dots, f_k]}^{L^2(\xi)} h_\eta\|_{L^2(\xi)}^2 = 1. \quad (13)$$

Using (11) for the construction of h_η (see the proof of Wiens) we obtain ' $<$ ' in (13) and we recognize that then $h_\eta \in B[0, 1]$.

Proof of the last statement of the theorem. For arbitrary $\xi, \eta \in \Xi_\lambda^{B[0,1]}$, $\xi \neq \eta$, condition (11) holds true and hence (12). By (12) the assertion follows.

Remark 2. The class Ξ_0 of competing designs can be enlarged. For example, all exact designs can be added.

Remark 3. The above result does not imply a specific optimality for the uniform design. The optimality depends on the chosen class of alternatives. This was already pointed out by Biedermann and Dette (2001) for the original LOF optimality.

Remark 4. In Bischoff and Miller (2006) practically attractive designs are established by a constrained optimal design approach. In a class of LOF-efficient designs the optimal design is determined with respect to a classical design criterion.

Acknowledgment. The author thanks an unknown referee for her/his careful reading of the first draft and for her/his suggestions.

4 Appendix

Let $\mathbf{0}_m = (0, \dots, 0)^\top \in \mathbb{R}^m$, $m \in \mathbb{N}$, and let $n \in \mathbb{N}$ with $n - 1 > k$ where k is the number of known regression functions f_1, \dots, f_k . We define

$$h(x) := 1 - \cos(2\pi \cdot (n-1)x), x \in [0, 1].$$

Then it holds that $h(i/(n-1)) = 0$, $i = 0, \dots, n-1$. Let $B_i = [(i-1)/(n-1), i/(n-1)]$, $i = 1, \dots, n-1$. Then we define

$$h_0(x) := \sum_{j=1}^{n-1} w_j \cdot \mathbf{1}_{B_j}(x) h(x) / \left(\sum_{j=1}^{n-1} w_j^2 \int_{B_j} h^2 d\lambda \right)^{1/2}.$$

where $\mathbf{w} = (w_1, \dots, w_{n-1})^\top \neq \mathbf{0}_{n-1}$ is a solution of the following k equations with $n-1 > k$ variables w_1, \dots, w_{n-1} :

$$\sum_{j=1}^{n-1} w_j \int_{B_j} h(x) f(x) \lambda(dx) = \mathbf{0}_k.$$

Thus $h_0 \in C[0, 1] \cap B[0, 1]$ and we have

$$\int_{\mathcal{E}} f(x) h_0(x) \lambda(dx) = \mathbf{0}_k,$$

hence $h_0 - pr_{[f_1, \dots, f_k]}^{L^2(\lambda)} h_0 = h_0$, and $\int_{\mathcal{E}} h_0^2(x) \lambda_n(dx) = \|h_0\|_{L^2(\lambda_n)}^2 = 0$, $\int_{\mathcal{E}} h_0^2(x) \lambda(dx) = \|h_0\|_{L^2(\lambda)}^2 = 1$. Therefore $h_0 \in \mathcal{F}_\lambda$ and $\|h - pr_{[f_1, \dots, f_k]}^{L^2(\lambda_n)} h\|_{L^2(\lambda_n)}^2 = 0$.

References

- Biedermann, S. and Dette, H. (2001). Optimal designs for testing the functional form of a regression via nonparametric estimation techniques. *Statistics and Probability Letters* 52, 215–224.
- Bischoff, W. (1998). A functional central limit theorem for regression models. *Annals of Statistics* 26, 1398–1410.
- Bischoff, W. (2002). The structure of residual partial sums limit processes of linear regression models. *Theory of Stochastic Processes* 8, 23–28.
- Bischoff, W., Miller, F. (2000). Asymptotically optimal tests and optimal designs for testing the mean in regression models with applications to change-point problems. *Annals of the Institute of Statistical Mathematics* 52, 658–679.
- Bischoff, W., Miller, F. (2006). Optimal designs which are efficient for lack of fit tests. *Annals of Statistics* 34, 2015–2025.
- Wiens, D.P. (1991). Designs for approximately linear regression: Two optimality properties of uniform design. *Statistics and Probability Letters* 12, 217–221.

Optimal Allocation Proportion for a Two-Treatment Clinical Trial Having Correlated Binomial Responses

Atanu Biswas and Saumen Mandal

Abstract Optimal allocation designs for the allocation proportion are obtained in the present paper for a two-treatment clinical trial, in the presence of possible correlation between the proportion of successes for two treatments. The possibility of such correlation is motivated by real data. It is observed that the optimal allocation proportions highly depend on the correlation.

1 Introduction

Consider a clinical trial with two treatments, A and B, with binary responses. Suppose n , possibly unknown, patients are treated in the trial, of which n_A and n_B patients are to be allocated to A and B, respectively, $n = n_A + n_B$. Here n_A and n_B are also typically unknown, each patient being randomized with certain probability among the competing treatments. The optimal design in this context may be to set a target ratio $R = n_A/n_B$ for the two competing treatments. Denote the number of successes under the two treatments by Y_A and Y_B . The resulting data can be presented as a 2×2 table as follows.

	Risk factor		
	Present	Absent	Total
Outcome Favourable	Y_A	Y_B	$Y_A + Y_B$
Adverse	$n_A - Y_A$	$n_B - Y_B$	$n - Y_A - Y_B$
Total	n_A	n_B	$n = n_A + n_B$

Suppose, marginally $Y_A \sim Bin(n_A, p_A)$ and $Y_B \sim Bin(n_B, p_B)$. The conditioning argument in the table on the two marginals n_A and n_B is crucial here.

The usual assumption is independence between Y_A and Y_B , conditionally fixing n_A and n_B . The optimal allocation design can be represented by the following optimization problem:

Atanu Biswas

Applied Statistics Unit, Indian Statistical Institute, 203 B. T. Road, Kolkata – 700 108, India, e-mail: atanu@isical.ac.in

Saumen Mandal

Department of Statistics, University of Manitoba, Winnipeg, MB, R3T 2N2, Canada, e-mail: saumen_mandal@umanitoba.ca

To minimize

$$\Psi_A n_A + \Psi_B n_B \tag{1}$$

with respect to $R = n_A/n_B$, subject to $Var(\widehat{\xi}) = K$, where $\xi = \xi(p_A, p_B)$ is a parameter combination of interest, and K is a preset positive constant. Here Ψ_A and Ψ_B are suitable weights which are functions of ξ . Note $Var(\widehat{\xi}) \leq K$ yields the same design.

Case 1: If $\xi = p_A - p_B$, the *treatment difference* (Ware (1989)), then $Var(\widehat{\xi}) = \frac{p_A q_A}{n_A} + \frac{p_B q_B}{n_B}$, with $q_k = 1 - p_k$ for $k = A, B$. For general Ψ_A and Ψ_B , the optimal solution for R is

$$R_{opt} = \sqrt{\frac{p_A q_A}{p_B q_B}} \sqrt{\frac{\Psi_B}{\Psi_A}}. \tag{2}$$

If $\Psi_A = \Psi_B = 1$, (1) indicates minimization of the total sample size, and this gives the well-known Neyman allocation. With $\Psi_k = q_k, k = A, B$, (1) indicates minimization of the total expected failures subject to the constant value of the variance of the estimate of treatment difference. This problem is discussed by Rosenberger, Stallard, Ivanova, Harper, and Ricks (2001). See also Rosenberger and Lachin (2002) (pp. 174-176).

Case 2: If ξ is the Relative Risk, $RR = q_B/q_A$, we have

$$Var(\widehat{\xi}) = \frac{p_A q_B^2}{n_A q_A^3} + \frac{p_B q_B}{n_B q_A^2},$$

and with $\Psi_k = q_k, k = A, B$, the solution for R is

$$R_{opt} = \sqrt{\frac{p_A}{p_B} \frac{q_B}{q_A}}. \tag{3}$$

See Rosenberger and Lachin (2002).

Case 3: If ξ is the Odds Ratio, $OR = \frac{p_A q_B}{p_B q_A}$, we get

$$Var(\widehat{\xi}) = \frac{p_A q_B^2}{n_A q_A^3 p_A^2} + \frac{p_A^2 p_B}{n_B q_A^2 p_B^3},$$

and with $\Psi_k = q_k, k = A, B$, the solution for R is

$$R_{opt} = \sqrt{\frac{p_B}{p_A} \frac{q_B}{q_A}}. \tag{4}$$

See Rosenberger and Lachin (2002).

Almost all studies for 2×2 tables are done assuming independence of Y_A and Y_B , conditionally fixing n_A and n_B . But, in reality, given n_A and n_B , often Y_A and Y_B are correlated due to common social/economic/environmental exposures. In a clinical

trial setting, the correlation appears due to the identical nature of care available to the trial subjects irrespective of receiving treatment or placebo.

Ergin and Ergin (2005) described a meta-analysis of 11 studies (resulting in 11 2×2 tables) of thrombolytic therapy versus placebo for the treatment of acute ischemic stroke patients. The total number of patients were 3709. The data is provided by Ergin and Ergin (2005). Here the observed correlation coefficient from the 11 pairs is 0.669, which is significantly away from 0 under the assumption of independence for 2×2 tables (P -value of 0.048). This correlation can be explained by similar medical care and common environmental effect within each trial, which are treated as unknown random effects. Thrombolytic therapy was associated with an insignificant increase in mortality using the standard uncorrelated model. Here we wish to study the situation by incorporating the correlation into account.

Hwang and Biswas (2008) discussed such a correlation in the context of a single 2×2 table, while Biswas and Hwang (2009) discussed the same for multiple tables. For describing the correlated binomial distribution, they considered the bivariate binomial model of Biswas and Hwang (2002) and some extension of that. See Biswas and Hwang (2002), Biswas and Hwang (2009) and Hwang and Biswas (2008) for the bivariate models. According to the models, marginally $Y_k \sim \text{Binomial}(n_k, p_k)$, $k = A, B$, and $\text{corr}(Y_A, Y_B) = \rho$. Consequently, $\text{corr}(\hat{p}_A, \hat{p}_B) = \rho$. The correlation is due to some common random effect associated with all the observations. The model is obtained as follows: Suppose $\min(n_A, n_B) = m$, and we assume m dependent Bernoulli pairs (Y_{Ai}, Y_{Bi}) where marginally $Y_{ki} \sim \text{Bernoulli}(p_k)$. The dependent structure between Y_{Ai} and Y_{Bi} is introduced either by

$$P(Y_{Bi} = 1 | Y_{Ai}) = p_B + \theta(Y_{Ai} - p_A)$$

or

$$P(Y_{ki} = 1 | Y_{0i}) = p_k + \theta(Y_{0i} - p_0), \quad k = A, B,$$

where $Y_{0i} \sim \text{Bernoulli}(p_0)$, independent of Y_{Ai} and Y_{Bi} .

In the present paper we consider the situation of bivariate binomial models where the correlation ρ is present. We obtain an optimal allocation proportion R_{opt} in the presence of such correlation. This is obtained in Section 2. In Section 3 we provide some numerical computational results. Finally Section 4 concludes.

2 Optimal Allocation Proportions in the Presence of Correlation

Under the correlated set up,

$$\text{Var}(\hat{p}_A - \hat{p}_B) = \frac{p_A q_A}{n_A} + \frac{p_B q_B}{n_B} - 2\rho \frac{\sqrt{p_A q_A p_B q_B}}{\sqrt{n_A n_B}}. \quad (5)$$

Now, consider the case of odds ratio, OR. Writing $\mu_A = n_A p_A$, $k = A, B$, we expand $\log\left(\frac{Y_A}{n_A - Y_A}\right)$ by retaining terms up to the first order. We get $\log \widehat{OR}$ as

$$\begin{aligned} \log \widehat{OR} &\simeq \log \left(\frac{\mu_A}{n_A - \mu_A} \right) + \frac{n_A}{\mu_A(n_A - \mu_A)} (Y_A - \mu_A) \\ &\quad - \log \left(\frac{\mu_B}{n_B - \mu_B} \right) - \frac{n_B}{\mu_B(n_B - \mu_B)} (Y_B - \mu_B). \end{aligned}$$

Hwang and Biswas (2008) showed that

$$Var(\log \widehat{OR}) \simeq \frac{1}{\mu_A} + \frac{1}{n_A - \mu_A} + \frac{1}{\mu_B} + \frac{1}{n_B - \mu_B} - f_\rho,$$

where

$$f_\rho = 2\rho \sqrt{\left(\frac{1}{\mu_A} + \frac{1}{n_A - \mu_A} \right) \left(\frac{1}{\mu_B} + \frac{1}{n_B - \mu_B} \right)},$$

where ρ is the correlation coefficient between Y_A and Y_B or between \widehat{p}_A and \widehat{p}_B .

So,

$$Var(\log \widehat{OR}) = \frac{1}{n_A p_A q_A} + \frac{1}{n_B p_B q_B} - 2\rho \sqrt{\frac{1}{n_A n_B p_A q_A p_B q_B}},$$

and consequently, under the presence of correlation,

$$Var(\widehat{OR}) = \frac{p_A^2 q_B^2}{p_B^2 q_A^2} Var(\log \widehat{OR}) = \frac{p_A q_B^2}{n_A p_B^2 q_A^3} + \frac{p_A^2 q_B}{n_B q_A^2 p_B^3} - \frac{2\rho p_A^{3/2} q_B^{3/2}}{\sqrt{n_A n_B p_B^{5/2} q_A^{5/2}}}. \quad (6)$$

In a similar fashion, under correlation,

$$Var(\widehat{RR}) = \frac{p_A q_B^2}{n_A q_A^3} + \frac{p_B q_B}{n_B q_A^2} - \frac{2\rho \sqrt{p_A p_B q_B^3}}{\sqrt{n_A n_B q_A^{5/2}}}. \quad (7)$$

It is interesting to note that all the correlation-adjusted variance expressions, (5), (6) and (7), are of the form

$$\frac{\sigma_A^2}{n_A} + \frac{\sigma_B^2}{n_B} - \frac{2\sigma_{AB}}{\sqrt{n_A n_B}}$$

for some choice of σ_A^2 , σ_B^2 and σ_{AB} .

Thus, the problem at hand is to minimize (1) subject to

$$\frac{\sigma_A^2}{n_A} + \frac{\sigma_B^2}{n_B} - \frac{2\sigma_{AB}}{\sqrt{n_A n_B}} = K. \quad (8)$$

Writing $n_A = nR/(1+R)$ and $n_B = n/(1+R)$ in (8), we have

$$nK = \frac{\sigma_A^2(1+R)}{R} + \sigma_B^2(1+R) - \frac{2\sigma_{AB}(1+R)}{\sqrt{R}}. \quad (9)$$

Using (9) in (1), we get

$$\frac{(\Psi_A R + \Psi_B)(\sigma_A^2 + \sigma_B^2 R - 2\sigma_{AB}\sqrt{R})}{RK},$$

which is to be minimized. Thus, essentially we have to minimize

$$\Psi_A \sigma_B^2 R - 2\Psi_A \sigma_{AB} \sqrt{R} - 2\Psi_B \sigma_{AB} R^{-1/2} + \Psi_B \sigma_A^2 R^{-1} + (\Psi_A \sigma_A^2 + \Psi_B \sigma_B^2).$$

Differentiating with respect to R and equating to zero, we have to solve

$$\Psi_A \sigma_B^2 R^2 - \Psi_A \sigma_{AB} R^{3/2} + \Psi_B \sigma_{AB} R^{1/2} - \Psi_B \sigma_A^2 = 0 \quad (10)$$

to find a solution for minimum R , denoted by R_{opt} . If $\rho = 0$, the equation (10) has two real roots of equal magnitude but of different signs. The positive root is the solution R_{opt} . For $\rho > 0$, (10) has three roots, two being complex and the other is positive real. The real root is the solution R_{opt} . Numerical computations are given in the next Section.

3 Numerical Computations

We carry out a detailed numerical study to find the optimal proportion R_{opt} in different situations. Part of our numerical results are provided in Tables 1-4. We provide results with (5) as the constraint in Tables 1-2, and (6) as the constraint in Tables 3-4. We consider four choices of ρ , namely $\rho = 0, 0.3, 0.6$ and 0.9 , $\rho = 0$ being the case under independence. The existing results in Section 1 correspond to $\rho = 0$. We always consider $p_B \geq p_A$ in our computations.

In Table 1, we consider $\Psi_k = q_k$, $k = A, B$. We observe that when $p_A = p_B$, $R_{opt} = 1$ always. This does not depend on ρ as the roles of A and B are exchangeable. However, for $p_A \neq p_B$, the optimal allocation proportion R_{opt} depends on ρ to quite an extent. When both p_A and p_B are smaller than 0.5, R_{opt} decreases as ρ increases. But, if either of p_A and p_B is greater than 0.5, R_{opt} is an increasing function of ρ . R_{opt} is so much influenced by ρ that it is possible that R_{opt} is less than 1 for some ρ , and it is greater than 1 for some other value of ρ . In Table 1, the row corresponding to $(p_A, p_B) = (0.4, 0.8)$ exhibits such a pattern.

Table 2 corresponds to Neyman allocation (where $\Psi_A = \Psi_B = 1$). Here also $R_{opt} = 1$ when $p_A = p_B$, irrespective of the value of ρ . In addition, here $R_{opt} = 1$ and independent of ρ , whenever $p_A = p_B$. In fact, here R_{opt} depends on ρ otherwise, but R_{opt} is the same for $(p_A, p_B) = (a, b), (a, 1-b), (1-a, b), (1-a, 1-b)$. Here R_{opt} is decreasing in ρ when $p_A < p_B$ and $p_A < 0.5$; however R_{opt} is increasing in ρ when $0.5 < p_A < p_B$.

In Tables 3 and 4 we present the results where the constraint corresponds to the variance of the estimator of the Odds Ratio. Table 3 is for $\Psi_k = q_k$, $k = A, B$, while Table 4 corresponds to $\Psi_A = \Psi_B = 1$. From Table 3, we observe that $R_{opt} = 1$ and does not depend on ρ when $p_A = p_B$. Otherwise R_{opt} is always an increasing function of ρ whenever $p_A < p_B$. However, even for $p_A < p_B$, R_{opt} can be less than

Table 1: Optimal allocation proportion (R_{opt}) for different $p_A < p_B$ and ρ corresponding to $\Psi_k = q_k$, $k = A, B$, and $Var(\hat{p}_A - \hat{p}_B) = K$ as the constraint.

p_A	p_B	$\rho = 0$	$\rho = 0.3$	$\rho = 0.6$	$\rho = 0.9$
0.2	0.2	1.000	1.000	1.000	1.000
0.2	0.4	0.707	0.700	0.689	0.674
0.2	0.6	0.577	0.592	0.614	0.650
0.2	0.8	0.500	0.567	0.680	0.892
0.4	0.4	1.000	1.000	1.000	1.000
0.4	0.6	0.816	0.846	0.891	0.964
0.4	0.8	0.707	0.811	0.990	1.328
0.6	0.6	1.000	1.000	1.000	1.000
0.6	0.8	0.866	0.956	1.102	1.366
0.8	0.8	1.000	1.000	1.000	1.000

Table 2: Optimal allocation proportion (R_{opt}) for different $p_A < p_B$ and ρ corresponding to $\Psi_A = \Psi_B = 1$, and $Var(\hat{p}_A - \hat{p}_B) = K$ as the constraint.

p_A	p_B	$\rho = 0$	$\rho = 0.3$	$\rho = 0.6$	$\rho = 0.9$
0.2	0.2	1.000	1.000	1.000	1.000
0.2	0.4	0.816	0.788	0.748	0.691
0.2	0.6	0.816	0.788	0.748	0.691
0.2	0.8	1.000	1.000	1.000	1.000
0.4	0.4	1.000	1.000	1.000	1.000
0.4	0.6	1.000	1.000	1.000	1.000
0.4	0.8	1.225	1.269	1.336	1.446
0.6	0.6	1.000	1.000	1.000	1.000
0.6	0.8	1.225	1.269	1.336	1.446
0.8	0.8	1.000	1.000	1.000	1.000

1 or more than 1 depending on the values of (p_A, p_B) . The computational results for $\Psi_A = \Psi_B = 1$ are given in Table 4. Here $R_{opt} = 1$ (and free of ρ) if $p_A = p_B$ or $p_A = q_B$. Also R_{opt} is increasing in ρ for small values of $p_A (< p_B)$, and is decreasing in ρ otherwise. Also, when $p_A + p_B = 1$, it is easy to check that both the constraints corresponding to the variance expressions (5) and (6) become the same, and will yield identical R_{opt} -values. This is also confirmed by our computations of Tables 1 and 3, and 2 and 4, corresponding to the rows where $p_A + p_B = 1$.

We also carried out numerical computations for the constraint corresponding to (7), which is for the variance of the estimator of the Risk Ratio, for $\Psi_k = q_k$, $k = A, B$, and also for $\Psi_A = \Psi_B = 1$. We do not provide tables given space limitations. However, here $R_{opt} = 1$ (and independent of ρ) only when $p_A = p_B$. Otherwise, for $p_A < p_B$, we observe R_{opt} always decreases as ρ increases. Hence, for $p_A < p_B$, R_{opt} is always less than 1. We observe an exactly similar feature for both choices of Ψ_k .

Table 3: Optimal allocation proportion (R_{opt}) for different $p_A < p_B$ and ρ corresponding to $\Psi_k = q_k$, $k = A, B$, and $Var(\widehat{OR}) = K$ as the constraint.

p_A	p_B	$\rho = 0$	$\rho = 0.3$	$\rho = 0.6$	$\rho = 0.9$
0.2	0.2	1.000	1.000	1.000	1.000
0.2	0.4	1.061	1.128	1.232	1.411
0.2	0.6	2.739	3.018	3.373	3.833
0.2	0.8	0.500	0.567	0.680	0.892
0.4	0.4	1.000	1.000	1.000	1.000
0.4	0.6	0.816	0.846	0.891	0.964
0.4	0.8	0.471	0.501	0.548	0.627
0.6	0.6	1.000	1.000	1.000	1.000
0.6	0.8	0.577	0.592	0.614	0.650
0.8	0.8	1.000	1.000	1.000	1.000

Table 4: Optimal allocation proportion (R_{opt}) for different $p_A < p_B$ and ρ corresponding to $\Psi_A = \Psi_B = 1$, and $Var(\widehat{OR}) = K$ as the constraint.

p_A	p_B	$\rho = 0$	$\rho = 0.3$	$\rho = 0.6$	$\rho = 0.9$
0.2	0.2	1.000	1.000	1.000	1.000
0.2	0.4	1.225	1.269	1.336	1.446
0.2	0.6	1.225	1.269	1.336	1.446
0.2	0.8	1.000	1.000	1.000	1.000
0.4	0.4	1.000	1.000	1.000	1.000
0.4	0.6	1.000	1.000	1.000	1.000
0.4	0.8	0.816	0.788	0.748	0.691
0.6	0.6	1.000	1.000	1.000	1.000
0.6	0.8	0.816	0.788	0.748	0.691
0.8	0.8	1.000	1.000	1.000	1.000

4 Concluding Remarks

In the present paper, we provide optimal design (allocation proportions) for a two-treatment trial, where we minimize some suitable objective function subject to some standard constraints. The focus of the present paper is to obtain an optimal design under a possible correlation between the two treatment responses, possibly due to some random effect. The optimal allocation proportion, R_{opt} , depends on the correlation ρ , and hence it is unwise to ignore such correlations.

The optimal design R_{opt} , like a standard nonlinear optimization problem, depends on the parameters p_A , p_B and ρ , which are typically unknown. One may work with some prior idea about these parameters, or use sequentially updated estimates of the parameters, as is done in adaptive designs (Rosenberger, Stallard, Ivanova, Harper, and Ricks 2001; Biswas and Mandal 2004; Biswas and Mandal 2007).

There are, though, issues with respect to estimation of ρ . If a prior estimate of ρ is available from past study, that can be used. If no past data can be available, and the trial yields only one pair of success counts, Y_A and Y_B , an estimate of ρ cannot be

determined. However, as in the example of multiple tables obtained in meta-analysis or some specially structured single table where the data can be partially partitioned, estimates of such correlations are obtained (see Biswas and Hwang 2009 and Hwang and Biswas 2008). This is possible for matched case-control studies. The estimation of ρ is possible for a multi-stage study, where the earlier stage data can be used to estimate ρ and also update that estimate. This is, of course a sequential approach.

However, details on the estimation is beyond the scope of the present paper. The optimal designs can be singular, ruling out estimation of some parameters. Here no one trial, optimal or not, can facilitate estimation of ρ . This in turn means that updated estimates of the variances of the estimates of D , RR , OR are not available.

Further extensions of the problem may be in the presence of covariates, multi-treatments, or for the case of continuous responses. These are topics of future study.

Acknowledgements The authors wish to thank two anonymous referees for their careful reading and constructive suggestions which led to some improvement over an earlier version of the manuscript. The research of S. Mandal is supported by a Discovery Grant from the Natural Sciences and Engineering Research Council (NSERC) of Canada.

References

- Biswas, A. and J.-S. Hwang (2002). A new bivariate binomial distribution. *Statistics and Probability Letters* 60, 231–240.
- Biswas, A. and J.-S. Hwang (2009). Distribution of odds ratio in 2×2 contingency tables: adjustment for correlation. (Submitted).
- Biswas, A. and S. Mandal (2004). Optimal adaptive designs in phase III clinical trials for continuous responses with covariates. In *mODa 7 - Advances in Model-Oriented Design and Analysis*. Heidelberg: Physica-Verlag, 51–59.
- Biswas, A. and S. Mandal (2007). Optimal three-treatment response-adaptive designs for phase III clinical trials with binary responses. In *mODa 8 - Advances in Model-Oriented Design and Analysis*, Heidelberg: Physica-Verlag, 33–40.
- Ergin, A. and N. Ergin (2005). Is thrombolytic therapy associated with increased mortality? *Archives of Neurology* 62, 362–366.
- Hwang, J.-S. and A. Biswas (2008). Odds ratio for a single 2×2 table with correlated binomial for two margins. *Statistical Methods and Applications* 17, 483–497.
- Rosenberger, W. and J. Lachin (2002). *Randomization in Clinical Trials: Theory and Practice*. New York: Wiley.
- Rosenberger, W., N. Stallard, A. Ivanova, C. Harper, and M. Ricks (2001). Optimal adaptive designs for binary response trials. *Biometrics* 57, 909–913.
- Ware, J. (1989). Investigating therapies of potentially great benefit: ECMO. *Statistical Science* 4, 298–306.

Sample Size Determination for Multivariate Performance Analysis with Complex Designs

Stefano Bonnini, Livio Corain and Luigi Salmaso

Abstract The literature of multiple comparison methods addresses the problem of ranking treatment groups from best to worst. However, there is no clear indication of how to deal with the information from pairwise multiple comparisons, particularly in the case of blocking (or stratification) or in the case of multivariate response variables. In the present paper we take three methods into consideration to produce a performance ranking of C treatments under study. By means of a simulation study, it is possible to calculate the percentages of correct classifications of the compared methods and study their performances. The proposed simulation study also allows us to determine the minimum sample size useful for detecting performance differences among treatments.

1 Introduction

In experimental design and analysis of variance the topic of defining a treatment ranking from a multivariate point of view seems to be quite recent: it has been addressed for the first time by Bonnini et al. (2009). The literature of multiple comparison methods addresses the problem of ranking treatment groups from best to worst. However there is no clear indications of how to deal with the information from pairwise multiple comparisons, particularly in the case of blocking (or stratification) or in the case of multivariate response variables.

Dr. S. Bonnini

Department of Mathematics, University of Ferrara, Italy, e-mail: stefano.bonnini@unife.it

Dr. L. Corain

Department of Management and Engineering, University of Padova, Italy, e-mail: livio.corain@unipd.it

Prof. L. Salmaso

Department of Management and Engineering, University of Padova, Italy, e-mail: salmaso@gest.unipd.it

In industrial research a global ranking in terms of performance of all investigated products/prototypes is very often a natural goal. As proof, in 2008 an international industrial organization called AISE formally incorporated such a method as the official standard for industrial research on house cleaning products (see www.aise.eu/).

Let Y be the multivariate numeric variable related to the response of any experiment of interest and let us assume, without loss of generality, that high values of each Y univariate element correspond to better performance and therefore to a higher degree of treatment preference. The experimental design of interest is defined by the comparison of C groups or treatments with respect to S different variables, where n replications of a single experiment are performed by a random assignment of a statistical unit to a given group. The C -group multivariate statistical model (with fixed effects) can be represented as follows:

$$Y_{ijk} = \mu_{ij} + \varepsilon_{ijk}, \quad \varepsilon_{ijk} \sim \text{IID}(0, \sigma_{ij}^2), \quad i = 1, \dots, C; j = 1, \dots, S; k = 1, \dots, n; \quad (1)$$

where, in the case of a balanced design, n is equal to the number of replications and indices i and j label groups (treatments) and univariate response variables respectively.

The resulting inferential problem of interest is concerned with a set of S hypothesis testing procedures $H_{0j} : \mu_{1j} = \mu_{2j} = \dots = \mu_{Cj}$ vs. $H_{1j} : \bar{H}_{0j}$. If H_{0j} is rejected, a further possible set of $C \times (C - 1)/2$ all pairwise comparisons are performed:

$$\begin{cases} H_{0j(ih)} : \mu_{ij} = \mu_{hj} \\ H_{1j(ih)} : H_{0j(ih)} \text{ is not true} \end{cases} .$$

In the framework of parametric methods, when assuming the hypothesis of normality for random error components, the inferential problem can be solved by means of the ANOVA F -test and a further set of pairwise tests using Fisher's LSD or Tukey procedures, which are two of the most popular multiple comparison procedures. On the basis of inferential results achieved at the univariate C -group comparison stage, the next step consists in producing a ranking of the treatments in terms of performance.

2 Global Ranking Methods

All the methods we take into consideration in this paper consist in producing a ranking of the C treatments under study according to the following steps:

1. The starting point is the result of the multiple comparisons analysis, i.e. $S C \times C$ matrices of p -values;
2. For each response variable a suitable score matrix is then defined;
3. Through a synthesis procedure (sum, mean or some combination function) the scores are synthesized into $S C$ -dimensional score vectors;

4. The set of S score vectors is finally synthesized to perform one global score vector;
5. The rank of this final global score provides the multivariate global ranking of treatments, where rank 1 corresponds to the worst treatment/product/group and rank C corresponds to the best one (increasing rank).

According to the AISE Method, we can define a set of $S \times C \times C$ score matrices, where each element $x_{j(ih)}$ is related to the comparison between the treatments i and h for the j -th response variable, as described in Bonnini et al. (2006) and Corain and Salmaso (2007). Formally,

$$\begin{cases} \text{if } H_{0j(ih)} : \mu_{ij} = \mu_{hj} \text{ is not rejected, then } x_{j(ih)} = x_{j(hi)} = 0; \\ \text{if } H_{0j(ih)} : \mu_{ij} = \mu_{hj} \text{ is rejected, then } x_{j(ih)} = \text{sgn}(\bar{y}_{ij} - \bar{y}_{hj}) = -x_{j(hi)}; \end{cases} \quad (2)$$

where \bar{y}_{ij} and \bar{y}_{hj} , $i, h = 1, \dots, C$, $i \neq h$, are the sample means of groups i and h for response variable Y_j , $j = 1, \dots, S$. Note that pairwise comparisons and the valid score assignments are performed only when the C -sample test has rejected the null hypothesis H_{0j} , $j = 1, \dots, S$. In order to obtain a final global score, following the AISE procedure, an additive function is applied:

$${}_A x_i = \sum_{j=1}^S \sum_{h=1}^C x_{j(ih)}, i = 1, \dots, C. \quad (3)$$

The global ranking obtained following the AISE Method is ${}_A R_i = \text{Rank}({}_A x_i)$, $i = 1, \dots, C$.

Instead of summing ± 1 scores, a function of the p -values of one-sided tests of pairwise comparisons could be used. Let us use $p_{j(ih)}$ to indicate the p -value of the two-sample test where the null hypothesis is $H_{0j(ih)}$ and the alternative is $H_{1j(ih)} : \mu_{ij} > \mu_{hj}$. The final global score, according to the NPC Method, can be calculated as follows:

$${}_N x_i = \sum_{j=1}^S \psi(p_{j(i1)}, \dots, p_{j(iC)}), \quad (4)$$

$i = 1, \dots, C$, where $\psi(\cdot)$ is a suitable combining function satisfying some weak properties described in Lago and Pesarin (2000). An example of a combining function is Fisher's function:

$$\psi(p_{j(i1)}, \dots, p_{j(iC)}) = \sum_{h=1, h \neq i}^C \log(p_{j(ih)}). \quad (5)$$

It is worth noting that this combining method is non-parametric with respect to the underlying dependence structure among p -values from different univariate response variables, in that many kinds of monotonic relations are implicitly captured. Indeed, no explicit model for this dependence structure is needed and no dependence parameters (i.e. covariances) have to be estimated directly from the data. The global ranking obtained according to the NPC Method is ${}_N R_i = \text{Rank}({}_N x_i)$, $i = 1, \dots, C$.

Another interesting method to calculate a global score is the so-called GPS Method. Let us use $\bar{y}_{(1)j} \geq \bar{y}_{(2)j} \geq \dots \geq \bar{y}_{(C)j}$ to indicate the ordered observed sample means for Y_j and assume that higher values correspond to better performances. The algorithm to calculate the GPS score is the following:

1. For each of the S variables a $C \times C$ matrix W is created where the elements under the main diagonal are null and those over the main diagonal take value 0 or 1 according to the following rule:

$$w_{j(uv)} = g \left[\bar{y}_{(u)j}, \bar{y}_{(v)j} \right] = \begin{cases} 1 & \text{if } \bar{y}_{(u)j} \text{ is significantly not equal to } \bar{y}_{(v)j} \\ 0 & \text{otherwise;} \end{cases} \quad (6)$$

2. A rank table is created, where each column corresponds to a treatment and treatments are ordered according to the values of the sample means, according to the following steps:
 - a. In row 1, rank C is assigned to the treatment with the highest mean (first column), indicated by (1), and to all the other products with mean performances which are not significantly different from that of (1);
 - b. In row 2, rank $C - 1$ is assigned to the treatment with the highest mean from among those excluded from rank C assignation, and to all the other products with mean performances which are not significantly different from that of (2);
 - c. In row r , rank $C - r + 1$ is assigned to the treatment with the highest mean from among those excluded from rank $C - r$ assignation, and to all the other products with mean performances which are not significantly different from that of (r);
 - d. The iterated procedure stops when a rank is assigned to the product (C);
3. For each treatment, the arithmetic mean of the values from the rank table (mean by columns) gives a partial performance score: $w_{j(i)}$;
4. In order to obtain the final global GPS score, the partial scores are summed:

$${}_G x_i = \sum_{j=1}^S w_{j(i)}, \quad (7)$$

$i = 1, \dots, C$, and the global combined ranking is obtained through the usual rank transformation: ${}_G R_i = \text{Rank}({}_G x_i)$, $i = 1, \dots, C$.

3 Simulation Study and Sample Size Determination

Let us consider a simulation study based on real case data to compare the performances of the three described global ranking methods. The treatments of interest are 4 dosages (P1: 100%, P2: 95%, P3: 90%, P4: 85%) of a given detergent ($C = 4$). A priori, we know the true ranking: $P1 \succ P2 \succ P3 \succ P4$. Detergent performances are

assessed by measuring the percentages of removed stain (so-called reflectance) from a piece of fabric, previously soiled with 25 different stains ($S = 25$). The sample size for these experiments is usually equal to four or five (AISE, 2009), but since we are interested in the determination of a suitable sample size the experiment has been replicated 24 times. Indeed, it is possible to randomly extract samples of different sizes (i.e. 4,8,12,16,20) and to perform a Monte Carlo simulation study using data randomly extracted from a huge real dataset. Then for each simulation, consisting of a random selection of n experimental data (i.e. random selection of n replications among the 24 observed) the ranking is established according to the score, calculated as described in section 2 and the rates of correct classification can be calculated, marginally for each product and jointly for all the treatments.

The advantage of this simulation approach is that we do not need to set the unknown values of parameters (e.g. means and variance/covariance of model 1) and results will be directly related to the real experimental framework. The stains can be classified by their degree of importance (discrimination capability: 1, 2 or 3) and by their main chemical properties (Bleachable, Enzymatic, General detergency).

Using *R* software, 1000 data-sets were simulated, and, for each of the three methods, the percentages of correct global rankings were calculated for all the products jointly considered and for each product. For $n = 4$ the percentage of correct joint classification, i.e. the percentage of estimated joint rankings exactly equal to the true one, for the NPC Method is 40.7%, while the performances of the AISE and GPS Methods are definitely worse (6.2% and 6.1%). Anyway in this specific study we are mainly interested in the correct classification rates of one product, in particular the best one, so the marginal rates are more informative.

In Table 1, it is possible to see the results of these simulations in the case of sample size $n = 4$. Each row corresponds to the true rank of a given treatment, and the columns correspond to the “empirical” ranks obtained applying the three methods. Hence for each method, frequencies of a given row correspond to the marginal frequencies of classification for the related treatment and they sum to the number of simulations (1000). Even some columns sum to 1000 but this is always true only for the NPC method. This is due to ties, which are more frequent in the AISE and GPS Methods, because of the loss of information using ± 1 scores, according to the statistical significance of the tests, instead of p -values for each pair-wise comparison. The fact that the AISE and GPS rankings tend to overclassify to rank 1 depends on the ranking assignment rules in the presence of ties: first of all rank 1 is assigned to the worst treatment (more than one in the case of ties); rank 2 is assigned to the worst of the remaining treatments; etc.

It is evident that the NPC Method is the best procedure because the percentages of correct classifications are higher than those of the other two methods.

Table 2 shows the 2.5-th and 97.5-th percentiles ($q_{0.025}$ and $q_{0.975}$ respectively) of the NPC global index for each treatment as a function of the sample size. As expected, the ranges of the pseudo-confidence intervals (i.e. the intervals from $q_{0.025}$ to $q_{0.975}$) tend to decrease when the sample sizes increase. Suppose that the difference between the upper limit of the interval for a given treatment and the lower limit of the interval for a better treatment is negative, it can be considered empirical evidence

Table 1: Counts and percentages of right and wrong classifications (ranks) of the four treatments.

true ranking	NPC				AISE				GPS			
	4	3	2	1	4	3	2	1	4	3	2	1
4	614	313	73	0	327	164	160	349	484	117	82	317
3	348	469	155	28	225	238	164	373	32	258	360	350
2	38	213	681	68	3	66	476	455	24	11	298	667
1	0	5	91	904	2	2	65	931	9	51	179	761
% right ranking	61%	47%	68%	90%	33%	24%	48%	93%	48%	26%	30%	76%

that the former is worse than the latter. In this way it is possible to deduce the sample size which allows us to detect differences in the performance of the treatments. For example, in Table 2 it is possible to see that from $n = 12$ the NPC Method can detect the difference between the performances of $P3$ and $P4$. The differences between the performances of $P2$ and $P3$ start to be significant when $n = 20$. Hence less information is needed to detect differences between the two best treatments. The reason is due to the nonlinear relation between the dosage and the treatment effect, i.e. the difference between the main effects of $P3$ and $P4$ is less than the difference between the main effects of $P2$ and $P3$ which is less than the difference between the main effects of $P1$ and $P2$ as well.

Table 2: Percentiles and pseudo confidence intervals of the distribution of the NPC score.

sample size (n)	P1		P2		P3		P4	
	$q_{0.025}$	$q_{0.975}$	$q_{0.025}$	$q_{0.975}$	$q_{0.025}$	$q_{0.975}$	$q_{0.025}$	$q_{0.975}$
4	3.9	37.3	6.1	60.3	17.5	70.4	43.5	94.2
8	5.1	24.6	7.5	42.3	21.8	55.7	48.3	82.6
12	6.1	19.1	9.9	32.7	24.1	47.2	51.5	77.3
16	7.0	16.1	11.2	27.6	26.6	42.8	54.6	73.1
20	8.4	14.2	13.3	23.5	29.3	39.5	57.0	68.7
	$q_{0.975} - q_{0.025}$		$q_{0.975} - q_{0.025}$		$q_{0.975} - q_{0.025}$		$q_{0.975} - q_{0.025}$	
4	31.2		42.8		26.8			
8	17.1		20.5		7.5			
12	9.2		8.7		- 4.4			
16	4.9		1.0		- 11.8			
20	0.9		- 5.8		- 17.5			

4 Conclusions

AISE, NPC and GPS Methods are reliable tools to obtain a ranking of treatments within an experimental design framework. The NPC Method may be preferable because the percentages of correct classifications are higher than those of the other

methods. The proposed simulation study makes it possible to determine the minimum sample size useful for detecting performance differences among treatments and to design an experiment from the point of view of number of replications, when it is possible to have data from a preliminary experiment on the same treatments under study.

Acknowledgements The authors wish to thank the University of Padova (CPDA088513/08 and CPDA092350/09) and the Italian Ministry for University and Research (2008WKHJPK.002) for providing the financial support for this research.

References

- AISE (2009). A.I.S.E. detergent test protocol. *WWW page, URL <http://www.aise.eu/aisedetergenttestprotocol2009/>*.
- Bonnini, S., L. Corain, A. Cordellina, A. Crestana, R. Musci, and L. Salmaso (2009). A novel global performance score with application to the evaluation of new detergents. In *Statistical Methods for the Evaluation of Educational Services and Quality of Products* (M. Bini, P. Montanari, D. Piccolo and L. Salmaso eds). Heidelberg: Physica-Verlag, (to appear).
- Bonnini, S., L. Corain, and L. Salmaso (2006). A new statistical procedure to support industrial research into new product development. *Quality and Reliability Engineering International* 22, 555–566.
- Corain, L., L. Salmaso (2007). A nonparametric method for defining a global preference ranking of industrial products. *Journal of Applied Statistics* 34, 203–216.
- Lago, A., F. Pesarin (2000). Non parametric combination of dependent rankings with application to the quality assessment of industrial products. *Metron* 58, 39–52.

Optimal Design for Compositional Data

Roelof L. J. Coetzer and Linda M. Haines

Abstract In this paper optimal designs for experiments involving compositional data, specifically locally D -optimal designs for the additive logistic normal model and locally D_S -optimal designs for Dirichlet regression, are investigated. The theory underpinning the construction of these designs is based on the appropriate information matrices and the development, while new, is relatively straightforward. The ideas are illustrated by means of a simple example, that of two consecutive reactions.

1 Introduction

Models for compositional data, that is data comprising proportions, and more specifically models based on the additive logistic transformation developed by Aitchison (2004), are well researched and widely used. In contrast however there appears to have been no attention given in the literature to the design of experiments for compositional data, with the exception of a small comment in Atkinson, Donev and Tobias (2007). The present study therefore represents a first attempt to investigate the construction of optimal designs for proportions. Specifically two models together with the attendant designs are considered, the additive logistic normal in Section 2 and the Dirichlet in Section 3. The ideas developed in these sections are illustrated by means of a simple example, that of two consecutive reactions, in Section 4 and some broad conclusions and pointers for future research are given in Section 5.

Roelof L. J. Coetzer

Sasol Technology, PO Box 1, Sasolburg 1947, South Africa, e-mail: roelof.coetzer@sasol.co.za

Linda M. Haines

Department of Statistical Sciences, University of Cape Town, Rondebosch 7700, South Africa, e-mail: linda.haines@uct.ac.za

2 Additive Logistic Normal

2.1 Model

Consider the proportions $y = (y_1, \dots, y_D)$ constrained in the usual way so that $0 \leq y_j \leq 1, j = 1, \dots, D$, and $\sum_{j=1}^D y_j = 1$. Then, following Aitchison (2004), an immediate approach to modelling such proportions is to introduce the additive logratio transformation, that is to define

$$z_j = \ln \frac{y_j}{y_D}, \quad j = 1, \dots, d,$$

where $d = D - 1$, and to take the vector $z = (z_1, \dots, z_d)$ to be normally distributed with mean μ and variance matrix Σ . The proportions y can be recovered by invoking the logistic transformation

$$y_i = \frac{\exp z_i}{1 + \sum_{i=1}^d \exp z_i}, \quad i = 1, \dots, d,$$

with $y_D = 1 - y_1 - \dots - y_d$, and are said to follow an additive logistic normal distribution.

Suppose now that the expected values of the proportions are specified by a mechanistic model and can be written as

$$E(y_i) = \eta_i(x; \theta), \quad i = 1, \dots, D,$$

where $x = (x_1, \dots, x_p)$ is a vector of explanatory variables defined on a design space \mathcal{X} , $\theta = (\theta_1, \dots, \theta_k)$ is a vector of unknown parameters, $\eta_i(x; \theta)$ represents a function nonlinear in and differentiable with respect to θ , and, necessarily, $\sum_{i=1}^D \eta_i(x; \theta) = 1$. Since the means $E(y_i)$ cannot be expressed explicitly in terms of the multivariate normal parameters μ and Σ , it is not a straightforward matter to translate their specification into the logratio formulation. An obvious, albeit somewhat cavalier, strategy is therefore to “transform both sides of the model” to give

$$E(z_i) = \mu_i(x; \theta) = \ln \left(\frac{E(y_i)}{E(y_D)} \right) = \ln \left(\frac{\eta_i(x; \theta)}{\eta_D(x; \theta)} \right), \quad i = 1, \dots, d,$$

and this is the approach adopted here.

The additive logistic normal approach to the modelling of proportions offers considerable flexibility in that all the advantages of the multivariate normal environment accrue. There are however some drawbacks. In particular the logratio transformation stabilizes variance, a feature which is arguably not appropriate, and in addition the resultant variance matrix Σ does not have an immediate interpretation. Further, as explained above, it is a straightforward matter to model the means of the log transformed variables but not to recover the means of the proportions themselves.

2.2 Design

Consider now an approximate design ξ which puts weight w_u on the vector of explanatory variables x_u , where $0 < w_u < 1$, $\sum_{u=1}^n w_u = 1$ and $u = 1, \dots, n$. Then it follows immediately from standard multivariate normal theory that the information matrix for the parameters θ at the design ξ is given by

$$M(\xi; \theta, \Sigma) = \sum_{u=1}^n w_u D_u \Sigma^{-1} D_u^T$$

where D_u is the $k \times d$ matrix with i th column equal to the vector of first-order derivatives $\frac{\partial \mu_i(x_u, \theta)}{\partial \theta}$, $i = 1, \dots, d$, $u = 1, \dots, n$ (Draper and Hunter, 1966). Note that $M(\xi; \theta, \Sigma)$ is necessarily nonlinear in the parameters θ and also that $M(x; \theta, \Sigma)$, the information matrix evaluated at a single design point x , may well have rank greater than 1.

Locally D -optimal designs, that is designs which minimize the generalized variance of $\hat{\theta}$, the maximum likelihood estimate of θ , at best guesses of the unknown parameters θ and Σ , can now be constructed by maximizing the determinant of the information matrix $M(\xi; \theta; \Sigma)$ at those best guesses. Since the information matrix depends in a complicated nonlinear manner on the design points and on the model parameters, algebraically tractable expressions for its determinant are not available, at least in general, and the required designs must therefore be obtained numerically. An Equivalence Theorem for the locally D -optimal criterion within the context of the present multi-response setting can be readily formulated following results presented in Fedorov (1972). Specifically the directional derivative of the criterion at a design ξ in the direction of a single design point x is given by

$$d(x, \xi; \theta, \Sigma) = \text{tr} \{ M^{-1}(\xi; \theta, \Sigma) M(x; \theta, \Sigma) \} - s$$

and the design ξ is globally optimal if this derivative is less than or equal to zero for all values of x in the design space \mathcal{X} , with equality holding at the support points of the design. However checking that the condition for global optimality is satisfied for a particular candidate design must again be done numerically.

3 Dirichlet Model

3.1 Model

The vector of proportions $y = (y_1, \dots, y_D)$ with $E(y_i) = \eta_i(x, \theta)$ can also be modelled directly by invoking the Dirichlet distribution with p.d.f.

$$g(y) = \frac{\Gamma(a_1 + \dots + a_D)}{\Gamma(a_1)\Gamma(a_2)\dots\Gamma(a_D)} y_1^{a_1-1} y_2^{a_2-1} \dots y_D^{a_D-1}$$

for $y_i > 0, \sum_{i=1}^n y_i = 1$ and unknown parameters $a_i > 0, i = 1, \dots, D$. The means of the proportions are given by $E(y_i) = \frac{a_i}{a_1 + a_2 + \dots + a_D}$ and are also specified through a mechanistic model by $\eta_i(x, \theta)$. Thus the relation

$$\eta_i(x, \theta) = \frac{a_i}{\phi}, \quad i = 1, \dots, D,$$

where $\phi = a_1 + a_2 + \dots + a_D$, holds, with the parameters a_i identified as meta-parameters. Furthermore the variances and covariances of the proportions can now be expressed succinctly as

$$\text{Var}(y_i) = \frac{\eta_i(x, \theta)(1 - \eta_i(x, \theta))}{(\phi + 1)} \quad \text{and} \quad \text{Cov}(y_i, y_j) = -\frac{\eta_i(x, \theta)\eta_j(x, \theta)}{(\phi + 1)}$$

for $i \neq j, i, j = 1, \dots, D$ and the parameter ϕ can thus be regarded as a precision parameter.

The Dirichlet distribution necessarily induces negative correlations between individual proportions and for this reason has been rarely used to model compositional data. It does however offer a more holistic approach to modelling proportions than that provided by the additive logistic normal and there have in fact been a few isolated studies within the context of Dirichlet regression. Specifically Campbell and Mosimann (1987), and more recently Hijazi (2006) and Hijazi and Jernigan (2007), introduced a constrained linear model to describe the Dirichlet parameters, while Gueorguieva, Rosenheck and Zelterman (2008) present a log linear model for that purpose. The strategy introduced in the present study is somewhat different in that the means of the proportions themselves are modelled directly, albeit in a nonlinear way. In fact the approach mirrors and extends that formulated and developed by Ferrari and Cribari-Neto (2004) for beta regression.

3.2 Design

The information matrix for the parameters $a = (a_1, \dots, a_D)$ of the Dirichlet distribution can be expressed succinctly as

$$M(a) = \text{Diag}\{\psi'(a_i)\} - \psi'(a_1 + a_2 + \dots + a_D)J$$

where $\psi'(\cdot)$ represents the trigamma function, that is the second derivative of the logarithm of the gamma function, and J is a square matrix of ones of order D . It thus follows that, within the present regression context, the information matrix for the parameters θ and ϕ at a single design point x is given by

$$M(x; \theta, \phi) = F M(a) F^T$$

where $M(a)$ is expressed in terms of $\eta_i(x; \theta)$, the parameter ϕ through the relationships $a_i = \phi \eta_i(x; \theta)$ and F is a $(k+1) \times D$ matrix with i th column

$$\begin{bmatrix} \frac{\partial a_i}{\partial \theta} \\ \frac{\partial a_i}{\partial \phi} \end{bmatrix} = \begin{bmatrix} \phi \frac{\partial \eta_i(x; \theta)}{\partial \theta} \\ \eta_i(x; \theta) \end{bmatrix}$$

for $i = 1, \dots, D$.

Consider again an approximate design ξ which puts weights w_u on the design points x_u , $u = 1, \dots, n$. Then the information matrix for the parameters θ and ϕ at the design ξ is given by

$$M(\xi; \theta, \phi) = \sum_{u=1}^n w_u M(x_u; \theta, \phi)$$

and depends in a complicated way on the design points, in particular through the trigamma functions embedded in the matrix $M(a)$. Locally D_S -optimal designs, that is designs which maximize information on the regression parameters θ with ϕ regarded as a nuisance parameter and which are evaluated at best guesses of those parameters, can now be constructed based on the matrix $M(\xi; \theta, \phi)$. Specifically, the D_S -optimality criterion can be formulated quite straightforwardly as

$$\ln |M(\xi; \theta, \phi)| - \ln [M_{\phi\phi}(\xi; \theta)],$$

where $M_{\phi\phi}(\xi; \theta)$ is the scalar submatrix of the full information matrix $M(\xi; \theta, \phi)$ relating to ϕ , and the directional derivative used in the associated Equivalence Theorem is given by

$$d(x, \xi; \theta, \phi) = \left[\text{tr} \{ M^{-1}(\xi; \theta, \phi) M(x; \theta, \phi) \} - \frac{M_{\phi\phi}(x; \theta)}{M_{\phi\phi}(\xi; \theta)} \right] - k.$$

Thus the requisite locally D_S -optimal designs can be constructed numerically and the global optimality or otherwise of a candidate design can be confirmed by invoking the appropriate Equivalence Theorem, again numerically.

4 Example

Consider two consecutive reactions represented schematically as $A \xrightarrow{\theta_1} B \xrightarrow{\theta_2} C$ where θ_1 and θ_2 are rate constants. Then the expected proportions of the reactants A, B and C at time x are given, respectively, by

$$\begin{aligned}\eta_1(x; \theta) &= \exp(-\theta_1 x) \\ \eta_2(x; \theta) &= \frac{\theta_1}{(\theta_2 - \theta_1)} [\exp(-\theta_1 x) - \exp(-\theta_2 x)] \\ \eta_3(x; \theta) &= 1 - \left[\frac{\theta_2 \exp(-\theta_1 x) - \theta_1 \exp(-\theta_2 x)}{(\theta_2 - \theta_1)} \right],\end{aligned}$$

where x represents time and $\theta = (\theta_1, \theta_2)$. The best guesses for the parameters are taken to be $\theta_1 = 0.7$ and $\theta_2 = 0.2$ and the design space to be the interval $[1, 10]$. (See for example Draper and Hunter, 1966, and Atkinson, 2003).

4.1 Additive Logistic Normal

Consider first modelling the proportion of a single reactant, either A , B or C . Then the additive logistic normal model comprises a single response with mean taken to be equal to the logit transformation of the mean of the proportion, that is $\ln \frac{\eta_i(x; \theta)}{1 - \eta_i(x; \theta)}$ where $i = 1, 2$ or 3 . Locally D -optimal designs for the three reactants, considered separately, are somewhat curious. In particular the optimal design for reactant A comprises a single point at the maximum time, that is 10, the design for B puts equal weights on the points 1.229 and the maximum time, and the design for C places equal weights on the minimum and maximum times of 1 and 10. The results are similar in form to those presented by Atkinson (2003) for log transformed means.

Consider now invoking the additive logistic normal to model the proportions of the three reactants. In particular suppose that the proportions of reactants A and B , with C as the benchmark, are modelled as $\mu_1(x; \theta) = \ln \frac{\eta_1(x; \theta)}{\eta_3(x; \theta)}$ and $\mu_2(x; \theta) = \ln \frac{\eta_2(x; \theta)}{\eta_3(x; \theta)}$ for some choice of the variance matrix Σ . Locally D -optimal designs for this setting were constructed for a range of Σ and, for example, for $\Sigma = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ and $\Sigma = \begin{pmatrix} 1 & 0.1 \\ 0.1 & 1 \end{pmatrix}$ single point designs comprising the maximum time point were obtained while for $\Sigma = \begin{pmatrix} 1 & -0.9 \\ -0.9 & 1 \end{pmatrix}$ the optimal design places weights of 0.288 and 0.712 on the extreme time points 1 and 10 respectively. It is thus clear that the designs depend sensitively on the choice of Σ . Furthermore suppose now that proportions of B and C , with A as the benchmark, are modelled, that is $\mu_1(x; \theta) = \ln \frac{\eta_2(x; \theta)}{\eta_1(x; \theta)}$ and $\mu_2(x; \theta) = \ln \frac{\eta_3(x; \theta)}{\eta_1(x; \theta)}$. Then it is not immediately clear as to how the choice of the variance matrix Σ in this case relates to that used to model A and B with the proportion of C as the benchmark. The fact that the designs depend crucially on Σ merely exacerbates this problem.

4.2 Dirichlet

Consider first modelling the proportions of the reactants separately by invoking the beta distribution as a special case of the Dirichlet. Locally D_S -optimal designs for this setting can be readily constructed numerically and the results for $\phi = 1$ and 100 are summarized in Table 1. The dependence of the optimal designs on precision, as

Table 1: Optimal designs for beta regression. (Weights in brackets.)

	$\phi = 1$		$\phi = 100$	
A	1 (0.537)	10 (0.463)	2.277	
B	1.227 (0.578)	10 (0.422)	1.227 (0.501)	9.743 (0.499)
C	1 (0.356)	4.950 (0.087)	10 (0.557)	2.582 (0.501) 10 (0.499)

captured by the parameter ϕ , is marked. Overall, and in contrast to the results for the additive logistic normal model, the designs tend to favour times towards the centre of the design space, that is times at which the variances of the responses are large within the Dirichlet context, as well as times at the extremes of the design space.

Now consider extending the modelling context and using the Dirichlet distribution to model the proportions of the three reactants A , B and C simultaneously. Locally D_S -optimal designs for selected values of the precision parameter ϕ are presented in Table 2 and again, and not surprisingly, exhibit strong dependence on

Table 2: Optimal designs for Dirichlet regression.

	$\phi = 1$		$\phi = 10$		$\phi = 100$		$\phi = 1000$
x	1	10	1	10	3.443	10	4.056
w	0.242	0.758	0.105	0.895	0.635	0.365	1

the precision parameter. Indeed the results essentially mirror those obtained for individual proportions.

5 Conclusions

The main aim of the present study has been to construct optimal designs for experiments involving compositional data, specifically locally D -optimal designs for the additive logistic normal model and locally D_S -optimal designs for Dirichlet regression. The theory underpinning the construction of these designs is based on the appropriate information matrices and the development is quite straightforward. The results for the selected example, that of two consecutive reactions, are however somewhat disappointing. In particular the locally D -optimal designs for both models of interest tend to put weights on the minimum and maximum time points. In

addition designs for the additive logistic normal setting are not easily formulated and interpreted.

There is much scope for further research. From a modelling perspective, strategies for fitting compositional data based, for example, on the Liouville distribution, following Iyengar and Dey (2002), and on the simplex dispersion model of Barndorff-Nielson and Jorgensen (1991), could well be examined. From a design perspective, extensions to the present setting relating to the model, the functions invoked to describe the means of the proportions and the optimality criteria are immediately attractive. At present work is in progress to explore in some detail the construction of optimal designs for the base scenario, that of beta regression.

Acknowledgements Linda Haines would like to thank the University of Cape Town, the National Research Foundation and SASOL, South Africa, for financial support.

References

- Aitchison, J. (2004). *The Statistical Analysis of Compositional Data*. New Jersey: Blackburn Press.
- Atkinson, A. C. (2003). Horwitz's rule, transforming both sides and the design of experiments for mechanistic models. *Applied Statistics* 52, 261–278.
- Atkinson, A. C., A. N. Donev, and R. D. Tobias (2007). *Optimum Experimental Designs, with SAS*. Oxford: Oxford University Press.
- Barndorff-Nielson, O. E. and B. Jorgensen (1991). Some parametric models on the simplex. *Journal of Multivariate Analysis* 39, 106–116.
- Campbell, G. and J. E. Mosimann (1987). Multivariate analysis of size and shape: modelling with the dirichlet distribution. *ASA Proceedings of Section on Statistical Graphics*, 93–101.
- Draper, N. R. and W. G. Hunter (1966). Design of experiments for parameter estimation in multiresponse situations. *Biometrika* 53, 525–533.
- Fedorov, V. V. (1972). *Theory of Optimum Experiments*. New York: Academic Press.
- Ferrari, S. L. and F. Cribari-Neto (2004). Beta regression for modelling rates and proportions. *Journal of Applied Statistics* 31, 799–815.
- Gueorguieva, R., R. Rosenheck, and D. Zelterman (2008). Dirichlet component regression and its applications to psychiatric data. *Computational Statistics and Data Analysis* 52, 5344–5355.
- Hijazi, R. H. (2006). Residuals and diagnostics in dirichlet regression. *ASA Proceedings of the General Methodology Section*, 1190–1196.
- Hijazi, R. H. and R. W. Jernigan (2007). Modelling compositional data using dirichlet regression models. Technical report, Preprint.
- Iyengar, M. and D. K. Dey (2002). A semiparametric model for compositional data analysis in the presence of covariates on the simplex. *Test* 11, 303–315.

Dose Finding Experiments: Responses of Mixed Type

Valerii V. Fedorov, Yuehui Wu and Rongmei Zhang

Abstract Multiple-endpoint models are widely used in drug development and other fields. It is common that the endpoints have different characteristics such as all continuous, all binary, or a mixture of them. This study investigates mixed responses, one continuous and one binary, correlated and observed simultaneously. It is an extension of our previous studies based on a bivariate probit model for two binary endpoints. We quantify the study goal with a utility function, construct locally two-stage D-optimal designs under the constraints, and use them as benchmarks for the two-stage designs with interim adjustment and fully adaptive designs. The simulation results suggest that the two-stage design is almost as efficient as the locally two-stage optimal design, as well as being logistically simpler than the fully adaptive design. We do not analyze asymptotic properties but confine ourselves to Monte-Carlo simulations to evaluate their properties for reasonable (practical) sample sizes.

1 Introduction

This article is an extension of our research on designs based on a bivariate probit model (Dragalin, Fedorov, and Wu 2008; Fedorov and Wu 2007) and is focused on mixed correlated responses, one continuous and one binary, observed simultaneously.

Valerii V. Fedorov, Yuehui Wu
GlaxoSmithKline, Collegeville, PA, USA e-mail: Valeri.V.Fedorov@gsk.com, Yuehui.2.Wu@gsk.com

Rongmei Zhang
CCEB, University of Pennsylvania, Philadelphia, PA, USA e-mail: rongmei@mail.med.upenn.edu

2 Model

Assume that the two responses (e.g. efficacy and toxicity) follow an underlying bivariate normal distribution $Z \sim N(\eta, \Sigma)$, where Z is a vector of responses, with mean $\eta = (\eta_1, \eta_2)$ and variance-covariance matrix $\Sigma = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix}$. In this study, we focus on mixed observed responses when Y_1 (efficacy) is continuous and Y_2 (toxicity) is binary (dichotomized) such that

$$Y_1 = Z_1, \quad Y_2 = \begin{cases} 1, & \text{if } Z_2 \geq c_2 \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

The marginal probability for toxicity is $P(Y_2 = 1) = p_{.1} = 1 - F(v_2)$ with $v_2 = (c_2 - \eta_2)/\sigma_2$, where $F(x)$ denotes the c.d.f. of the standard normal distribution. The conditional probability of $Y_2 = 1$ given Y_1 is $p_{1|y_1} = P(Y_2 = 1|Y_1 = y_1) = 1 - F(u_2)$, where $u_2 = \frac{v_2 - \rho \frac{y_1 - \eta_1}{\sigma_1}}{\sqrt{1 - \rho^2}}$. We denote the marginal probability density of efficacy as $\varphi(y_1)$, the probability density for a normal distribution with mean η_1 and standard deviation σ_1 . The other probabilities can easily be derived. The log likelihood for a single observation is:

$$l(y_1, y_2; \vartheta) \propto y_2 \log \{1 - F(u_2)\} + (1 - y_2) \log \{F(u_2)\} - \log \sigma_1 - \frac{(y_1 - \eta_1)^2}{2\sigma_1^2},$$

where $\vartheta = (\eta_1, v_2, \rho, \sigma_1)^T$ are the elemental parameters. In practice, these parameters may depend on some covariates such as doses of various compounds (drugs), age, gender, etc. Although η_1 and σ_1 can be estimated separately, their counterparts η_2 and σ_2 cannot. Only $(c_2 - \eta_2)/\sigma_2$ is estimable (Dragalin et al. 2008).

2.1 Information Matrix for a Single Observation

In experimental design, the information matrix plays a crucial role since it is the basis for the formulation of the optimality criterion and determines the allocation of patients (Fedorov and Hackl 1997). Because the information matrix of independent observations is the sum of the information matrices of all single observations, the derivation of the information matrix for a single observation becomes a central step in optimal design construction.

Elemental Information Matrix: The information matrix of a single observation under model (1) can be found in (4.4) of Tate (1955):

$$\mu(\vartheta) = \begin{pmatrix} \frac{1-\rho^2+\rho^2a_0}{\sigma_1^2(1-\rho^2)} & \frac{\rho a_0}{\sigma_1(1-\rho^2)} & \frac{\rho(\rho v_2 a_0 - a_1)}{\sigma_1(1-\rho^2)^2} & \frac{\rho^2 a_1}{\sigma_1^2(1-\rho^2)} \\ \frac{a_0}{(1-\rho^2)} & \frac{\rho v_2 a_0 - a_1}{(1-\rho^2)^2} & \frac{\rho v_2 a_1 - a_2}{\sigma_1(1-\rho^2)} & \frac{\rho^2 a_1}{\sigma_1^2(1-\rho^2)} \\ \text{Symmetric} & \frac{a_2 - 2\rho v_2 a_1 + \rho^2 v_2^2 a_0}{(1-\rho^2)^3} & \frac{\rho(\rho v_2 a_1 - a_2)}{\sigma_1(1-\rho^2)^2} & \frac{2(1-\rho^2) + \rho^2 a_2}{\sigma_1^2(1-\rho^2)} \end{pmatrix}, \quad (2)$$

where

$$a_k(v_2, \rho) = \int_{-\infty}^{+\infty} \frac{t^k \varphi(t) \varphi^2\left(\frac{v_2 - \rho t}{\sqrt{1-\rho^2}}\right)}{F\left(\frac{v_2 - \rho t}{\sqrt{1-\rho^2}}\right) \left[1 - F\left(\frac{v_2 - \rho t}{\sqrt{1-\rho^2}}\right)\right]} dt \quad \text{and } k = 0, 1, 2. \quad (3)$$

In (3), except for $\rho = 0$, we have to use numerical integration to find a_k . If Σ is known, then $\vartheta = (\eta_1, v_2)^T$ and their information matrix is the upper-left 2×2 submatrix of (2).

Assume that $\eta_1 = \theta_1^T f_1(x)$ and $v_2 = (c_2 - \eta_2)/\sigma_2 = \theta_2^T f_2(x)$, i.e. $\theta = (\theta_1, \theta_2, \rho, \sigma_1)^T$, then the information matrix for unknown parameters θ of a single observation can be easily calculated using the Jacobian transformation.

2.2 Utility and Penalty Functions

For the mixed responses of continuous efficacy and binary toxicity, observed simultaneously, we define our utility function as the expected value of efficacy given no toxicity multiplied by the probability of having no toxicity,

$$\begin{aligned} \zeta(x, \theta) &= E(Y_1 | Y_2 = 0) P(Y_2 = 0) = \eta_1 F(v_2) - \rho \sigma_1 \varphi(v_2) \\ &= \theta_1^T f_1(x) F(\theta_2^T f_2(x)) - \rho \sigma_1 \varphi(\theta_2^T f_2(x)). \end{aligned} \quad (4)$$

In drug development studies, there are always ethical concerns and cost constraints associated with different doses. These needs can be quantified by introducing a penalty function $\phi(x, \theta)$. We use a penalty function similar to that of Lai and Robbins (1978):

$$\phi(x, \theta) = r\{x - X^*(\theta)\}^2 + c, \quad (5)$$

where r is the constant controlling the magnitude of the penalty and $c > 0$ represents the cost per subject; $X^*(\theta) = \arg \max_{x \in \mathcal{X}} \zeta(x, \theta)$ is the best dose. Note that only the ratio of r/c in (5) affects the optimal design; see (7). We assume that $\gamma = r/c = 0.5$. When $\gamma = 0$, the optimal design coincides with the optimal design without constraints.

Figure 1 illustrates a possible relationship between the dose-response curves, utility function and penalty function under the proposed model. Model parameters were estimated from the VTE trial (see details in Dragalin et al. 2008).

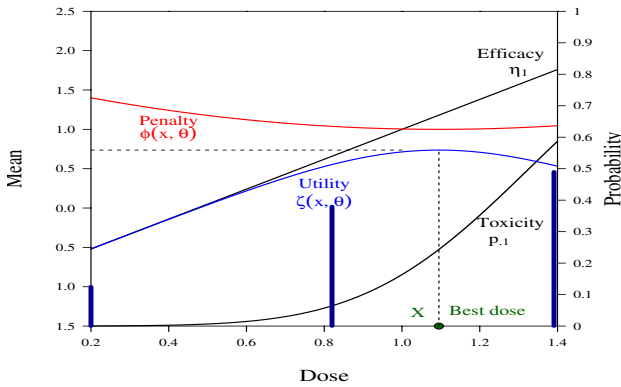


Fig. 1: Efficacy η , toxicity p_1 , utility $\zeta(x, \theta)$ and penalty $\phi(x, \theta)$ for the mixed-outcome model with $\theta = (-0.9, 1.9, 3.98, -3)$, $\sigma_1 = 1$, $\rho = 0.5$. The left y-axis is for efficacy, utility, and penalty; the right y-axis is for toxicity and the bars denote the locally optimal design.

3 Optimal Designs

How to construct a locally optimal design is well known and discussed in detail in Fedorov and Hackl (1997). Here we only focus on the construction of a two-stage design. Compared to the administrative complexity of conducting fully adaptive designs, two-stage designs have the clear advantages of simplifying the patient enrollment and drug supply procedure and providing results which statistically (in terms of the bias and variability) are very close to, or even better than, fully adaptive designs.

Let $\xi = \{x_i, \lambda_i\}_1^k$ denote the design, within the design region \mathcal{X} . The sample size at design point $x_i \in \mathcal{X}$ is n_i , and the weight for design point x_i is $\lambda_i = n_i/N$, where $i = 1, \dots, k$, $N = \sum_{i=1}^k n_i$.

The idea of the two-stage design is that after N_0 observations in the initial design stage, the researcher obtains initial estimates for the unknown parameters, $\hat{\theta}_0$, and then at the second design stage, the $\hat{\theta}_0$ are treated as the “true” parameters on which the locally optimal design is constructed. At the end of the second stage, we re-estimate the unknown parameters using all $N_0 + N_1$ observations.

To get a benchmark for the actual two-stage design (i.e. with adaptation after the first stage), we start with the locally two-stage optimal design. To build this design we assume that the “true” values of the estimated parameters are known.

The study goal is to identify the target dose $X^*(\theta)$ and the behaviour of response functions $\eta_1(x, \theta_1)$ and $v_2(x, \theta_2)$ which involve all the unknown parameters. Therefore we select the D-criterion and define a locally optimal design as:

$$\xi^*(\theta) = \arg \min_{\xi} |\{\pi M(\xi_0, \theta) + (1 - \pi)M(\xi, \theta)\} / \Phi_T(\xi, \theta)|^{-1}, \quad (6)$$

where ξ_0 denotes the initial design, $\pi = N_0 / (N_0 + N_1)$, $\Phi(\xi, \theta) = \sum_{i=1}^k \lambda_i \phi(x_i, \theta)$, $\Phi_T(\xi, \theta) = \pi \Phi(\xi_0, \theta) + (1 - \pi)\Phi(\xi, \theta)$.

The necessary and sufficient condition for optimality of ξ^* (see Fedorov and Hackl 1997) is

$$\begin{aligned} \psi(x, \xi^*, \theta) = \operatorname{tr} \left\{ \frac{\mu(x, \theta)}{\phi(x, \theta)} [\pi M(\xi_0, \theta) + (1 - \pi)M(\xi^*, \theta)]^{-1} \right\} \leq \\ \operatorname{tr} \left\{ \frac{M(\xi^*, \theta)}{\Phi(\xi^*, \theta)} [\pi M(\xi_0, \theta) + (1 - \pi)M(\xi^*, \theta)]^{-1} \right\}, \end{aligned} \quad (7)$$

where x is any design point in the design region \mathcal{X} . The forward and backward steps in the first-order exchange algorithm, Fedorov (1972), are defined as

$$x_s^+ = \arg \max_{x \in \mathcal{X}} [\psi(x, \xi_s, \theta)] \quad \text{and} \quad x_s^- = \arg \min_{x \in \mathcal{X}} [\psi(x, \xi_s, \theta)]. \quad (8)$$

In the actual two-stage (composite) design, the true parameters θ are replaced by their estimates $\hat{\theta}_0$ found after the analysis of the first stage data. The matrix $[\pi M(\xi_0, \hat{\theta}_0) + (1 - \pi)M(\xi^*, \hat{\theta}_0)]$ is not the exact normalized information matrix any more because ξ^* depends on $\hat{\theta}_0$. However, at an intuitive level, we can make the following conjecture. If the initial design ξ_0 has a regular information matrix for any $\theta \in \Omega$, where the admissibility set Ω includes the true value of θ as an internal point, then the maximum likelihood estimator $\hat{\theta}_0 = \hat{\theta}(N_0)$ will be strongly consistent, i.e. converge almost surely to θ_{true} when $N_0 \rightarrow \infty$. Consequently the sensitivity function $\psi\{x, \xi, \hat{\theta}(N_0)\}$ will converge almost surely to $\psi(x, \xi, \theta_{true})$, uniformly with respect to x and ξ . Obviously, some smoothness of $f_1(x)$ and $f_2(x)$ is needed together with the compactness of \mathcal{X} (Rao 1973, §2.c). Consequently, the solution $\xi^*\{\hat{\theta}(N_0)\}$ will converge to $\xi^*(\theta_{true})$ and $\pi M\{\xi_0, \hat{\theta}(N_0)\} + (1 - \pi)M\{\xi^*, \hat{\theta}(N_0)\}$ will converge to the “true” information matrix $\pi M(\xi_0, \theta_{true}) + (1 - \pi)M(\xi^*, \theta_{true})$. Of course, the above statement is only a conjecture without any rigorous mathematical proof, which is why we resort to Monte-Carlo simulations to confirm its validity for our specific case.

3.1 Adaptive Designs

Fully D-adaptive designs under the penalty function (5) are also constructed for comparison with the two-stage design. In an adaptive design, the estimates are updated as new observations arrive. At each step, the next patient will be assigned to the dose which maximizes the current sensitivity function $\psi(x, \xi_s, \hat{\theta}_N)$. Compare this with the forward step in (8) with θ being replaced by $\hat{\theta}_N$. Note that, as in the two-stage design, in fully adaptive designs, the observations are not independent and the likelihood function should be conditioned on the previous observations. Again,

$\pi M(\xi_0, \hat{\theta}_N) + (1 - \pi)M(\xi^*, \hat{\theta}_N)$ is random and it is not the actual information matrix. However, the distribution of each element of the matrix heavily gravitates towards to the corresponding element in $[\pi M(\xi_0, \theta_{true}) + (1 - \pi)M(\xi^*, \theta_{true})]$. See our simulation results and Lai (2001) and Rosenberger and Hughes-Oliver (1999).

4 Examples

For illustration and comparison purposes, we used the model of Dragalin et al. (2008) for trials on the prevention of VTE. In model (1), assume that the parameters $\rho = 0.5$ and $\sigma_1 = 1$ are known. For the other two elemental parameters η_1 and v_2 , assume $\eta_1 = \theta_{11} + \theta_{12}x$ and $v_2 = \frac{c_2 - \eta_2}{\sigma_2} = \theta_{21} + \theta_{22}x$, with $\theta = (\theta_{11}, \theta_{12}, \theta_{21}, \theta_{22}) = (-0.9, 1.9, 3.98, -3)$. The design region is $[0.2, 1.4]$. The locally D-optimal design is $\{0.2, 0.13; 0.81, 0.38; 1.4, 0.49\}$, marked with bars in Figure 1.

The second-stage optimal design depends on the random vector $\hat{\theta}_0$. We conducted 1000 simulations of the two-stage design with $N_0 = 80$ and $N_1 = 120$ from an evenly spaced, equally weighted, five-point first-stage design $\{0.2, 0.2; 0.5, 0.2; 0.8, 0.2; 1.1, 0.2; 1.4, 0.2\}$. The left plot of Figure 2 gives the location and frequency of the design points in the second stage from the simulations. It shows that most of the points are located close to or on the locally two-stage optimal design points $\{0.2, 0.02; 0.81, 0.35; 1.4, 0.63\}$, marked as black dots (with the size of the dots corresponding to the weight). Note that the locally two-stage D-optimal design contains the same design points as the locally D-optimal design while the weights are different; in the simulation study, the two-stage D-optimal design will be used as the benchmark. The central plot illustrates the distribution for the predicted best dose \hat{X}^* after 200 observations. The reference line indicates the best dose ($X^*(\theta) = 1.095$). The curve denotes the fitted normal density with mean X^* and variance coinciding with the asymptotic variance of X^* for the locally two-stage optimal design (6).

The simulation shows that for the selected sample sizes (80 + 120), the two-stage procedure defined in Section 1.3 provides results that are hardly distinguishable from the benchmarks derived for the locally two-stage optimal designs.

Practitioners may be interested in the choice of the proportion of the sample sizes for the two stages in the two-stage designs. A larger N_0 at the first stage leads to a more accurate initial estimate of θ , but leaves a smaller sample size for the optimized stage. Overall, we may have an inefficient design. On the other hand, a small value of N_0 can result in a less accurate $\hat{\theta}$, which consequently may lead to an optimal design far from the true one. Monte-Carlo simulations for different N_0 can be helpful to identify the acceptable range of N_0 in terms of efficiency and accuracy. As N_0 varied from 20 to 120, with the other conditions fixed, the information per penalty $|M(\xi, \theta) / \Phi(\xi, \theta)|^{1/m}$, where m is the number of unknown parameters, was used to compare the different partitions of the total sample size. The values for locally D-optimal and uniform designs are drawn in the right panel of Figure 2 as reference lines. The solid lines represent the values under true parameters; the dashed lines represent the values for the 10th, 50th and 90th quantiles for 1000




Fig. 2: Simulations for the two-stage design. **Left:** Locations of design points in the second stage. **Center:** Distributions of the predicted “best dose” \hat{X}^* . **Right:** Information per penalty unit for two-stage designs with different N_0 .

simulations. The simulations suggest that a moderate sample size (60 or 80) in the first stage has relatively high accuracy and low variation.

Fully adaptive design simulations were conducted under the same setting as the two-stage design. Due to space limitations, these results are not shown here. The allocations of the doses at the final stage are distributed quite close to those of the locally D-optimal design and the distribution of the “best dose” X^* has a normal density similar to that of Figure 2.

5 Conclusions

Based on a bivariate model for continuous efficacy and discrete toxicity, we propose a dose-finding procedure based on the theory of optimal experimental design. A utility function is defined to quantify the targeted treatment effects, and a rather flexible penalty function is used to address ethical issues and cost constraints in a drug development study.

Two-stage design and the fully adaptive design for various scenarios were constructed and compared to the locally two-stage optimal designs via Monte-Carlo simulations. The results confirm our conjectures about properties of two-stage and fully-adaptive designs for reasonably large sample sizes. In both cases, the distributions of assignment of patients gravitate to the support points of the locally two-stage optimal design. As for the accuracy of the estimation of the target dose X^* , the two-stage design is quite close to the fully adaptive design with the same initial sample size. Actually, the two-stage designs in our example have slightly smaller variance than the fully adaptive designs, while the bias is minuscule. We recommend selec-

tion of the sample size for the initial stage using information about quantiles of the targeted “precision” matrices. These facts indicate that two-stage design is an efficient and practical procedure that should be strongly recommended.

References

- Dragalin, V., V. Fedorov, and Y. Wu (2008). Two-stage design for dose-finding that accounts for both efficacy and safety. *Statistics in Medicine* 27, 5156–5176.
- Fedorov, V. (1972). *Theory of Optimal Experiments*. New York: Academic Press.
- Fedorov, V. and Y. Wu (2007). Dose finding designs for continuous responses and binary utility. *Journal of Biopharmaceutical Statistics* 17, 1085–1096.
- Fedorov, V. V. and P. Hackl (1997). *Model-oriented Design of Experiments*. Berlin: Springer-Verlag.
- Lai, T. L. (2001). Sequential analysis: Some classical problems and new challenges. *Statistica Sinica* 11, 303–351.
- Lai, T. L. and H. Robbins (1978). Adaptive design in regression and control. *Proceedings of the National Academy of Sciences of the United States of America* 75, 586–587.
- Rao, C. R. (1973). *Linear Statistical Inference and Its Applications*. New York: Wiley.
- Rosenberger, W. F. and J. M. Hughes-Oliver (1999). Inference from a sequential design: Proof of a conjecture by Ford and Silvey. *Statistics & Probability Letters* 44, 177–180.
- Tate, R. (1955). The theory of correlation between two continuous variables when one is dichotomized. *Biometrika* 42, 205–216.

Pharmacokinetic Studies Described by Stochastic Differential Equations: Optimal Design for Systems with Positive Trajectories

Valerii V. Fedorov, Sergei L. Leonov and Vyacheslav A. Vasiliev

Abstract In compartmental pharmacokinetic (PK) modelling, ordinary differential equations (ODE) are traditionally used with two sources of randomness: measurement error and population variability. In this paper we focus on intrinsic (within-subject) variability modelled with stochastic differential equations (SDE), and consider stochastic systems with positive trajectories which are important from a physiological perspective. We derive mean and covariance functions of solutions of SDE models, and construct optimal designs, i.e. find sampling schemes that provide the most precise estimation of model parameters under cost constraints.

1 Introduction

In this paper we continue the research presented at the mODa-8 Conference, see Anisimov et al. (2007), and discuss PK models described by stochastic differential equations. We concentrate on within-subject variability or an “intrinsic variability of the metabolic system”, as stated in Picchini et al. (2006). This intrinsic variability suggests a move from ODE to SDE. Once expressions for the mean and covariance functions of the solution of the SDE system are derived, one may address the optimal design problem, i.e. selection of sequences of sampling times that “minimize”, in some sense, the variance-covariance matrix of parameter estimates.

Anisimov et al. (2007) explored a popular PK model which is described by a system of ODE, introduced its SDE analogue and derived the mean and covariance functions of the solutions together with optimal designs for several examples. To guarantee a meaningful physiological behaviour of the stochastic systems, a rather restrictive constraint of diminishing variability of the noise process was imposed in an attempt to have positive solutions of the SDE system. However, this goal has not been fully accomplished; see Section 2 for details. In this paper we consider less restrictive models that still guarantee the positiveness of the SDE solution. Note that

Valerii V. Fedorov, Sergei L. Leonov

GlaxoSmithKline, 1250 So. Collegeville Rd, PO Box 5089, Collegeville, PA, 19426, U.S.A., e-mail: Valeri.V.Fedorov@gsk.com, Sergei.2.Leonov@gsk.com

Vyacheslav A. Vasiliev

Tomsk State University, Department of Applied Mathematics and Cybernetics, Tomsk, Russia, e-mail: vas@mail.tsu.ru

in a number of previous publications on PK models, no attention was given to the fact that trajectories of the continuous stochastic systems may become negative with positive probability; see Overgaard et al. (2005) or Tornøe et al. (2004).

The paper is organized as follows. In Section 2 an outline of the earlier work on optimal designs for stochastic PK models is presented and an SDE system with positive trajectories is introduced and its covariance function derived. In Section 3 we present examples of optimal sampling schemes. It is shown that when costs are taken into account, designs with fewer samples may be preferred to more “dense” sampling schemes. Potential extensions are discussed in Section 4.

2 Response Models

To illustrate optimal design techniques for stochastic PK models, Anisimov et al. (2007) considered a one-compartment model with first-order absorption,

$$d\eta_1 = -\theta_1\eta_1 dt, \quad d\eta_2 = \theta_1\eta_1 dt - \theta_2\eta_2 dt; \quad \eta_1(0) = D; \quad \eta_2(0) = 0, \quad (1)$$

where $\eta_1(t)$ is the drug amount at the site of administration, $\eta_2(t)$ is the drug amount in the central compartment, D is the dose; θ_1, θ_2 are absorption and elimination rates, respectively; $\theta_1 > \theta_2$. The solution of the system (1) is positive for any $t > 0$,

$$\eta_1(t) = De^{-\theta_1 t}, \quad \eta_2(t) = [\theta_1 D / (\theta_1 - \theta_2)] (e^{-\theta_2 t} - e^{-\theta_1 t}). \quad (2)$$

The simplest stochastic analog of the system (1) may be introduced by

$$\begin{aligned} dy_1(t) &= -\theta_1 y_1(t) dt + \sigma_1(t) dw_1(t), \quad y_1(0) = D, \\ dy_2(t) &= \theta_1 y_1(t) dt - \theta_2 y_2(t) dt + \sigma_2(t) dw_2(t), \quad y_2(0) = 0, \end{aligned} \quad (3)$$

where $\sigma_r(t) \geq 0$ are deterministic functions, and $w_r(t)$ are independent Wiener processes for which $\mathbf{E}[w_r(t)] = 0$, $\mathbf{Cov}[w_r(t)w_r(s)] = \min(t, s)$; $r = 1, 2$. Using properties of the Itô integral and the independence of $w_1(t)$ and $w_2(t)$, it is straightforward to show the unbiasedness of the solution, i.e. $\mathbf{E}[y_r(t)] = \eta_r(t)$, $r = 1, 2$, and to derive an expression for the covariance function $\tilde{S}(t, t+s) = \mathbf{Cov}[y_2(t), y_2(t+s)]$ via weighted integrals of $\sigma_r^2(t)$. Anisimov et al. (2007) provided a closed-form solution for $\sigma_r(t) = \sigma_r e^{-\nu_r t}$ and showed that when $\nu_r > 0$, then $\mathbf{Var}[y_2(t)] \rightarrow 0$ as $t \rightarrow \infty$, and $\tilde{S}(t, t+s) \rightarrow 0$ as $s \rightarrow \infty$ for any fixed $t \geq 0$. However, the probability \tilde{P}_t that paths $y_2(t)$ may become negative, is nonzero for any fixed t . Moreover, if at least one $\nu_r = 0$, then $\mathbf{Var}[y_2(t)] \rightarrow \nu^* > 0$, and $\tilde{P}_t \rightarrow 0.5$ as $t \rightarrow \infty$ since $\mathbf{E}[y_2(t)] \rightarrow 0$.

Nevertheless models similar to (3) are popular due to their relative simplicity and the Gaussian distribution of the solutions. For instance, Overgaard et al. (2005) considered an exponential decay model analogous to the first equation in (3) with $\sigma_1(t) \equiv \sigma_w > 0$, and then took logarithms of the solution in the discrete-time model

of observations. Because of potential negativeness of y_1 , such a model seems counterintuitive from physiological considerations, and taking logs is formally incorrect.

2.1 Stochastic Systems with Positive Trajectories

It would be desirable to consider stochastic models with positive solutions $y_r(t)$, for example replacing the noise terms $\sigma_r(t)dw_r(t)$ on the right-hand side of (3) with $\sigma_r y_r(t)dw_r(t)$, $\sigma_r > 0$, i.e. making them proportional to the signal,

$$\begin{aligned} dy_1(t) &= -\theta_1 y_1(t)dt + \sigma_1 y_1(t)dw_1(t), \quad y_1(0) = D, \\ dy_2(t) &= \theta_1 y_1(t)dt - \theta_2 y_2(t)dt + \sigma_2 y_2(t)dw_2(t), \quad y_2(0) = 0. \end{aligned} \quad (4)$$

Such variance terms are analogs of a proportional component of variance. The problem of deriving the covariance function \tilde{S} for the system (4) was posed in Anisimov et al. (2007), and in this section we report its closed-form solution.

Lemma 1. (a) *The solution of system (4) is given by*

$$\begin{aligned} y_1(t) &= D e^{-(\theta_1 + \frac{\sigma_1^2}{2})t + \sigma_1 w_1(t)}, \\ y_2(t) &= \theta_1 D e^{-(\theta_2 + \frac{\sigma_2^2}{2})t + \sigma_2 w_2(t)} \cdot \int_0^t e^{(\theta_2 - \theta_1 + \frac{\sigma_2^2 - \sigma_1^2}{2})s + \sigma_1 w_1(s) - \sigma_2 w_2(s)} ds. \end{aligned} \quad (5)$$

(b) *The solution is unbiased: $\mathbf{E}y_r(t) = \eta_r(t)$, $r = 1, 2$, with $\eta_r(t)$ defined in (2).*

(c) *The covariance function $\tilde{S}(t, t+s)$ satisfies*

$$\begin{aligned} \tilde{S}(t, t+s) &= \theta_1^2 D^2 \{ a_1 e^{(-2\theta_1 + \sigma_1^2)t - \theta_2 s} + a_2 e^{(-2\theta_1 + \sigma_1^2)t - \theta_1 s} + a_3 e^{(-2\theta_2 + \sigma_2^2)t - \theta_2 s} \\ &\quad + a_4 e^{-(\theta_1 + \theta_2)t - \theta_1 s} + a_5 e^{-(\theta_1 + \theta_2)t - \theta_2 s} + a_6 (e^{-2\theta_1 t - \theta_1 s} + e^{-2\theta_2 t - \theta_2 s}) \}, \end{aligned} \quad (6)$$

where

$$\begin{aligned} a_1 &= \frac{\sigma_2^2 - \sigma_1^2}{\Delta\theta(\Delta\theta + \sigma_1^2)[2\Delta\theta - \sigma_2^2 + \sigma_1^2]}, \quad \Delta\theta = \theta_2 - \theta_1, \\ a_2 &= \frac{1}{\Delta\theta(\Delta\theta + \sigma_1^2)}, \quad a_3 = \frac{2}{(\Delta\theta - \sigma_2^2)[2\Delta\theta - \sigma_2^2 + \sigma_1^2]}, \\ a_4 &= \frac{\sigma_1^2}{(\Delta\theta)^2(\Delta\theta + \sigma_1^2)}, \quad a_5 = \frac{(\sigma_1^2 - 2\sigma_2^2)\Delta\theta - \sigma_1^2\sigma_2^2}{(\Delta\theta)^2(\Delta\theta + \sigma_1^2)(\Delta\theta - \sigma_2^2)}, \quad a_6 = \frac{-1}{(\Delta\theta)^2}. \end{aligned}$$

The derivation is postponed to the Appendix. Note that (5) immediately implies that y_1, y_2 are positive, and it follows from (6) that $\tilde{S}(t, t+s) \rightarrow 0$ as $t \rightarrow \infty$ for any fixed $s \geq 0$ if $\theta_1 > \sigma_1^2/2$ and $\theta_2 > \sigma_2^2/2$; and $\tilde{S}(t, t+s) \rightarrow 0$ as $s \rightarrow \infty$ for any fixed $t \geq 0$.

Measurement Model. Let t_{i1}, \dots, t_{i,k_i} be a sequence of k_i sampling times for subject i and let $\tilde{\mathbf{S}} = \{\tilde{S}(t_{ij_1}, t_{ij_2}), j_1, j_2 = 1, \dots, k_i\}$. The following model is often used for measurements $\{Y_{ij}\}$ of drug concentration: $Y_{ij} = y_2(t_{ij})/V_i + \varepsilon_{ij}$, where V_i is the volume of distribution, ε_{ij} are i.i.d. measurement errors with zero mean and variance σ_{obs}^2 . Population variability is often taken into account in PK modelling, e.g. $\gamma_i \sim N(\gamma, \Lambda)$, with $\gamma = (\theta_1, \theta_2, V)^T$. If all three sources of variability (measurement, population, stochastic) are considered, then the first-order approximation technique may be used to derive the variance-covariance matrix \mathbf{S} of the vector $\mathbf{Y} = [Y_{i1}, \dots, Y_{i,k_i}]^T$. In this paper, we consider a fixed effects model for parameters γ , so that $\mathbf{S} \approx \tilde{\mathbf{S}}/V^2 + \sigma_{obs}^2 \mathbf{I}_{k_i}$, and \mathbf{I}_{k_i} is a $(k_i \times k_i)$ identity matrix; cf. formula (8) in Anisimov et al. (2007) which included the population variability term.

3 Optimal Designs

Once the mean and variance-covariance matrix \mathbf{S} of the observed $(k \times 1)$ -vector \mathbf{Y} are derived, one can approximate the Fisher information matrix $\mu(\mathbf{x}, \vartheta)$ of a properly defined single observational unit \mathbf{x} , where $\vartheta = (\theta_1, \theta_2, V; \sigma_1^2, \sigma_2^2, \sigma_{obs}^2)$ includes all estimated parameters and, in the context of this paper, $\mathbf{x} = (t_1, t_2, \dots, t_k)$ is a $(k \times 1)$ sequence of sampling times; see formula (9) in Anisimov et al. (2007). We consider normalized designs and two types of normalization of the information matrix $\mathbf{M}_N(\vartheta) = \sum_i n_i \mu(\mathbf{x}_i, \vartheta)$ if n_i subjects are assigned to sequence \mathbf{x}_i , $\sum_i n_i = N$:

(1) Traditional normalization, or “*information per observation*”, with the normalized design ξ and the normalized information matrix $\mathbf{M}(\xi, \vartheta)$,

$$\xi = \{(\mathbf{x}_i, w_i), w_i = n_i/N, \mathbf{x}_i \in \mathbf{X}\}, \mathbf{M}(\xi, \vartheta) = \sum_i w_i \mu(\mathbf{x}_i, \vartheta). \quad (7)$$

where \mathbf{X} is a set of admissible sampling sequences.

(2) Cost-based normalization, or “*information per unit cost*”,

$$\mathbf{M}_C(\xi, \vartheta) = \mathbf{M}_N(\vartheta)/C = \sum_i p_i \tilde{\mu}(\mathbf{x}_i, \vartheta), \quad (8)$$

with $p_i = n_i c(\mathbf{x}_i)/C$, $\tilde{\mu}(\mathbf{x}, \vartheta) = \mu(\mathbf{x}, \vartheta)/c(\mathbf{x})$, $c(\mathbf{x}_i)$ is a cost of taking measurements at sequence \mathbf{x}_i and $C = \sum_i n_i c(\mathbf{x}_i)$ is the total cost.

To construct locally D-optimal designs, we use the first-order optimization algorithm with forward and backward steps; see Fedorov and Hackl (1997), Ch. 2, or Fedorov et al. (2007), Section 7.1.5. For references on optimization of sampling schemes for population models described by ODE, see Fedorov and Leonov (2007); for a discussion of various software tools, see Mentré et al. (2007).

3.1 Sampling Times and Examples of Optimal Design

As in Fedorov and Leonov (2007), for candidate sampling times we first take a uniform grid on the Y-axis (response) and then project points on the response curve

to the X-axis to obtain times \mathbf{x}_1 , or an “inverse linear” grid as in López-Fidalgo and Wong (2002); see Fig.1, top panel. Then we select splits $\tilde{\mathbf{x}}_i$ of different order:

- For some number n_1 of subjects, samples are taken at all times from \mathbf{x}_1 , i.e. $\tilde{\mathbf{x}}_1 = \mathbf{x}_1$.
- Second-order split: for n_{21} subjects samples are taken at times $\mathbf{x}_{2,1} = \{x_1, x_3, \dots, x_{2i-1}, \dots\}$, and for n_{22} subjects - at $\mathbf{x}_{2,2} = \{x_2, x_4, \dots, x_{2i}, \dots, x_n\}$ (without loss of generality we may assume that n is even). Then a generalized sequence $\tilde{\mathbf{x}}_2$ is the combination of two ‘half’-sequences $\mathbf{x}_{2,1}$ and $\mathbf{x}_{2,2}$ with, in general, different weights. We use the notation w_{21} and w_{22} [not to be confused with $w_r(t)$] for weights of standard normalized designs, and p_{21} and p_{22} for cost-based designs.
- Third-order: sampling times are $\mathbf{x}_{3,1} = \{x_1, \dots, x_{3i+1}, \dots\}$, $\mathbf{x}_{3,2} = \{x_2, \dots, x_{3i+2}, \dots\}$ and $\mathbf{x}_{3,3} = \{x_3, \dots, x_{3i}, \dots\}$ for n_{31} , n_{32} and n_{33} subjects, respectively.

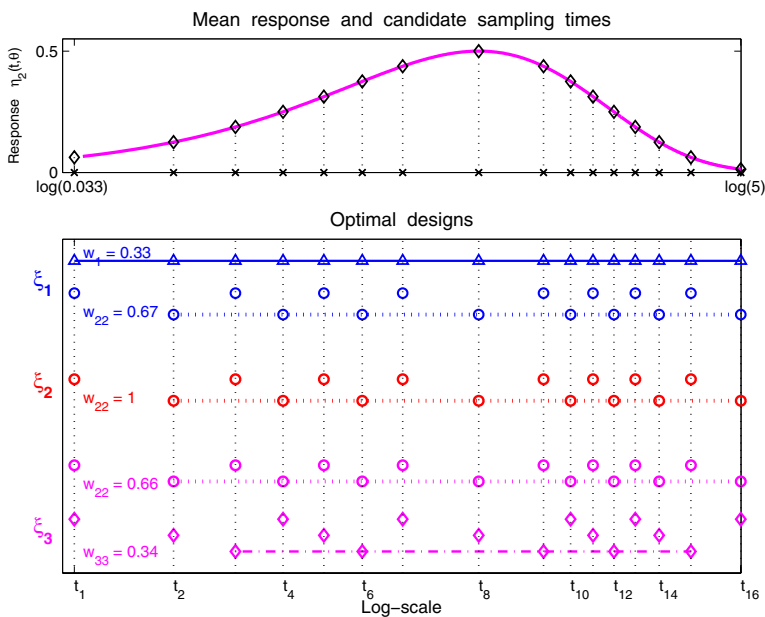


Fig. 1: Sampling times (top panel) and examples of optimal designs (bottom panel)

The information matrix in this case is $\mathbf{M}(\xi, \vartheta) = \sum_i \sum_{j=1}^i w_{ij} \mu(\mathbf{x}_{i,j}, \vartheta)$, where, to calculate $\mu(\mathbf{x}_{i,j}, \vartheta)$, we use the techniques described in Gagnon and Leonov (2005). If the cost function is selected as $c(\mathbf{x}) = C_v + C_s k$, where k is the length of a sequence \mathbf{x} , C_v is the cost of enrollment and C_s is the cost of analyzing a single sample, then

$$c(\mathbf{x}_{ij}) = C_v + C_s n/i, \quad j = 1, \dots, i, \tag{9}$$

where n is the length of the full sequence \mathbf{x}_1 . Similarly to Anisimov et al. (2007),

$\mathbf{x}_1 = [t_1, \dots, t_{16}] = [0.033, 0.069, 0.111, 0.158, 0.215, \dots, 2, 703, 3.433, 5]$, $n = 16$, and the design region is $\mathbf{X} = \{\mathbf{x}_1; \mathbf{x}_{21}, \mathbf{x}_{22}; \mathbf{x}_{31}, \mathbf{x}_{32}, \mathbf{x}_{33}\}$, i.e. all splits up to the 3rd order. For all examples we use $\dim(\vartheta) = 6$; $\theta_1 = 2$, $\theta_2 = 1$, $V = 1$; $\sigma_1 = \sigma_2 = 0.4$, $\sigma_{obs} = 0.05$; $D = 1$, and $C_v = 1$. When standard normalized designs, as in (7), are constructed, then, as expected, the optimal design is the full sequence \mathbf{x}_1 . Once costs are incorporated, as in (8), the full sequence may lose its optimality.

- If $C_s = 0.1$, then the cost-based D-optimal design is the full sequence \mathbf{x}_1 .
- If $C_s = 0.15$, then the optimal design $\xi_1 = \{(\mathbf{x}_1, p_1 = 0.43), (\mathbf{x}_{22}, p_{22} = 0.57)\}$. Recall that to calculate frequencies n_i for cost-based designs as in (8), the cost function $c(\mathbf{x}_{ij})$ has to be taken into account. According to (9), this leads to $w_{22}/w_1 = [p_{22}/p_1] \times [(1 + 0.15 \times 16)/(1 + 0.15 \times 8)] = 2.05$. Thus about 33% of subjects should be randomized to the full sequence \mathbf{x}_1 and 67% - to $\mathbf{x}_{2,2}$.
- If $C_s = 0.25$, then the optimal design ξ_2 is the sequence \mathbf{x}_{22} .
- If $C_s = 0.3$, then $\xi_3 = \{(\mathbf{x}_{22}, p_{22} = 0.73), (\mathbf{x}_{33}, p_{33} = 0.27)\}$, so that $w_{22}/w_{33} = (p_{22}/p_{33}) \times (2.5/3.4) = 2.15$, and $w_{22} \approx 0.66$, $w_{33} \approx 0.34$; see Fig.1, bottom.

It is interesting to compare the standard D-efficiency of our constructed designs (i.e. when $C_s = 0$) to the ratio of relative costs Rc_ξ : $\text{Deff}_{\xi_1} = 0.77$, $\text{Rc}_{\xi_1} = 0.76$; $\text{Deff}_{\xi_2} = 0.65$, $\text{Rc}_{\xi_2} = 0.6$; $\text{Deff}_{\xi_3} = 0.59$, $\text{Rc}_{\xi_3} = 0.53$, where

$$\text{Deff}_\xi = [|\mathbf{M}(\xi, \vartheta)| / |\mu(\mathbf{x}_1, \vartheta)|]^{1/6}, \quad \text{Rc}_\xi = \sum_{i=1}^i w_{ij} c(\mathbf{x}_{ij}) / c(\mathbf{x}_1).$$

For practical reasons, one may force equal weights for subsequences in the same split, i.e. $w_{21} = w_{22}$, $w_{31} = w_{32} = w_{33}$ etc. This restriction has essentially no effect in our example since D-efficiency of such restricted designs drops by less than 0.01. As shown in Fedorov and Leonov (2007), the use of split grids with equal weights of subsequences may be beneficial for non-compartmental analysis of clinical data.

4 Discussion

The law of conservation of mass/matter entails that the term $\theta_1 y_1(t) dt$ appears on the right-hand side of both equations in (4), though with opposite signs. By the same token it seems reasonable to consider the following analog of the second equation:

$$dy_2(t) = -dy_1(t) - \theta_2 y_2(t) dt + \sigma_2 y_2(t) dw_2(t).$$

However, using the same technique as in the proof of Lemma 1, it can be shown that the solution y_2 of such a modified system can be negative with positive probability.

Also it is worthwhile noting that the solution $y_1(t)$ in (4) is not monotonic. To overcome this practical deficiency, one may consider the following model:

$$\begin{aligned} dy_1(t) &= -\theta_1 e^{\sigma_1 w_1(t)} y_1(t) dt, \quad y_1(0) = D, \\ dy_2(t) &= \theta_1 e^{\sigma_1 w_1(t)} y_1(t) dt - \theta_2 e^{\sigma_2 w_2(t)} y_2(t) dt, \quad y_2(0) = 0. \end{aligned} \quad (10)$$

The solution of system (10) is positive and is monotonic:

$$y_1(t) = D e^{-\theta_1 \int_0^t e^{\sigma_1 w_1(s)} ds}, \quad y_2(t) = \theta_1 \cdot \int_0^t e^{-\theta_2 \int_s^t e^{\sigma_2 w_2(u)} du} \cdot e^{\sigma_1 w_1(s)} y_1(s) ds. \quad (11)$$

Though we can derive the mean of processes $y_1(t), y_2(t)$ defined in (11), the calculation of higher moments represents a formidable technical problem. In general, the presence of “white noise” w leads to continuous but non-differentiable solutions of the SDE models which may complicate their physiological interpretation. However, their mean and variance that are of practical interest are smooth functions.

Appendix: Proof of Lemma 1

(a) The derivation of (5) is a simple exercise in the application of general results of stochastic calculus and Itô's formula, see Gardiner (2003), Ch. 4.4.7.

(b) The unbiasedness of solutions y_r follows from the independence of increments $w_r(v) - w_r(u)$ and $w_r(u) - w_r(s)$, $0 \leq s < u < v$, and the equality

$$\mathbf{E} e^{a\xi} = e^{a^2/2}, \text{ for normal } \xi \sim N(0, 1) \text{ and any } a. \quad (12)$$

(c) It follows from (5) that for any $s > 0$,

$$y_2(t+s) = \alpha(t, s) y_2(t) + \beta(t, s) e^{\sigma_1 w_1(t)}, \text{ where } \alpha(t, s) = e^{-(\theta_2 + \frac{\sigma_2^2}{2})s + \sigma_2[w_2(t+s) - w_2(t)]},$$

$$\beta(t, s) = \theta_1 D e^{-(\theta_2 + \frac{\sigma_2^2}{2})(t+s)} \cdot \int_t^{t+s} e^{(\theta_2 - \theta_1 + \frac{\sigma_2^2}{2} - \frac{\sigma_1^2}{2})u + \sigma_1[w_1(u) - w_1(t)] + \sigma_2[w_2(t+s) - w_2(u)]} du.$$

Using (12), one can calculate the following conditional mathematical expectations:

$$\mathbf{E}[\alpha(t, s) | y_2(t)] = \tilde{\alpha}(s), \quad \mathbf{E}[\beta(t, s) | w_1(t), y_2(t)] = \tilde{\beta}(t, s),$$

where $\tilde{\alpha}(s) = e^{-\theta_2 s}$, $\tilde{\beta}(t, s) = (\theta_1 D / \Delta \theta) \cdot e^{-(\theta_1 + \frac{\sigma_1^2}{2})t} \cdot (e^{-\theta_1 s} - e^{-\theta_2 s})$. Thus

$$\tilde{S}(t, t+s) = \tilde{\alpha}(s) \cdot \mathbf{E} y_2^2(t) + \tilde{\beta}(t, s) \cdot \mathbf{E} e^{\sigma_1 w_1(t)} y_2(t) - \mathbf{E} y_2(t) \cdot \mathbf{E} y_2(t+s). \quad (13)$$

Next it follows from (5), (12) and the independence of increments of $w_1(t)$ that

$$\mathbf{E} e^{\sigma_1 w_1(t)} y_2(t) = \theta_1 D e^{-(\theta_2 + \frac{\sigma_2^2}{2})t} \cdot \int_0^t e^{(\theta_2 - \theta_1 + \frac{\sigma_2^2 - \sigma_1^2}{2})s} \cdot \mathbf{E} e^{\sigma_1 [w_1(t) - w_1(s)]}$$

$$\cdot \mathbf{E} e^{2\sigma_1 w_1(s)} \cdot \mathbf{E} e^{\sigma_2 [w_2(t) - w_2(s)]} ds = \frac{\theta_1 D}{\Delta \theta + \sigma_1^2} \left[e^{-(\theta_1 - \frac{3\sigma_1^2}{2})t} - e^{-(\theta_2 - \frac{\sigma_1^2}{2})t} \right]. \quad (14)$$

Similarly, to calculate $\mathbf{E} y_2^2(t)$, introduce the notation $M = \max(s, u)$ and $m = \min(s, u)$:

$$\begin{aligned} \mathbf{E}y_2^2(t) &= \theta_1^2 D^2 e^{-2(\theta_2 + \frac{\sigma_2^2}{2})t} \cdot \int_0^t \int_0^t e^{(\theta_2 - \theta_1 + \frac{\sigma_2^2 - \sigma_1^2}{2})(s+u)} \cdot \mathbf{E}e^{\sigma_1[w_1(M) - w_1(m)]} \\ &\quad \cdot \mathbf{E}e^{2\sigma_1 w_1(m)} \cdot \mathbf{E}e^{2\sigma_2[w_2(t) - w_2(M)]} \cdot \mathbf{E}e^{\sigma_2[w_2(M) - w_2(m)]} ds du = \\ &= \frac{\theta_1^2 D^2}{\Delta\theta + \sigma_1^2} \left[\frac{2e^{(-2\theta_1 + \sigma_1^2)t}}{2\Delta\theta + \sigma_1^2 - \sigma_2^2} - \frac{2e^{-(\theta_1 + \theta_2)t}}{\Delta\theta - \sigma_2^2} + \frac{2e^{(-2\theta_2 + \sigma_2^2)t}(\Delta\theta + \sigma_1^2)}{(\Delta\theta - \sigma_2^2)(2\Delta\theta + \sigma_1^2 - \sigma_2^2)} \right]. \end{aligned}$$

The expression for the function $\tilde{S}(t, t+s)$ now follows from (13)-(14). \square

References

- Anisimov, V., V. Fedorov, and S. Leonov (2007). Optimal design of pharmacokinetic studies described by stochastic differential equations. In J. López-Fidalgo, J. Rodriguez-Diaz, and B. Torsney (Eds.), *mODa 8 - Advances in Model-Oriented Design and Analysis*, pp. 9–16. Physica-Verlag, Heidelberg.
- Fedorov, V., R. Gagnon, S. Leonov, and Y. Wu (2007). Optimal design of experiments in pharmaceutical applications. In A. Dmitrienko, C. Chuang-Stein, and R. D'Agostino (Eds.), *Pharmaceutical Statistics Using SAS. A Practical Guide*, pp. 151–195. SAS Press, Cary, NC.
- Fedorov, V. and P. Hackl (1997). *Model-Oriented Design of Experiments*. Springer-Verlag, New York.
- Fedorov, V. and S. Leonov (2007). Population PK measures, their estimation and selection of sampling times. *J. Biopharmaceutical Statistics* 17(5), 919–941.
- Gagnon, R. and S. Leonov (2005). Optimal population designs for PK models with serial sampling. *J. Biopharmaceutical Statistics* 15(1), 143–163.
- Gardiner, C. (2003). *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*. Springer-Verlag, Berlin.
- López-Fidalgo, J. and W. Wong (2002). Design issues for Michaelis-Menten model. *J. Theoretical Biology* 215, 1–11.
- Mentré, F., S. Duffull, I. Gueorguieva, A. Hooker, S. Leonov, K. Ogungbenro, and S. Retout (2007). Software for optimal design in population PK/PD: a comparison. In *Abstracts of the Annual Meeting of the Population Approach Group in Europe (PAGE)*. ISSN 1871-6032, <http://www.page-meeting.org/?abstract=1179>.
- Overgaard, R., N. Jonsson, C. Tornøe, and H. Madsen (2005). Non-linear mixed-effects models with stochastic differential equations: implementation of an estimation algorithm. *J. Pharmacokinetics and Pharmacodynamics* 32(1), 85–107.
- Picchini, U., S. Ditlevsen, and A. De Gaetano (2006). Modeling the euglycemic hyperinsulinemic clamp by stochastic differential equations. *J. Theoretical Biology* 53, 771–796.
- Tornøe, C., J. Jacobsen, O. Pedersen, T. Hansen, and H. Madsen (2004). Grey-box modelling of pharmacokinetic/pharmacodynamic systems. *J. Pharmacokinetics and Pharmacodynamics* 31(5), 401–417.

On Testing Hypotheses in Response-Adaptive Designs Targeting the Best Treatment

Nancy Flournoy, Caterina May, Jose A. Moler and Fernando Plo

Abstract We consider a sequential, response-adaptive design for clinical trials which is characterized by the fact that it assigns patients to the best treatment with a probability converging to one. This property is optimal from an ethical point of view; in this paper we analyze some inferential problems related to the design. In particular, we want to establish, by means of a test of hypothesis, which treatment is superior, in the sense that it has greater mean response. Together with the natural generalization of the classical t-statistic, we introduce a statistic based on the probability of assigning patients to a treatment conditional on past observations. Theoretical properties of the tests are studied, together with numerical evaluations of the power for dichotomous responses.

1 Introduction

Response-adaptive designs have been the subject of increasing attention by many researchers in the area of sequential procedures for several years. This is due to the fact that they permit skewing the allocation probabilities during the experiment, on the bases of the accrued information (previous allocations and/or previous re-

Nancy Flournoy

University of Missouri, 146 Middlebush Hall Columbia, MO 65211-6100, USA, e-mail: flournoyn@missouri.edu

Caterina May

Università degli Studi del Piemonte Orientale, Via Perrone 18, 28100 Novara, Italy, e-mail: caterina.may@eco.unipmn.it

Jose A. Moler

Universidad Publica de Navarra, Campus de Arrosadia s/n. 31006-Pamplona, Spain, e-mail: jmoler@unavarra.es

Fernando Plo

Universidad de Zaragoza, Pedro Cerbuna, 12 50009-Zaragoza, Spain, e-mail: fplo@unizar.es

sponses). In a clinical trial, this characteristic is fundamental from an ethical point of view; moreover, these procedures present advantages also in other areas of investigation, including industrial and economic contexts. A recent review of a broad class of such designs has been provided in Hu and Rosenberger (2006).

Let us consider an experiment conducted to compare two or more treatments. Patients enter the experiment sequentially and they are randomly allocated to one of the treatments, according to the adopted design. Most response-adaptive procedures considered in the literature allocate patients with a probability converging to a target allocation $\rho \in (0, 1)$, which can depend on the unknown parameters and more recent procedures are chosen based on some optimality properties. Admissible target allocations in $(0, 1)$, both from inferential and ethical points of view, have been recently discussed in Baldi Antognini and Giovagnoli (2009).

In addition to these procedures, a class of response-adaptive designs based on a randomly reinforced urn (*RRU*) which targets the best treatment with a probability converging to $\rho = 1$ has been studied, from the early papers of Durham and Yu (1990), Li, Durham, and Flournoy (1996), Durham, Flournoy, and Li (1998) and then, among others, in Muliere, Paganoni, and Secchi (2006) and May, Paganoni, and Secchi (2007), up to the recent work of May and Flournoy (2009). This property is very desirable from an ethical point of view.

For tests of hypotheses to compare the mean responses between treatments, the Wald t-test statistic is usually considered in the literature. In particular, the power of the test between different response-adaptive designs targeting a $\rho \in (0, 1)$ has been compared, for instance, in Hu and Rosenberger (2003) or Zhang and Rosenberger (2006). Such a comparative study cannot be applied to the *RRU*-designs, where the target proportion is $\rho = 1$, since the t-test statistic for *RRU* designs has different asymptotic properties, as shown in May and Flournoy (2009).

In this work, we simulate the number of patients assigned to the superior treatment in a finite-sample experiment modelled by a *RRU*-design, showing the ethical advantage in comparison to some other designs proposed in the literature. Then we analyze the power of the two sample t-test in a *RRU*-design and we propose also a different test statistic based on the proportion of black balls contained in the urn. We study the asymptotic power of each test and we perform numerical evaluations of the power assuming fixed sample sizes. We perform our numerical simulation for the basic case where treatments have dichotomous (success/failure) responses, but the analysis may be extended to a more general situation. In the final Discussion, we present further developments of this promising area of research.

2 *RRU*-Designs and Test of Hypotheses

Consider an experiment conducted to compare two treatments, say *B* and *W*. An urn contains initially *b* balls of colour *B* and *w* balls of colour *W*. At each time $n = 1, 2, \dots$ a patient enters the experiment and is allocated to treatment *B* or *W*, according to the colour of a ball randomly sampled from the urn. The response of

the patient is then observed before the arrival of the $(n + 1)$ -th patient. The ball extracted is replaced in the urn along with a random number of balls that are of the same colour as the ball that has been extracted. This random number is equal to the response of the patient to the treatment or it is a suitable function of it. (For further details on this model see May and Flournoy 2009).

Let us denote by δ_n the indicator of the event representing the extraction of a ball of colour B , that is, the assignment of treatment B . The response of the patient at time n is a random variable $Y_B(n)$ with law \mathcal{L}_B if $\delta_n = 1$, or a random variable $Y_W(n)$ with law \mathcal{L}_W if $\delta_n = 0$. Moreover, we denote by

$$N_B(n) = \sum_{i=1}^n \delta_i \quad N_W(n) = \sum_{i=1}^n (1 - \delta_i)$$

the random numbers of patients assigned until time n to B and W , respectively.

Let Z_n be the proportion of black balls contained in the urn at time n ; Muliere, Paganoni, and Secchi (2006) generalizes Durham, Flournoy, and Li (1998) in proving that, if treatment B has a higher mean response than W , Z_n converges almost surely to one. This means that, when treatment B is more favourable, the probability of allocating a patient to B converges to one, and this property is very desirable from an ethical point of view. It has been also proved that, when the treatments have the same mean responses, the process Z_n converges almost surely to a random variable without point masses in $[0, 1]$.

In order to compare the performance of the treatments, we consider the hypothesis test on the mean responses $m_B = \int y \mathcal{L}_B(dy)$ and $m_W = \int y \mathcal{L}_W(dy)$:

$$H_0 : m_B = m_W, \quad \text{versus} \quad H_1 : m_B > m_W. \tag{1}$$

Let us now discuss two different test statistics to perform the test.

2.1 Test Based on the t-Statistic

In May and Flournoy (2009) the asymptotic properties of the following natural extension of the t-statistic

$$\zeta_0(n) = \frac{\hat{Y}_B(n) - \hat{Y}_W(n)}{\sqrt{\frac{\hat{\sigma}_B^2}{N_B(n)} + \frac{\hat{\sigma}_W^2}{N_W(n)}}, \tag{2}$$

where

$$\hat{Y}_B(n) = \frac{\sum_{i=1}^n \delta_i Y_B(i)}{N_B(n)}, \quad \hat{\sigma}_B^2 = \frac{\sum_{i=1}^n \delta_i (Y_B(i) - \hat{Y}_B(n))^2}{N_B(n)}$$

and

$$\hat{Y}_W(n) = \frac{\sum_{i=1}^n (1 - \delta_i) Y_W(i)}{N_W(n)}, \quad \hat{\sigma}_W^2 = \frac{\sum_{i=1}^n (1 - \delta_i) (Y_W(i) - \hat{Y}_W(n))^2}{N_W(n)}$$

have been studied. In particular, it is proved that $\zeta_0(n)$ converges, under the null hypotheses, to a standard normal distribution, while under the alternative hypothesis it converges conditionally on η to a normal distribution with unit variance and mean equal to $\sqrt{n^{m_W/m_B}} \eta \frac{m_B - m_W}{\sigma_W}$, where η is the positive square root of the random variable η^2 is defined by

$$\eta^2 = \lim_{n \rightarrow \infty} \frac{N_W(n)}{n^{m_W/m_B}}, \quad a.s.$$

Hence, if we fix an asymptotic significance level α (and denoting by $z_{1-\alpha}$ the quantile of order $1 - \alpha$ of a standard normal distribution), we can then consider the following critical region C_α^1 :

$$C_\alpha^1 = \{\zeta_0(N) > z_{1-\alpha}\}.$$

Moreover, the power of the test, $1 - \beta_1$, can be approximated, for a large number N of patients, by

$$1 - \beta_1 = P\left(\mathcal{N} + N^{m_W/(2m_B)} \eta \frac{m_B - m_W}{m_W(1 - m_W)} > z_{1-\alpha}\right), \quad (3)$$

where \mathcal{N} is a standard normal random variable independent of η^2 (see Proof of Corollary 3 in May and Flournoy 2009), and $N = N_B(N) + N_W(N)$.

2.2 Test Based on the ‘Proportion of Black Balls’ Statistic

The proportion of black balls Z_n has relevant information about the performance of both treatments. In a clinical trial with N patients, the behaviour of the trajectories of the stochastic process $\{Z_n\}$ should be reflected in the value of Z_N . In fact, for an asymptotic significance level α we can consider the critical region

$$C_\alpha^2 = \{Z_N > c_{1-\alpha}\},$$

where $c_{1-\alpha}$ is the quantile of order $1 - \alpha$ for the limit Z_∞ of the process Z_n . From Corollary 2 in May and Flournoy (2009) we are also able to approximate the value of the power $1 - \beta_2 = P(Z_N > c_{1-\alpha} | H_1)$ in this case. In fact, from

$$\lim_{n \rightarrow +\infty} \frac{1 - Z_n}{n^{m_W/m_B - 1}} = \frac{m_W}{m_B} \eta^2, \quad a.s., \quad (4)$$

we can approximate, for large N ,

$$1 - \beta_2 = P\left(\eta^2 < (1 - c_{1-\alpha}) \frac{m_B}{m_W} N^{1 - m_W/m_B}\right). \quad (5)$$

We remark that the distribution of η^2 is in general unknown, but we can simulate it by using (4) (a simulation study for continuous responses has been considered in May, Paganoni, and Secchi 2007). Moreover, we can compare the two test-statistics in terms of the rates with which their power functions converge to one, as the number of patients N increases to infinity. In fact, if we compare $1 - \beta_1$ and $1 - \beta_2$, we note that the rate is the same when the ratio m_W/m_B is equal to $2/3$, while for $m_W/m_B < 2/3$ the power of the proportion-test converges faster.

3 Numerical Results

In this section several simulation studies are reported. We focus on a clinical trial with two treatments, B and W , with dichotomous responses (success/failure). Success probabilities are denoted p_B and p_W respectively. This obviously means that, according to the notation of Section 2, $p_B = m_B$ and $p_W = m_W$. In the sequel, the sample size, i.e., the number of patients in the experiment, is taken to be $N = 400$. As in practical situations the best treatment is unknown, a balanced initial urn composition, $b = w$, is chosen. When a treatment is successful, we reinforce the urn with one ball of the colour associated with this treatment.

First, we compare the ethical performance of the *RRU*-design with other designs for a finite number N of patients. In fact, we simulate the number of patients $N_B(n)$ allocated to the best treatment and the number of failures in the trial for the complete randomization design (*CR*), the Efron biased coin design (*Efron*) (see Chapter 3 in Rosenberger and Lachin 2002), the randomized play-the-winner rule (*PTW*), the drop the loser rule (*DL*) (see Ivanova and Flournoy 2001 and Ivanova 2003) and the *RRU*-design when it is initialized with $b=w=1$ and with $b=w=3$.

Figure 1 displays the box-plots of the random variable $N_B(400)$ and the total number of failures after 100 repetitions for each design in a clinical trial for which $p_B = 0.8$ and $p_W = 0.4$. We can observe that equal allocation is obtained on average with the non-response adaptive designs and the variance is smaller with the *Efron* design. However, response-adaptive designs skew the allocation of patients towards the best treatment and diminish the number of failures. The *RRU*-design provides the empirical distribution of $N_B(400)$ most skewed towards the best treatment and the empirical distribution of the number of failures most skewed towards zero. These good properties imply a greater variability. Increasing the initial number of balls (from $b=w=1$ to $b=w=3$) reduces both variability and skewness, but even in this case the *RRU*-design is competitive with the other designs from an ethical point of view. Note, also, that the variability in this case would decrease for smaller values of p_W/p_B .

We remark also that the inverse relation between the power of the t-statistic and the variability of allocations is proved in Hu and Rosenberger (2003) when the target allocation is $\rho \in (0, 1)$, but this is not the case for the *RRU*-design. Simulation studies about the performance of the t-test in response-adaptive designs targeting a $\rho \in (0, 1)$ can be found, for instance, in Ivanova (2003).

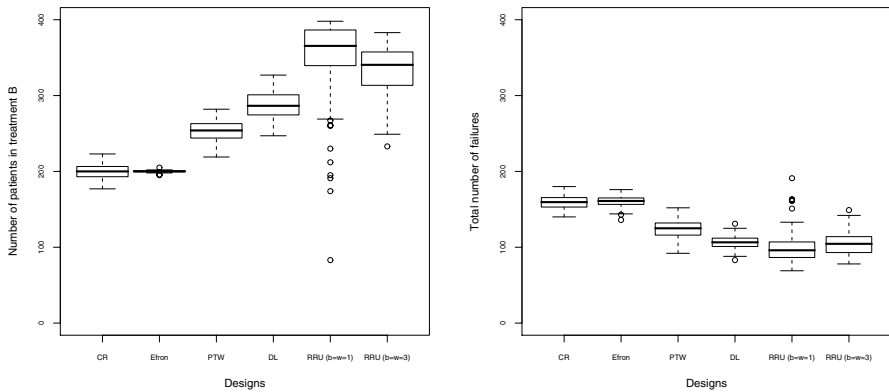


Fig. 1: Box-plots of $N_B(400)$ and of the total number of failures for some different designs, when $p_B = 0.8$ and $p_W = 0.4$. The RRU -designs have been simulated for $b = w = 1$ and $b = w = 3$.

In Table 1, a comparative simulation study of the power of the Z_N -test and the t -test is carried out when the RRU -design is applied. This study shows two scenarios which are reflected in two rows: the top panels result from the RRU -design initialized with $b = w = 1$, and the bottom panels with $b = w = 3$. In each row, the left panel gives results for the Z_N -test, and the right panel for the t -test. For each combination of success probabilities given in the tables, we simulate the clinical trial 1000 times and obtain the proportion of simulated clinical trials for which the corresponding statistic is in the critical region C_α^i , $i = 1, 2$. In the dichotomous situation here considered, when $p_B = p_W$ the limit of Z_n is a $Beta(b, w)$ distribution, so $c_{1-\alpha}$ is the $(1 - \alpha)$ -percentile of the corresponding Beta distribution. The proportion obtained is the simulated significance level of the test in the diagonal elements of each panel and the simulated power of the test is reflected in the off-diagonal terms.

Observe that for $b = w = 1$ the empirical power of the t -test is higher than the empirical power of the Z_n -test. However, the significance level of the t -test is far from 0.05 whereas this is not the case for the Z_n -test. The power of the t -test is inflated possibly because the normal approximation is not good enough when the number of patients allocated to the worst treatment is small. On the other hand, the Beta approximation for the Z_n -test seems adequate to estimate the significance level but the empirical power of this test is too small when the success probabilities are very close. When $b = w = 3$ the RRU -design performs worse than $b = w = 1$ from an ethical point of view, but it is still competitive with the other designs. The Z_n -test improves its empirical power but the t -test increases its empirical power too, and its empirical significance level approaches 0.05 except for extremal success probabilities.

Table 1: Each row provides two tables with the power of the Z_N statistic (on the left) and the power of the t-test (on the right). The first block is obtained with the initial condition $b = w = 1$ and the second with the initial condition $b = w = 3$. The empirical level of significance is on the diagonal.
Initial condition: $b = w = 1$

$p_B p_W$	0.1	0.2	0.3	0.5	0.7	0.8	0.9	$p_B p_W$	0.1	0.2	0.3	0.5	0.7	0.8	0.9
0.1	0.05	-	-	-	-	-	-	0.1	0.09	-	-	-	-	-	-
0.2	0.30	0.05	-	-	-	-	-	0.2	0.80	0.11	-	-	-	-	-
0.3	0.63	0.19	0.03	-	-	-	-	0.3	0.98	0.64	0.09	-	-	-	-
0.5	0.95	0.61	0.29	0.05	-	-	-	0.5	1.00	0.98	0.92	0.08	-	-	-
0.7	1.00	0.91	0.63	0.15	0.05	-	-	0.7	1.00	0.99	0.99	0.89	0.06	-	-
0.8	1.00	0.94	0.78	0.34	0.07	0.04	-	0.8	1.00	0.99	0.99	0.96	0.59	0.06	-
0.9	1.00	0.98	0.85	0.40	0.15	0.10	0.04	0.9	1.00	1.00	0.99	0.98	0.92	0.67	0.08

Initial condition: $b = w = 3$

$p_B p_W$	0.1	0.2	0.3	0.5	0.7	0.8	0.9	$p_B p_W$	0.1	0.2	0.3	0.5	0.7	0.8	0.9
0.1	0.03	-	-	-	-	-	-	0.1	0.06	-	-	-	-	-	-
0.2	0.56	0.06	-	-	-	-	-	0.2	0.84	0.08	-	-	-	-	-
0.3	0.95	0.32	0.05	-	-	-	-	0.3	1.00	0.71	0.05	-	-	-	-
0.5	1.00	0.94	0.58	0.05	-	-	-	0.5	1.00	0.94	0.58	0.05	-	-	-
0.7	1.00	1.00	0.95	0.39	0.04	-	-	0.7	1.00	1.00	1.00	0.98	0.06	-	-
0.8	1.00	1.00	0.98	0.65	0.12	0.04	-	0.8	1.00	1.00	1.00	1.00	0.70	0.05	-
0.9	1.00	1.00	1.00	0.81	0.30	0.13	0.04	0.9	1.00	1.00	1.00	1.00	1.00	0.83	0.04

4 Discussion and Further Developments

In this paper we have investigated hypothesis tests for RRU -designs, which are characterized by optimality from an ethical point of view. They can't be approached with the usual method for comparing designs having target allocations ρ in $(0, 1)$. We have considered the well-known two sample t-test statistic and we have introduced a new test statistic, the proportion of black balls. The numerical results on the power of the tests suggest that a combination of the two tests could be a useful way to obtain more efficient statistical inference. We leave this as a further research topic.

Comparisons between the performance of different designs and between different statistics are influenced by the choice of the initial parameters b, w and by the sample size N . We think that it will be necessary to develop a comparison study focusing on the initial parameters and on different sample sizes. Moreover, the study here presented needs to be extended to the case of general outcomes (even continuous) and to the case of K treatments.

We also believe that a further improvement of inference in response-adaptive designs could be obtained by constructing sequential tests (as is also claimed in Rosenberger 2002); we hope to deal in future with this promising area of research.

Acknowledgements Thanks to Piercesare Secchi and Anna Maria Paganoni for the ideas and discussions which have stimulated this work.

References

- Baldi Antognini, A. and A. Giovagnoli (2009). Ethics and inference in binary clinical trials: admissible allocations and response-adaptive randomization. *Preprint*.
- Durham, S. D., N. Flournoy, and W. Li (1998). A sequential design for maximizing the probability of a favourable response. *Canadian Journal of Statistics* 26, 479–495.
- Durham, S. D. and K. F. Yu (1990). Randomized play-the leader rules for sequential sampling from two populations. *Probability in Engineering and Information Science* 26, 355–367.
- Hu, F. and W. Rosenberger (2003). Optimality, variability, power: evaluating response-adaptive randomization procedures for treatment comparisons. *Journal of the American Statistical Association* 98, 671–678.
- Hu, F. and W. F. Rosenberger (2006). *The Theory of Response-Adaptive Randomization in Clinical Trials*. Wiley Series in Probability and Statistics. New York: Wiley.
- Ivanova, A. (2003). A play-the-winner-type urn design with reduced variability. *Metrika* 58 1–13.
- Ivanova, A and N. Flournoy (2001). A birth and death urn for ternary outcomes: stochastic processes applied to urn models. In: Charalambides C, Koutras MV, Balakrishnan (eds) *Probability and Statistical Models with Applications.*, 583–600. Boca Raton: Chapman and Hall/CRC.
- Li, W., S. D. Durham, and N. Flournoy (1996). Randomized Polya urn designs. *Proceedings of the Biometric Section of the American Statistical Association*, 166–170.
- May, C. and N. Flournoy, N. (2009). Asymptotics in response-adaptive designs generated by a two-color, randomly reinforced urn. *The Annals of Statistics* 32, 1058–1078.
- May, C., A. Paganoni, and P. Secchi (2007). Response-adaptive designs targeting the best treatment for clinical trials with continuous responses. In *S.Co.2007 Fifth Conference - Complex Models and Computational intensive methods for estimation and prediction - Book of short papers*, pp. 326–331. Cluep.
- Muliere, P., A. Paganoni, and P. Secchi (2006). A randomly reinforced urn. *Journal of Statistical Planning and Inference* 136, 1853–1874.
- Rosenberger, W. F. (2002). Randomized urn models and sequential design. *Sequential Analysis* 21, 1–28.
- Rosenberger, W. F. and J. M. Lachin (2002). *Randomization in Clinical Trials*. New York: Wiley.
- Zhang, L. and W. F. Rosenberger (2006). Response-adaptive randomization for clinical trials with continuous responses. *Biometrics* 62, 562–569.

Towards Gaussian Process-based Optimization with Finite Time Horizon

David Ginsbourger and Rodolphe Le Riche

Abstract During the last decade, Kriging-based sequential optimization algorithms have become standard methods in computer experiments. These algorithms rely on the iterative maximization of sampling criteria such as the *Expected Improvement (EI)*, which takes advantage of Kriging conditional distributions to make an explicit trade-off between promising and uncertain points in the search space. We have recently worked on a multipoint *EI* criterion meant to choose simultaneously several points for synchronous parallel computation. The results presented in this article concern sequential procedures with a fixed number of iterations. We show that maximizing the usual *EI* at each iteration is suboptimal. In essence, the latter amounts to considering the current iteration as the last one. This work formulates the problem of optimal strategy for finite horizon sequential optimization, provides the solution to this problem in terms of a new multipoint *EI*, and illustrates the suboptimality of maximizing the 1-point *EI* at each iteration on the basis of a first counter-example.

1 Introduction

The Gaussian Process (GP) has become a major tool in *metamodeling* for computer experiments (Rasmussen and Williams 2006). When studying a simulator with scalar output, $y : \mathbf{x} \in D \subset \mathbb{R}^d \rightarrow y(\mathbf{x}) \in \mathbb{R}$, GP metamodeling consists in assuming that y is one path of a GP Y . The main focus of this paper is on metamodel-based optimization with finite time horizon. In GP-based optimization, it is common to sequentially enrich the current Design of the Experiment (DoE) $\mathbf{X} = \{\mathbf{x}^1, \dots, \mathbf{x}^n\} \in D^n$ ($n \in \mathbb{N}^*$) —denoted by $\mathbf{X} = \mathbf{X}^0$ and $n = n_0$ in the initial state— by maximizing a probabilistic criterion of interest, updating the GP model, and iterating. The *Expected Improvement (EI)* is now one of the most popular GP-based optimization criteria:

$$EI(\mathbf{x}) = \mathbb{E} [(\min\{Y(\mathbf{X})\} - Y\{\mathbf{x}\})^+ | Y(\mathbf{X}) = \mathbf{Y}] = \mathbb{E}[I(\mathbf{x})|A], \quad (1)$$

Dr. David Ginsbourger
CHYN, Emile-Argand 11, CH-2009 Neuchâtel, Switzerland, e-mail: david.ginsbourger@unine.ch

Dr. Rodolphe Le Riche
CROCUS, Ecole des Mines, 158 cours Fauriel, Saint-Etienne, France, e-mail: leriche@emse.fr

where $I(\mathbf{x}) := [\min\{Y(\mathbf{X})\} - Y(\mathbf{x})]^+ = \max[0, \min\{Y(\mathbf{X})\} - Y(\mathbf{x})]$ is the random *improvement* at \mathbf{x} , and the event A sums up all available observations. EI is appreciated for trading off exploitation of known information and exploration of unvisited search space areas. Furthermore, EI is known in closed form (Jones, Schonlau, and Welch 1998), which allows fast evaluations, and calculation of its derivatives. Such a criterion, though updated by integrating new data, is typically considered at each iteration without structural change. In fact, in EI algorithms like EGO , the point visited at the j^{th} iteration is determined by maximizing a conditional expectation:

Algorithm 1 EI algorithm with known Kriging parameters and $r \in \mathbb{N}^*$ iterations

- 1: for $j \leftarrow 1, r$ do
 - 2: $A_{j-1} = \{Y(\mathbf{x}^1) = y(\mathbf{x}^1), \dots, Y(\mathbf{x}^{n+j-1}) = y(\mathbf{x}^{n+j-1})\}$
 - 3: $\mathbf{x}^{n+j} = \arg \max_{\mathbf{x} \in D} \{\mathbb{E}[I(\mathbf{x})|A_{j-1}]\}$
 - 4: end for
-

Example 1. We consider an objective function defined by $y_1 : x \in [0, 1] \rightarrow y_1(x) = \sin(10x + 1)/(1 + x) + 2 \cos(5x)x^4 \in \mathbb{R}$, where $D = [0, 1]$. Fig. 1 illustrates y_1 and its actual minimizer, the initial DoE $\mathbf{X}^0 = \{0.1, 0.2, 0.85\}$, as well as the associated 1-point and 2-point EI functions. Simple Kriging is performed using a Matern covariance ($\nu = \frac{3}{2}$; see Stein (1999) for details), with a unit variance and a range of $\frac{0.3}{\sqrt{3}}$. Further comments can be found in the caption of Fig. 1 and in Section 3.

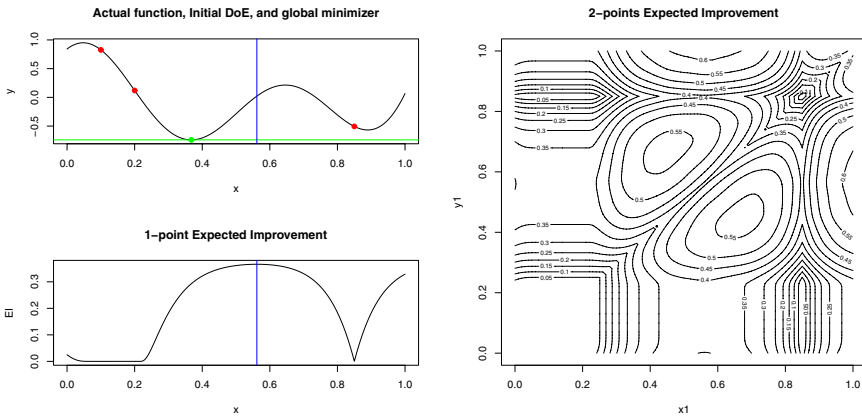


Fig. 1: y_1 (upper left) with its global minimizer (horizontal line, with a dot at the minimum) and the design \mathbf{X}^0 (three dots above the former horizontal line); 1-point EI and 2-point EI criteria (lower left, and right) corresponding to the Kriging model of Example 1. The vertical lines show the 1-point EI maximizer, $\mathbf{x} \approx 0.57$. The maximum of the 2-point EI is reached with one point as previously, and one point at the boundary point $\mathbf{x} = 1$.

2 What is a Strategy and How to Measure its Performance?

Sequential Deterministic Strategies for Optimization with Finite Horizon. Assume that one has a budget of r evaluations after having evaluated y at an arbitrary n -point design, \mathbf{X} . One step of a sequential strategy consists in looking for the next point where to evaluate y , say \mathbf{x}^{n+1} . In some sampling procedures like crude Monte Carlo, \mathbf{x}^{n+1} is determined without taking (\mathbf{X}, \mathbf{Y}) into account. However, in the considered case of adaptive strategies, \mathbf{x}^{n+1} is determined on the basis of the available information. Furthermore, we restrict ourselves to the case of deterministic strategies, i.e. where \mathbf{x}^{n+1} only depends on the past and does not involve any random operator (like mutations in genetic algorithms). So \mathbf{x}^{n+1} is defined as a function:

$$s_1 : (\mathbf{X}, \mathbf{Y}) \in (D \times \mathbb{R})^n \longrightarrow \mathbf{x}^{n+1} = s_1(\mathbf{X}, \mathbf{Y}) \in D. \quad (2)$$

For instance, $s_1(\cdot)$ is defined in Algorithm 1 as $\arg \max_{\mathbf{x} \in D} \mathbb{E}[I(\mathbf{x})|A_0]$. To use the previous notation, one can similarly define a function s_j for every $j \in [2, r]$:

Definition 1. We call a *deterministic strategy with horizon r* ($r \in \mathbb{N}^*$) any finite sequence $\mathcal{S} = (s_j)_{j \in [1, r]}$ of measurable functions $s_j(\cdot) : (D \times \mathbb{R})^{n_0+j-1} \longrightarrow D$ ($j \in [1, r]$), and denote by \mathbb{S}_r the space of such \mathcal{S} .

In Algorithm 1, the s'_j 's are implicitly taken as $\arg \max_{\mathbf{x} \in D} \mathbb{E}[I(\mathbf{x})|\mathbf{X}^{j-1}, Y(\mathbf{X}^{j-1})]$ for all $j \in [2, r]$, where $\mathbf{X}^{j-1} = \mathbf{X}^0 \cup \{\mathbf{x}^{n_0+1}, \dots, \mathbf{x}^{n_0+j-1}\}$ and $\mathbf{Y}^{j-1} = Y(\mathbf{X}^{j-1})$ denote the augmented design and the vector of observations. The only change in the criteria of such an *EI* algorithm is the updated information. Here we consider more general strategies, where the s'_j 's may be subject to structural changes at each iteration.

After r function evaluations, it is possible to evaluate the success of $\mathcal{S} \in \mathbb{S}_r$ by comparing the initial best response, $m_0 := \min(y(\mathbf{X}^0))$, to the best response observed during the additional runs, $m_{1,r} := \min(y(\mathbf{x}^{n_0+1}), \dots, y(\mathbf{x}^{n_0+r}))$. The corresponding performance measure $(m_0 - m_{1,r})^+$ can be written in terms of multipoint improvement (see e.g. Schonlau 1997 or Ginsbourger, Le Riche, and Carraro 2010):

Definition 2. The improvement of $\mathcal{S} \in \mathbb{S}_r$ seen from the initial state is defined as

$$i^0(\mathcal{S}) := (\min\{y(\mathbf{X}^0)\} - \min\{y\{s_1(\mathbf{X}^0, \mathbf{Y}^0)\}, \dots, y\{s_r(\mathbf{X}^{r-1}, \mathbf{Y}^{r-1})\}\})^+. \quad (3)$$

More generally ($0 \leq j \leq r$), $i^j(\mathcal{S}) := (m_j - m_{(j+1):r})^+$ stands for the cumulative improvement obtained between the j^{th} step and the end of strategy \mathcal{S} .

Our purpose is to find strategies that produce the largest possible *a posteriori* improvement in a given number of iterations. However, evaluating $i^0(\mathcal{S})$ obviously requires already knowing $(\mathbf{X}^r, \mathbf{Y}^r)$, i.e. being at the end of the algorithm. So we need a criterion that takes a strategy $\mathcal{S} = (s_j)_{j \in [1, r]}$ as argument while not explicitly depending on the design points and response values to be observed during the

algorithm. This is what we will propose in the next subsection with the adaptation of the *Expected Improvement* criterion to *sequential strategies*.

EI of a Finite Time Sequential Strategy. The quantities \mathbf{X}^j and \mathbf{Y}^j are deterministic for an observer having collected information at or after the j^{th} iteration. We now detail the case where the latter are seen from the past of iteration j , and hence inherit from an epistemic random nature: \mathcal{X}^{n_0+j} denotes the random variable corresponding to \mathbf{x}^{n_0+j} , and $\mathbb{X}^j = \mathbf{X}^0 \cup \{\mathcal{X}^{n_0+1}, \dots, \mathcal{X}^{n_0+j}\}$ the random design corresponding to \mathbf{X}^{n_0+j} with known \mathbf{X}^0 (in all cases, $j \in [1, r]$). Similarly, $\mathbb{Y}^j = \mathbf{Y}^0 \cup \{Y(\mathcal{X}^{n_0+1}), \dots, Y(\mathcal{X}^{n_0+j})\}$. Note that $\mathcal{X}^{n_0+1} = s_1(\mathbf{X}^0, \mathbf{Y}^0)$ is non-random. However, $\mathcal{X}^{n_0+2} = s_2(\mathbb{X}^1, \mathbb{Y}^1)$ is random, and is more precisely $\sigma\{Y(\mathbf{X}^{n_0+1})\}$ - or $\sigma(\mathbb{Y}^1)$ -measurable. More generally, each \mathcal{X}^{n_0+j} is a $\sigma(\mathbb{Y}^{j-1})$ -measurable random variable. We are now ready to introduce the random variables

$$I^j(\mathcal{S}) = (\min\{Y(\mathbf{X}^j)\} - \min\{Y\{s_{j+1}(\mathbf{X}^0, \mathbf{Y}^0)\}, \dots, Y\{s_r(\mathbb{X}^{r-1}, \mathbb{Y}^{r-1})\}\})^+, \quad (4)$$

where $0 \leq j \leq r-1$. Finally, let $A_j = \{\mathbb{X}^j = \mathbf{X}^j, Y(\mathbf{X}^j) = \mathbf{Y}^j\}$ ($0 \leq j \leq r$) denote the information available right after the calculation of \mathcal{X}^{n_0+j} and $y(\mathcal{X}^{n_0+j})$.

Definition 3. The Expected Improvement of a strategy $\mathcal{S} = (s_j)_{j \in [1, r]}$ seen from its initial state is given by $EI^0(\mathcal{S}) = \mathbb{E}[I^0(\mathcal{S})|A_0]$

3 Towards Deriving the Optimal Finite Time Strategy

We restrict ourselves here to the case where D is a compact subset of \mathbb{R}^d , and assume for convenience that each considered $\mathbb{E}[I^j(\mathbf{x}, \dots)|A_j]$ ($0 \leq j \leq r$) possesses one unique global maximizer over D . Under these working assumptions, we denote by \mathcal{P}_r the problem: find $\mathcal{S}_r^* = (s_j^*)_{j \in [1, r]}$ maximizing EI^0 . Let us first write a trivial property of strategies with horizon 1 which will nevertheless be useful in the sequel:

Lemma 1. *The solution of \mathcal{P}_1 is given by $s_1^*(\mathbf{X}^0, \mathbf{Y}^0) = \arg \max_{\mathbf{x} \in D} \mathbb{E}[I^0(\mathbf{x})|A_0]$.*

Proof. Follows directly from the definition of \mathcal{P}_1 .

Lemma 2. $\forall (a, b, c) \in \mathbb{R}^3, (a - \min(b, c))^+ = (a - b)^+ + (\min(a, b) - c)^+.$

Proof. If $a = \min(a, b, c)$, then both left and right terms are 0. If $b = \min(a, b, c)$, both terms equal $(a - b)$ since $\min(b, c) = b$ and $(\min(a, b) - c)^+ = 0$. Finally, if $c = \min(a, b, c)$, the left term equals $(a - c)$ and the right one equals $0 + (a - c)$ if $b \geq a$ and $(a - b) + (b - c) = (a - c)$ else. \square

Theorem 1. *In \mathcal{P}_r , choosing \mathbf{x}^{n_0+r} amounts to maximizing $\mathbb{E}[I^{r-1}(\cdot)|A_{r-1}]$.*

Proof. After $r-1$ iterations, $\{\mathbf{X}^{r-1}, \mathbf{Y}^{r-1}\}$ is known, and the maximization of EI over \mathbb{S}_r reduces to a simpler problem over \mathbb{S}_1 . Noting $M_0 = \min(\{Y(\mathbf{X}^0)\})$ and $M_{1:(r-1)} = \min\{Y(\mathcal{X}^{n_0+1}), \dots, Y(\mathcal{X}^{n_0+r-1})\}$, we have:

$$\begin{aligned} \mathbf{x}^{n_0+r} &= \arg \max_{\mathbf{x} \in D} \mathbb{E}[(M_0 - \min\{Y(\mathcal{X}^{n_0+1}), \dots, Y(\mathcal{X}^{n_0+r-1}), Y(\mathbf{x})\})^+ | A_{r-1}] \\ &= \arg \max_{\mathbf{x} \in D} \mathbb{E}[(M_0 - \min\{M_{1:(r-1)}, Y(\mathbf{x})\})^+ | A_{r-1}]. \end{aligned} \quad (5)$$

We then use Lemma 2 with $a = \min\{Y(\mathbf{X}^0)\}$, $b = M_{1:(r-1)}$, $c = Y(\mathbf{x})$ and obtain:

$$\mathbb{E}[(M_0 - \min\{M_{1:(r-1)}, Y(\mathbf{x})\})^+ | A_{r-1}] = (m_0 - m_{1:(r-1)})^+ + \mathbb{E}[I^{r-1}(\mathbf{x}) | A_{r-1}]. \quad (6)$$

As $(m_0 - m_{1:(r-1)})^+$ does not depend on \mathbf{x} , \mathbf{x}^{n_0+r} maximizes $\mathbb{E}[I^{r-1}(\mathbf{x}) | A_{r-1}]$. \square

Corollary 1. *The solution (s_1^*, \dots, s_r^*) of \mathcal{P}_r is given by the following recursion:*

$$\begin{cases} s_r^*(\mathbf{X}^{r-1}, \mathbf{Y}^{r-1}) = \arg \max_{\mathbf{x} \in D} \mathbb{E}[I^{r-1}(\mathbf{x}) | A_{r-1}], \\ s_{r-1}^*(\mathbf{X}^{r-2}, \mathbf{Y}^{r-2}) = \arg \max_{\mathbf{x} \in D} \mathbb{E}[I^{r-2}(\mathbf{x}, s_r^*\{\mathbb{X}^{r-1}(\mathbf{x}), \mathbb{Y}^{r-1}(\mathbf{x})\}) | A_{r-2}], \dots, \\ s_1^*(\mathbf{X}^0, \mathbf{Y}^0) = \arg \max_{\mathbf{x} \in D} \mathbb{E}[I^0(\mathbf{x}, s_2^*\{\mathbb{X}^1(\mathbf{x}), \mathbb{Y}^1(\mathbf{x})\}, \dots, s_r^*\{\mathbb{X}^{r-1}(\mathbf{x}), \mathbb{Y}^{r-1}(\mathbf{x})\}) | A_0]. \end{cases}$$

Proof. The first equality directly follows from Theorem (1). Now, the point \mathbf{x}^{n_0+r-1} is obtained after observation of $\mathbf{X}^{r-2}, \mathbf{Y}^{r-2}$ by maximizing the overall criterion

$$\begin{aligned} &\mathbb{E}[(M_0 - \min\{Y(\mathcal{X}^{n_0+1}), \dots, Y(\mathcal{X}^{n_0+r-2}), Y(\mathbf{x}), Y(\mathcal{X}^{n_0+r})\})^+ | A_{r-2}] \\ &= \mathbb{E}[(m_0 - \min\{y(\mathbf{x}^{n_0+1}), \dots, y(\mathbf{x}^{n_0+r-2}), Y(\mathbf{x}), Y[s_r^*\{\mathbf{X}^{r-1}(\mathbf{x}), \mathbb{Y}^{r-1}(\mathbf{x})\}]\})^+ | A_{r-2}], \end{aligned}$$

where equality is due to the facts that \mathcal{X}^{n_0+j} and that $Y(\mathcal{X}^{n_0+j})$ ($1 \leq j \leq r-2$) are known conditionally on A_{r-2} , and $\mathcal{X}^{n_0+r} = s_r^*(\mathbb{X}^{r-1}(\mathbf{x}), \mathbb{Y}^{r-1}(\mathbf{x}))$ by the last result. When we apply Lemma 2 with $a = m_0$, $b = m_{1:r-2}$,

$$c = \min(Y(\mathbf{x}), Y[s_r^*(\{\mathbf{X}^{r-1}(\mathbf{x}), \mathbb{Y}^{r-1}(\mathbf{x})\})])$$

leads to maximizing $\mathbb{E}[I^{r-2}(\mathbf{x}, s_r^*\{\mathbb{X}^{r-1}(\mathbf{x}), \mathbb{Y}^{r-1}(\mathbf{x})\}) | A_{r-2}]$. The remaining points are determined by Dynamic Programming (see Theorem 5.2 of Auger and Teytaud (2010) for a more general result, and Chapter 1 of Bertsekas (2007) for the basics).

Example: Decomposing the EI of a Two-Iterations Strategy. We consider a family of elementary 2-iterations strategies defined as follows ($\mathbf{a} \in D$):

$$\mathcal{S}(\mathbf{a}) = \text{“choose } \mathbf{a} \text{ at iteration 1, and maximize the 1-point EI at iteration 2.”} \quad (7)$$

Our purpose is to show that, in some cases, there exists a better strategy than sequentially maximizing the 1-point EI like Algorithm 1 does. Let us fix $\mathbf{a} \in D$ and develop $EI\{\mathcal{S}(\mathbf{a})\}$. The second point is given by

$$\mathcal{X}^{n_0+2} = s_2^*(\mathbb{X}^1, \mathbb{Y}^1) = \arg \max_{\mathbf{x} \in D} \mathbb{E}[\{\min(\mathbb{Y}^1) - Y(\mathbf{x})\}^+ | A_0, Y(\mathbf{a})]. \quad (8)$$

Lemma (2) then enables us once again to provide an interesting EI decomposition:

$$\begin{aligned}
EI^0\{\mathcal{S}(\mathbf{a})\} &= \mathbb{E} \left[\left(\min(\mathbf{Y}^0) - \min\{Y(\mathbf{a}), Y(\mathcal{X}^{n_0+2})\} \right)^+ | A_0 \right] \\
&= \mathbb{E} [I(\mathbf{a}) | A_0] + \mathbb{E} \left[\left\{ \min(\mathbb{Y}^1) - Y(\mathcal{X}^{n_0+2}) \right\}^+ | A_0 \right].
\end{aligned} \tag{9}$$

The latter hence appears as the sum of the 1-point EI at point \mathbf{a} and the expectation of the future 1-point EI at \mathcal{X}^{n_0+2} . Since $EI(\mathbf{a})$ is analytically known, calculating $EI\{\mathcal{S}(\mathbf{a})\}$ amounts to computing the second term of this sum. Now, seen from 0, $Y(\mathbf{a})$ is random. Under the usual assumptions of a centred GP with known kernel, the law of $Y(\mathbf{a})$ conditional on A_0 sends us back to the results of Simple Kriging:

$$Y(\mathbf{a}) | A_0 \sim \mathcal{N} \left(m_0(\mathbf{a}), s_0^2(\mathbf{a}) \right), \text{ where } \begin{cases} m_0(\mathbf{a}) := \mathbf{k}_0^T(\mathbf{a}) \mathbf{K}_0^{-1} \mathbf{Y}^0 \\ s_0^2(\mathbf{a}) := k(\mathbf{a}, \mathbf{a}) - \mathbf{k}_0(\mathbf{a})^T \mathbf{K}_0^{-1} \mathbf{k}_0(\mathbf{a}). \end{cases} \tag{10}$$

Algorithm 2 Computation of $EI^0\{\mathcal{S}(\mathbf{a})\}$ by Monte-Carlo

- 1: $\mathbf{X}^1 = \mathbf{X}^0 \cup \{\mathbf{a}\}$
 - 2: **for** $j \leftarrow 1, m$ **do**
 - 3: $y_{sim} \sim \mathcal{N} \{m_0(\mathbf{a}), s_0^2(\mathbf{a})\}$
 - 4: $\mathbf{Y}^1 = \mathbf{Y}^0 \cup \{y_{sim}\}$
 - 5: $\mathbf{x}_{sim}^{n_0+2} = \arg \max_{\mathbf{x} \in D} \{\mathbb{E} [I^1(\mathbf{x}) | A_1]\}$
 - 6: $v_j = \mathbb{E} [I^1(\mathbf{x}_{sim}^{n_0+2}) | A_1]$
 - 7: **end for**
 - 8: **return** $\widehat{EI^0} = \mathbb{E} [I^0(\mathbf{a}) | A_0] + \frac{1}{m} \sum_{j=1}^m v_j$
-

Numerical Application. Back to *Example 1*, $EI^0\{\mathcal{S}(\mathbf{a})\}$ is computed for the boundary point $\mathbf{a} = 1$ and compared to the EI value obtained with two iterations of Algorithm 1, i.e. twice maximizing the regular EI . As detailed in Algorithm 2, the computation of $\mathbb{E} \left[\left\{ \min(\mathbb{Y}^1) - Y(\mathcal{X}^{n_0+2}) \right\}^+ | A_0 \right]$ is based on:

$$EI^1\{\mathcal{S}(\mathbf{a})\} \approx \frac{1}{m} \sum_{i=1}^m \mathbb{E} \left[\left(\min\{\mathbb{Y}^1(\mathbf{a})\} - Y\{\mathcal{X}^{n_0+2}(\mathbf{a})\} \right)^+ | A_0, Y(\mathbf{a}) = y_a^i \right], \tag{11}$$

where the y_a^i ($1 \leq i \leq m$) are i.i.d. following $\mathcal{N}\{m_0(\mathbf{a}), s_0^2(\mathbf{a})\}$. Figure 2 sums up the results obtained by running Algorithm 2 with $m = 1000$, for both $\mathbf{a} = 1$ and \mathbf{a} fixed to the 1-point EI maximizer. We compared the $\widehat{EI^0}$ obtained for the two strategies by means of a Welch t-test —using the *t.test* function of the *stats* R package. With respective estimates of 0.6061 and 0.6132 for the means, the second corresponding to $\mathbf{a} = 1$, the t-test with alternative hypothesis “the true difference in means is less than 0” returned a p-value of 0.056 ($t = -1.589$ and $df = 1847$).

The slightly higher EI hence obtained with $\mathbf{a} = 1$ supports the belief that maximizing 1-point EI at each iteration is not (always) the best sequential strategy with fixed horizon. In this particular example, the phenomenon seems due to the delayed payoff associated with sampling at $\mathbf{a} = 1$. Indeed, evaluating y there at the first iteration leaves room to explore the most interesting zone with a little bit more information at

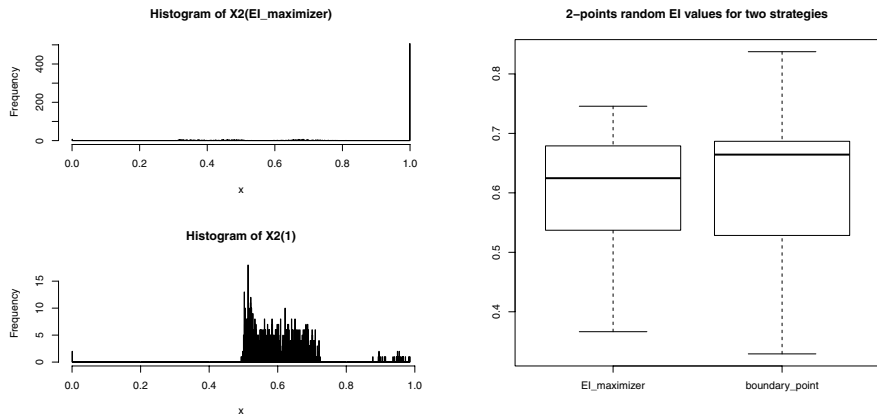


Fig. 2: The left graphics represents the two populations of \mathcal{X}^{n_0+2} points (1000 each) corresponding to both strategies, and the right one compares the samples of improvement values, $i^0(\mathcal{S})$ from eq. (3), obtained in both cases.

iteration 2 than initially. In the straightforward strategy however, one greedily visits the main bump of the 1-point EI at the first iteration and then almost systematically samples y at the boundary point $x = 1$ during the second shot (See Fig. (2), upper left).

Computational Cost of Finite Horizon Strategies. Evaluating the EI of a strategy is an issue: for $r = 2$, m^2 simulations are needed (m by using the 1-point EI formula at step 2), m 1-point EI maximizations, not to mention the m Kriging model updates and auxiliary computation costs. A direct extension of this computational scheme to horizons $r \geq 3$ results in costs of the order of m^r operations, hence exponentially increasing with the horizon. This is one instance of Bellman's famous *curse of dimensionality*. Now, deriving the optimal strategy is even worse: not only is the computational time exponentially increasing in r , but it is also exponential in the time needed to optimize one-point criteria. Considering for simplicity an exhaustive maximization of the EI 's and $EI\{\mathcal{S}(\mathbf{a})\}$'s over a p -point grid, finding the optimal strategy costs $p \times m \times (p - 1) \times m$ improvement computations, which becomes $A_p^r m^r$ in the general case ($1 \leq r \leq p - 1$). Finding the optimal strategy by such a method seems thus limited to an horizon of 2 or 3, and only makes sense in the case of costly objective functions for which the optimal strategy is likely to bring a higher improvement than the same computation time invested in additional evaluations of y . However, 2- or 3-step optimal strategies might remain attractive in the following contexts:

- Sequential evaluations of y with consecutive blocks of 2-step optimal strategies,
- Sequential-parallel evaluations of y with consecutive blocks of $q \in 2\mathbb{N}$ points, with blocks of two q -point designs optimal in the sense of a 2-step strategy.

Conclusion and perspectives. The presented results extend the *multipoint EI* to optimization strategies with finite time horizons. Thanks to the modeling of the future points and observations in terms of random variables, the latter criterion is proposed and analyzed to derive the sequence of decisions to be made by the optimal fixed horizon algorithm, obtained by dynamic programming. It has been illustrated, on the basis of a specific example, that the classical *EI* algorithm is suboptimal. To this end, the strategic value of a point is decomposed as the sum of its one-point *EI* plus a delayed payoff, estimated by Monte-Carlo based on GP conditional simulations.

Perspectives include a detailed study and improvements of the latter Monte-Carlo method. Dimension reduction techniques and well-tuned heuristics may be required to allow the computation of reasonable estimates for the *EI* of a strategy with horizon $r \geq 3$. Furthermore, both large-scale practical examples and deeper connections with existing work on sequential strategies in the fields of control theory and approximate dynamic programming (Powell 2007), are currently being considered. In particular, the close (but not similarly proven nor illustrated) results given in Mockus (1988), recently discovered by the authors, motivate revisiting this book two decades later with a contemporary scientific approach and increased computational capacity.

Acknowledgements This work was funded by the *Optimisation Multi-Disciplinaire* (OMD) project of the French Research Agency (ANR). The authors would like to thank Julien Bect (Ecole Supérieure d'Electricité) for providing them with the related results of Mockus (1988).

References

- Auger, A. and O. Teytaud (2010). Continuous lunches are free plus the design of optimal optimization algorithms. *Algorithmica* 57, 121–146.
- Bertsekas, D. (2007). *Dynamic Programming and Optimal Control, Vol. 1*. Athena Scientific, Belmont, MA.
- Ginsbourger, D., R. Le Riche, and L. Carraro (2010). *Computational Intelligence in Expensive Optimization Problems*, Chapter “Kriging is well-suited to parallelize optimization”. *Studies in Evolutionary Learning and Optimization*. Springer-Verlag.
- Jones, D., M. Schonlau, and W. Welch (1998). Efficient global optimization of expensive black-box functions. *Journal of Global Optimization* 13, 455–492.
- Mockus, J. (1988). *Bayesian Approach to Global Optimization*. Amsterdam: Kluwer.
- Powell, W. (2007). *Approximate Dynamic Programming. Solving the Curses of Dimensionality*. New York: Wiley.
- Rasmussen, C. and K. Williams (2006). *Gaussian Processes for Machine Learning*. Boston, MA: M.I.T. Press.
- Schonlau, M. (1997). *Computer Experiments and Global Optimization*. Ph.D. thesis, University of Waterloo, Canada.
- Stein, M. (1999). *Interpolation of Spatial Data, Some Theory for Kriging*. Springer.

Optimal Designs for Linear Logistic Test Models

Ulrike Graßhoff, Heinz Holling and Rainer Schwabe

Abstract An important class of models within item response theory are Linear Logistic Test Models (LLTM). These models provide a means for rule-based item generation in educational and psychological testing based upon cognitive theories. After a short introduction into the LLTM, optimal designs for the LLTM will be developed with respect to the item calibration step assuming that persons' abilities are known. Therefore, the LLTM is embedded in a particular generalized linear model. Finally, future developments are outlined.

1 Introduction

The following work is motivated by the construction of rule-based tests for measuring intelligence. Intelligence is a very important prerequisite for academic and vocational performance as it provides a basis for nearly all cognitive abilities and skills. Thus, measurement of intelligence has played an important role in psychological test theory for a long time.

Recent approaches to developing intelligence tests are driven by attempts to automate item generation based on rule based items. Essential to rule-based item design are cognitive models for “solving items” and linking the components affecting item complexity and difficulty to the cognitive processes in operation during problem

Ulrike Graßhoff

Institut für Mathematische Stochastik, Otto-von-Guericke-Universität Magdeburg, PF 4120, 39016 Magdeburg, Germany, e-mail: ulrike.grasshoff@mathematik.uni-magdeburg.de

Heinz Holling

Psychologisches Institut IV, Westfälische Wilhelms-Universität Münster, Fliegerstr. 21, 48149 Münster, Germany, e-mail: holling@uni-muenster.de

Rainer Schwabe

Institut für Mathematische Stochastik, Otto-von-Guericke-Universität Magdeburg, PF 4120, 39016 Magdeburg, Germany, e-mail: rainer.schwabe@mathematik.uni-magdeburg.de

solving. The main principle is to analyze components of items that influence item complexity and difficulty and then use these components to combine them and generate items of arbitrary complexity and difficulty level.

Typical examples for such rule-based items are figural analogy items. Figure 1 illustrates a sample for a figural analogy item with the basic test format A:B = C:? (Kirchhoff and Holling 2010). Cognitive operations, defined by the relation between the A and B term were specified based on theories of figural intelligence. These operations constitute construction rules according to which items were generated. Nine different rules transforming the elements from A to B and thus from C to D were applied. These rules refer to four general rule classes: Reflection, rotation, size, and sequence. The item exposed in the left of figure 1 contains two rules, “sequence plus four” and “rotation 180°”. When “sequence plus” is applied letters were consecutive according to the alphabet and digits were increased according to arithmetic addition. (The correct solution would be “f” for the problem in figure 1.)




Fig. 1: Example for an item of a rule based figural analogy test.

A statistical foundation for constructing such items is provided by LLTMs. LLTMs constitute an important class of models in item response theory (IRT) which is usually the foundation of modern educational and psychological testing. In IRT models the probability of solving an item depends on item and person parameters. The basic and very popular IRT model is the Rasch model which is based upon a binary response variable Y_{ij} where $Y_{ij} = 1$ if person i solves an item j correctly and $Y_{ij} = 0$ otherwise. The success probability $P(Y_{ij} = 1)$ is given in terms of two parameters θ_i and σ_j , which describe the person’s “ability” and the “difficulty” of the given item, respectively. These quantities are connected by the logistic link function as

$$P(Y_{ij} = 1) = \frac{\exp(\theta_i - \sigma_j)}{1 + \exp(\theta_i - \sigma_j)} = \frac{1}{1 + \exp\{-(\theta_i - \sigma_j)\}} \tag{1}$$

resulting from the log odds formulation

$$\log \frac{P(Y_{ij} = 1)}{P(Y_{ij} = 0)} = \theta_i - \sigma_j .$$

Thus the logarithmic odds ratio is modeled as the difference between the person ability and item difficulty parameters.

The number of correctly solved items is a sufficient statistic for the ability parameter. This is an important advantage of the Rasch model compared to other IRT mod-

els. Closely connected with this feature is what Rasch (1966, p. 104 - 105) called “specific objectivity”: “The comparison of any two subjects can be carried out in such a way that no parameters are involved other than those of the two subjects ... Similarly, any two stimuli can be compared independently of all other parameters than those of the two stimuli as well as the parameters of the subjects having been replaced with observable numbers. It is suggested that comparisons carried out under such circumstances be designated as specific objective”. Thus, the Rasch model stands out in IRT due the above mentioned features.

An important extension of the Rasch model especially for generating rule based tests are LLTM. In the LLTM the difficulty of an item j is composed of K item specific components or rules X_{jk} given as a weighted sum of basic parameters η_k , $k = 1, \dots, K$,

$$\sigma_j = \sum_{k=1}^K \eta_k X_{jk} + c$$

with a norming constant term c , so that

$$\sum_{j=1}^J \sum_{k=1}^K \eta_k X_{jk} + c = 0,$$

where the summation is taken over all possible items $j = 1, \dots, J$.

Inserting this representation in (1) leads to

$$P(Y_{ij} = 1) = \frac{\exp(\theta_i - \sum_{k=1}^K \eta_k X_{jk} - c)}{1 + \exp(\theta_i - \sum_{k=1}^K \eta_k X_{jk} - c)}. \quad (2)$$

Usually the model parameters in (1) and (2), i.e. σ_j and η_k respectively, are estimated in a two-step procedure. First, the item parameters are estimated by maximum likelihood conditioned on fixed person parameters or by using a Bayesian approach. This step is called item calibration. In the second step the person parameters θ_i can be estimated under the relevant response model given the estimated item parameters.

An important tool for estimating the parameters of the item characteristics is provided by optimal designs. Optimal design has a long history in psychological and educational testing starting with Birnbaum (1968). He was the first who realized the importance of optimal test design for applications of IRT. When the interaction between an examinee and a test item is described by a response model with parameters for the examinee and the item, the information in the response about the examinee’s parameter is dependent on the item. Thus, the items or, more precisely, their specific components may serve as design variables. An immediate question is which distribution of item specific components would be optimal to measure an examinee or a population of examinees. Birnbaum suggested optimizing Fisher’s information about θ in the responses. His approach to optimal test designing consisted of three steps: First, establishing the measurement goal of the test that is to be assembled (e.g., diagnosis; pass-fail decision making; evaluation of educational progress). Second, translating the goal into a target for the information function of

the test. Third, selecting a test from the pool of calibrated items with an information function which best approximates the target.

Meanwhile, innumerable applications of optimal design within IRT have been developed (e.g. van der Linden 2005). These approaches may be categorised into two types of design problems. The first type is known as test design and refers to the optimal selection of items for estimating ability parameters. The second type, sometimes called sampling design, deals with the sampling of test-takers for optimal estimation of item parameters. Both types of design problem may be further differentiated according to the kind of testing: testing using fixed-form tests vs. adaptive testing. Buyske (2005) gives a short overview about important results of applying optimal design to these four different situations.

A still open problem is the test design of an LLTM, i. e. how to select the items for estimating the parameters of the item characteristics. Every item in an LLTM is described by its properties and represented by a row in the design matrix. Usually, this matrix consists of binary elements (0: property is not given, 1: property is given). But, in some cases properties may be applied several times and the elements in the design matrix describe how often a property is given in this task. Furthermore, the collection of items has to be subjected to certain constraints to deal with the content specifications of the test, its format, and practical restrictions, such as the testing time available, uniform usage of items over time, and items which cannot be used in the same test because one item contains a clue to the solution of the other. In such situations a restricted design space is imposed.

2 Optimal Design

In the following optimal designs for the LLTM will be constructed with respect to the item calibration step, where persons' abilities are known and can be adjusted arbitrarily. De Boeck and Wilson (2004) have shown that most IRT models like the LLTM are special cases of generalized linear models. Following this approach we embed model (2) in such a particular generalized linear model.

$$E(Y(\mathbf{x}_1, x_2)) = \mu(\mathbf{f}(\mathbf{x}_1)^\top \boldsymbol{\beta} + x_2) \quad (3)$$

with $\mathbf{x}_1 \in \mathcal{X}$, $x_2 \in \mathbb{R}$ and $\mu(z) = \exp(z)/(1 + \exp(z))$.

In order to establish the embedding of the LLTM we have to set $\mathbf{f}(\mathbf{x}_{1j})^\top \boldsymbol{\beta} = -\sum_{k=1}^K \eta_k X_{jk} - c$ and $x_{2i} = \theta_i$. This means that in terms of \mathbf{x}_1 we can think of results from K attributes with the vector of regression functions $\mathbf{f} = -(1, \mathbf{f}_1^\top, \dots, \mathbf{f}_K^\top)^\top$ for $\mathbf{x}_1 = (x_{11}, \dots, x_{1K})^\top$ and x_{1k} is the value of the k th attribute. The regression functions \mathbf{f}_k will represent dummy variables for qualitative factors or real valued functions when quantitative factors are involved. Concerning the second part of the linear component x_2 is treated as a continuous factor and we can assume the corresponding effect as a known parameter with value 1 due to the assumptions of the Rasch model

and hence x_2 is a free additional variable to choose. We assume in what follows that each item is presented to exactly one person, this means $i = j$.

The vector of unknown parameters is denoted by $\beta = (\beta_0, \beta_1^\top, \dots, \beta_k^\top)^\top \in \mathbb{R}^p$, where $\beta_0 = c$ is the normalizing constant and β_k is related to the effect of the k th attribute. The mean response μ as well as the response variance $\sigma^2 = \mu(1 - \mu)$ depend on the linear effect $\mathbf{f}(\mathbf{x}_1)^\top \beta + x_2$ only; so let $\text{Var}(Y(\mathbf{x}_1, x_2)) = \sigma^2(\mathbf{f}(\mathbf{x}_1)^\top \beta + x_2)$. Denote by $\lambda(z) = \mu'(z)^2 / \sigma^2(z)$ the intensity function, which is evaluated at $z = \mathbf{f}(\mathbf{x}_1)^\top \beta + x_2$ gives the intensity of the information of an observation with settings (\mathbf{x}_1, x_2) . For the present logistic model this intensity occasionally coincides with the variance ($\lambda = \mu' = \sigma^2$). The Fisher information for any point can be approximated by local linearisation for fixed β

$$\begin{aligned} \mathbf{M}(\mathbf{x}_1, x_2; \beta) &= \sigma^2(\mathbf{f}(\mathbf{x}_1)^\top \beta + x_2)^{-1} \frac{\partial}{\partial \beta} \mu(\mathbf{f}(\mathbf{x}_1)^\top \beta + x_2) \frac{\partial}{\partial \beta} \mu(\mathbf{f}(\mathbf{x}_1)^\top \beta + x_2)^\top \\ &= \lambda(\mathbf{f}(\mathbf{x}_1)^\top \beta + x_2) \mathbf{f}(\mathbf{x}_1) \mathbf{f}(\mathbf{x}_1)^\top \end{aligned}$$

For measuring the quality of an exact design ξ on $\mathcal{X} \times \mathbb{R}$ represented by N design points $\xi = ((\mathbf{x}_{11}^\top, x_{21}), \dots, (\mathbf{x}_{1N}^\top, x_{2N}))$, we use the normalized information matrix namely

$$\mathbf{M}(\xi; \beta) = \frac{1}{N} \sum_{i=1}^N \mathbf{M}(\mathbf{x}_{1i}, x_{2i}; \beta) = \frac{1}{N} \sum_{i=1}^N \lambda(\mathbf{f}(\mathbf{x}_{1i})^\top \beta + x_{2i}) \mathbf{f}(\mathbf{x}_{1i}) \mathbf{f}(\mathbf{x}_{1i})^\top.$$

restrict the D -optimality criterion to local optimality at a given parameter vector β at a first stage. A design ξ^* is called D -optimal (locally) at β , if it maximizes the determinant of the information matrix $\mathbf{M}(\xi; \beta)$.

We denote for a design ξ by ξ_1 the marginal design with respect to the first component \mathbf{x}_1 and for the corresponding linear response $E(Y_1(\mathbf{x}_1)) = \mathbf{f}(\mathbf{x}_1)^\top \beta$ the ‘‘linear’’ information matrix by $\mathbf{M}_1(\xi_1) = 1/N \sum_{i=1}^N \mathbf{f}(\mathbf{x}_{1i}) \mathbf{f}(\mathbf{x}_{1i})^\top$, which does not depend on β .

Theorem 1. *Let ξ_1^* be an exact D -optimal design on \mathcal{X} for the marginal linear model $E(Y_1(\mathbf{x}_1)) = \mathbf{f}(\mathbf{x}_1)^\top \mathbf{1}$ given by N design points $(\mathbf{x}_{11}^*, \dots, \mathbf{x}_{1N}^*)$. For given $\mathbf{1}$ set $x_{2i}^* = -\mathbf{f}(\mathbf{x}_{1i}^*)^\top \mathbf{1}$. Then the combined design $\xi^* = ((\mathbf{x}_{11}^*, x_{21}^*), \dots, (\mathbf{x}_{1N}^*, x_{2N}^*))$ is D -optimal at $\mathbf{1}$ for the model (3).*

Proof of Theorem 1

Since $\lambda(z) = \exp(z)/(1 + \exp(z))^2 \leq \lambda(0)$ is bounded by $\lambda(0) = 1/4$, we have for the information at any setting (\mathbf{x}_1, x_2) that $\mathbf{M}(\mathbf{x}_1, x_2) \leq \frac{1}{4} \mathbf{M}_1(\mathbf{x}_1)$. Hence for every exact design $\xi = ((\mathbf{x}_{11}^\top, x_{21}), \dots, (\mathbf{x}_{1N}^\top, x_{2N}))$ the information is bounded $\frac{1}{4}$ times the marginal information of the first component,

$$\mathbf{M}(\xi; \beta) \leq \frac{1}{4} \mathbf{M}_1(\xi_1),$$

and equality is attained for ξ^* as $\mathbf{f}(\mathbf{x}_{1i}^*)^\top \boldsymbol{\beta} + x_{2i}^* = 0$. The monotonicity of the determinant yields

$$\det \mathbf{M}(\xi; \boldsymbol{\beta}) \leq \left(\frac{1}{4}\right)^p \det \mathbf{M}_1(\xi_1) \leq \left(\frac{1}{4}\right)^p \det \mathbf{M}_1(\xi_1^*) = \det \mathbf{M}(\xi^*; \boldsymbol{\beta}),$$

where the second inequality holds due to the D -optimality of ξ_1^* . This proves the D -optimality of ξ^* . \square

This theorem justifies the usual practice of using a D -optimal design to determine both the design matrix of the item components and furthermore a person parameter which corresponds to the difficulty of the given item. For the Rasch model it is well-known and easy to show that Fisher's information is optimized when $\theta_i = \sigma_j$. Since $\sigma_j = \sum_{k=1}^K \eta_k X_{jk} + c$ the result $\theta_i = \sum_{k=1}^K \eta_k X_{jk} + c$ is plausible.

Even though this theorem is formulated in terms of exact designs, it is valid for approximate designs as well. Note that the present theorem is in the same spirit as the standardization used in Graßhoff, Großmann, Holling, and Schwabe (2007) for a slightly different model.

3 Discussion

The derived optimal design above is a locally D -optimal design. For practical purposes this approach should be extended by developing sequential designs. Another possibility is to use a weight distribution for the unknown parameters in a semi-Bayesian spirit.

Furthermore, optimal designs should be derived for recent extensions of the LLTM. In the traditional LLTM item effects are considered to be fixed effects. This may be a rather strict assumption for certain test situations. A LLTM with relaxed assumptions often called LLTME is obtained, when item-related random effects are assumed (Janssen, Schepers, and Peres 2004). Thus, additional item variation is taken into account. A further extension and relaxation of the LLTM is the RWLLTM - random weights LLTM (Rijmen and De Boeck 2002) - in which the weights of the item characteristics are random. For a more detailed exposition of these models see De Boeck and Wilson (2004).

Applications of these LLTMs offer interesting opportunities compared to traditional, heuristic approaches to item construction. First, this model provides a means to analyze and evaluate cognitive theories in the realm of task design (e.g. Embretson 1998). The LLTM therefore constitutes a means which requires pre-experimental hypotheses on item structure characteristics which are then to be tested. Results of the LLTM may be used for testing the hypothesized cognitive model.

Second, it enables automatic item generation as shown by Freund, Hofer, and Holling (2008). Based on rationally constructed figural matrix items, these authors

implemented algorithms and generated figural matrix items by computer. The main advantages of such automatic item generation are increased economy, avoidance of construction errors, and higher comparability of items due to more stringent construction algorithms.

Third, rule based generated items are useful for adaptive testing. In adaptive testing, the test is not assembled prior to the testing session but in real time by the computer which runs the test session. Numerous studies have shown that the reduction in test length for an adaptive test relative to a traditional fixed test form can be expected to be some 50-60%. But, adaptive testing, which is usually based on item pools, suffered from severe security problems. In several cases, examinees remembered items and published them on the internet. However, using a rule based approach a nearly infinite number of items may be generated on demand and administered to the test takers.

Acknowledgements This research was supported by the Deutsche Forschungsgemeinschaft (DFG) under grant HO 1286/6-1.

References

- Birnbaum, A. (1968). Some latent trait models and their use in inferring an examinee's ability. In F. Lord and M. Novick (Eds.), *Statistical Theories of Mental Test Scores*, pp. 397–479. Reading: MIT Press.
- Buyske, S. (2005). Optimal design in educational testing. In W. Wong and M. Berger (Eds.), *Optimal Designs: Their Roles and Applications*, pp. 1–19. Chichester: Wiley.
- De Boeck, P. and M. Wilson (2004). *Explanatory Item Response Models: A Generalized Linear and Nonlinear Approach*. New York: Springer.
- Embretson, S. E. (1998). A cognitive design system approach to generating valid tests: Application to abstract reasoning. *Psychological Methods* 3, 300–396.
- Freund, P., S. Hofer, and H. Holling (2008). Explaining and controlling for the psychometric properties of computer-generated figural matrix items. *Applied Psychological Measurement* 32, 195–210.
- Graßhoff, U., H. Großmann, H. Holling, and R. Schwabe (2007). Design optimality in multi-factor generalized linear models in the presence of an unrestricted quantitative factor. *Journal of Statistical Planning and Inference* 137, 3882–3893.
- Janssen, R., J. Schepers, and D. Peres (2004). Models with item and item group predictors. In P. De Boeck and M. Wilson (Eds.), *Explanatory Item Response Models: A Generalized Linear and Nonlinear Approach*, pp. 189–212. New York: Springer.
- Kirchhoff, B. and H. Holling (2010). Figural analogies: Rule-based test construction and evaluation using random effect models. Manuscript under review.

- Rasch, G. (1966). An individualistic approach to item analysis. In P. Lazarsfeld and N. Henry (Eds.), *Readings in Mathematical Social Sciences*, pp. 89–107. Cambridge: MIT Press.
- Rijmen, F. and P. De Boeck (2002). The random weights linear logistic test model. *Applied Psychological Measurement* 26, 271–285.
- van der Linden, W. (2005). *Linear Models for Optimal Test Design*. New York: Springer.

A Class of Screening Designs Robust to Active Second-Order Effects

Bradley Jones and Christopher J. Nachtsheim

Abstract Screening designs are attractive for assessing the relative impact of a large number of factors on a response of interest. Engineers prefer factors with three levels over two-level factors because having three levels allows for some assessment of curvature in the factor-response relationship. Yet, the most familiar screening designs limit each factor to only two levels. We propose a new class of designs that have three levels, allow for the estimation of quadratic effects, and have the property that the linear effect of every factor is independent of all second-order effects. We also provide an algorithm for design construction.

1 Introduction

An undesirable property of resolution III fractional-factorial screening designs Box and Hunter (1961) is that they confound the main effects of the factors with one or more two-factor interactions. If one or more of these confounded effects is active, the experimenter is left with substantial ambiguity. Resolving this ambiguity generally requires the experimenter to perform additional processing runs.

If there is strong reason to suspect active two-factor interactions, a resolution IV fractional-factorial design is a desirable alternative. However, these designs require twice as many runs as the resolution III design and they have no capability for capturing curvature due to pure quadratic effects. Of course, it is traditional to add centre runs to two-level screening designs to get a global assessment of curvature. Still, these runs do not allow for separate estimation of the quadratic effects of each

Bradley Jones

SAS Institute, Cary, NC 27513, USA, e-mail: bradley.jones@sas.com

Christopher J. Nachtsheim

Carlson School of Management, University of Minnesota, Minneapolis, MN 55455, USA, e-mail: nacht001@umn.edu

factor. So, an indication of curvature in the analysis leads to still more ambiguity that can only be resolved with additional runs.

We introduce a class of screening designs for quantitative factors that have the following desirable properties:

1. Two-factor interactions and quadratic effects are uncorrelated with main effects.
2. All quadratic effects are estimable.
3. The number of required runs is only one more than twice the number of factors.
4. With four or more factors, the design projects to the 3×3 factorial design in each pair of factors.

The methodology proposed here is different from, but related to, several prior contributions. For example, Cheng and Wu (2001) develop a novel approach for factor screening and response surface estimation using fractions of 3^m experiments for $n = 27$, and by using fractions of mixed-level orthogonal arrays for $n = 18$ and $n = 36$. Designs produced by Cheng and Wu are related, in that they employ three levels and can provide estimates of first- and second-order effects. The designs proposed here differ in that they (1) generally allow for substantially fewer runs for the same number of factors, and (2) do not require orthogonality between main effects.

The approach of Tsai, et al. (2000) (TGM) is also related. They consider the design and analysis of three-level designs using a design strategy that considers the efficiencies of low-level projections. We note that this seemingly disparate approach did lead to nearly the same arrangement as ours in one instance. Design 1 of TGM's Table 5 essentially identical our design for six factors, discussed below, with the exception that TGM require two centre points to our one.

In Jones and Nachtsheim (2009), the current authors considered the construction of designs that minimize the squared norm of the alias matrix subject to constraints on the D-efficiency of the design. We found that designs similar to those discussed here were sometimes produced using the proposed constrained optimal design approach.

We describe the structure of our designs using a simple illustrative example in Section 2. In Section 3 we present an algorithm for generating these designs. Section 4 deals with design diagnostics and comparisons. Section 5 supplies some ideas for modeling data obtained, and we conclude with a short summary in Section 6.

2 Design Structure: An Example

For two-level designs, one way to make two-factor interactions independent of main effects involves mirroring each row in the design by another that reverses the signs of all the elements in that row. This technique is called folding over the design.

Table 1 shows an example design with six factors and 13 runs. Note that the 2nd row is obtained by multiplying each element of the first row by -1 . Similarly, the 4th row mirrors the 3rd row. This pattern repeats for each pair of rows through row 12. The last row is a centre run.

Another pattern in Table 1 is apparent by observing the location of the zero elements. The first pair of runs has zero elements in the first column and the second pair of runs has zero elements in the second column. This pattern repeats so that each column has a contiguous pair of zero elements in the first 12 rows. Adding the centre run in the last row results in a design can fit a model including an intercept term, all the main effects and all the pure quadratic effects of each factor.

Table 1: Robust Screening Design for Six Factors (*A* through *F*), with a Simulated Response, *y*

Run	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	<i>E</i>	<i>F</i>	<i>y</i>
1	0	1	-1	-1	-1	-1	21.04
2	0	-1	1	1	1	1	10.48
3	1	0	-1	1	1	-1	17.89
4	-1	0	1	-1	-1	1	10.07
5	-1	-1	0	1	-1	-1	7.74
6	1	1	0	-1	1	1	21.01
7	-1	1	1	0	1	-1	16.53
8	1	-1	-1	0	-1	1	20.38
9	1	-1	1	-1	0	-1	8.62
10	-1	1	-1	1	0	1	7.80
11	1	1	1	1	-1	0	23.56
12	-1	-1	-1	-1	1	0	15.24
13	0	0	0	0	0	0	19.91

The columns of this design are orthogonal to each other. Because of the mirroring in pairs of runs, the main effects are all independent of any active two-factor interaction. However, the two-factor interactions are correlated with the pure quadratic effects. Section 5 considers the resulting analytical complexity when both pure quadratic effects and two-factor interactions are active.

Compared to the 12 run Plackett-Burman design with one additional centre run, the above design has a D-efficiency of 85.5% for the main effects model. Both designs are orthogonal for the main effects. The relative variance of each main effect in our design is 1/10 compared to 1/12 for the Plackett-Burman design. The ability to estimate pure quadratic effects and the independence of the main effects and the two-factor interactions compensates for the loss of efficiency in fitting the main effects model. Note that each main effect in the Plackett-Burman design is partially aliased with several two-factor interactions.

3 Algorithm

The patterns illustrated in the previous section are common to each member of the class of designs. Each even numbered row is a mirror image of the previous row.

The last row contains all zero elements. Finally, for the k th column, rows $2k$ and $2k - 1$ contain zeros.

The design algorithm maximizes the determinant of the information matrix of the main effects model while enforcing this structure. The starting design includes zeros in all the required places. These elements are not allowed to change in the course of the algorithm. The other elements in the odd numbered rows of the starting design are chosen randomly on the interval $[-1, 1]$. The even numbered rows of the starting design are obtained from the odd numbered rows by multiplying each element by -1 .

The starting design is improved using a variant of the coordinate exchange algorithm of Meyer and Nachtsheim (1995). For each nonzero element of every row, the algorithm evaluates the effect of changing that element to 1 or -1 while simultaneously changing the element in the mirroring row to -1 or 1 respectively. If the determinant of the information matrix improves for either or both of these operations, then the current design is updated for the given row and the mirroring row for the better of the two possible exchanges. After the first pass through each element of the design, the algorithm makes a second pass through every nonzero element. If any element of the design changes in the second pass, then the algorithm performs another pass. This process continues until there are no changes in any pass through the design or when a maximum iteration limit is reached. The resulting design, having been obtained from one random starting design, may not be globally optimal, so multiple random starting designs are used in an effort to avoid local maxima. A JMP scripting language (JSL) code for creating any design in this class of designs is available from the authors.

Note that to create a randomized design, the rows of the design generated by the algorithm should be randomly shuffled.

4 Design Diagnostic Comparisons

The six factor design in Table 1 has orthogonal columns. This is not the case in general. For example, the best seven factor design we found, shown in Table 2, had column correlations of $\pm 1/6$. This correlation has the effect of increasing the relative variance of the main effects from a theoretical minimum of 0.083 to 0.098.

The best eight factor design, shown in Table 3, also has orthogonal columns and is a strong alternative to the standard 2^{8-4} regular fractional-factorial design with one center run.

By construction, both designs have main effects independent of two-factor interactions. The relative coefficient variance for each main effect in the fractional-factorial design is $1/16$ compared to $1/14$ for the robust screening design. However the robust screening design can estimate the pure quadratic effect of every factor.

Table 4 compares the D-efficiency of the robust screening design to the exact D-optimal design for the main effects model. The D-optimal design was created using the coordinate exchange algorithm in JMP. We added a center run to the D-optimal

Table 2: Robust Screening Design for Seven Factors

Run	X_1	X_2	X_3	X_4	X_5	X_6	X_7
1	0	1	-1	1	-1	1	-1
2	0	-1	1	-1	1	-1	1
3	-1	0	1	-1	1	1	-1
4	1	0	-1	1	-1	-1	1
5	1	-1	0	1	1	1	1
6	-1	1	0	-1	-1	-1	-1
7	1	-1	-1	0	1	-1	-1
8	-1	1	1	0	-1	1	1
9	-1	-1	1	1	0	-1	-1
10	1	1	-1	-1	0	1	1
11	-1	1	-1	1	1	0	1
12	1	-1	1	-1	-1	0	-1
13	1	1	1	1	1	-1	0
14	-1	-1	-1	-1	-1	1	0
15	0	0	0	0	0	0	0

Table 3: Robust Screening Design for Eight Factors

Run	X_1	X_2	X_3	X_4	X_5	X_6	X_7	X_8
1	0	-1	1	1	-1	1	1	1
2	0	1	-1	-1	1	-1	-1	-1
3	-1	0	-1	1	1	1	1	-1
4	1	0	1	-1	-1	-1	-1	1
5	-1	-1	0	1	1	-1	-1	1
6	1	1	0	-1	-1	1	1	-1
7	1	-1	1	0	1	1	-1	-1
8	-1	1	-1	0	-1	-1	1	1
9	-1	-1	1	-1	0	-1	1	-1
10	1	1	-1	1	0	1	-1	1
11	1	-1	-1	-1	1	0	1	1
12	-1	1	1	1	-1	0	-1	-1
13	-1	1	1	-1	1	1	0	1
14	1	-1	-1	1	-1	-1	0	-1
15	1	1	1	1	1	-1	1	0
16	-1	-1	-1	-1	-1	1	-1	0
17	0	0	0	0	0	0	0	0

design with $2m$ runs (where m is the number of factors) to provide a fair comparison. That is, in each case both designs have the same number of runs including one center run.

Table 4: Relative Efficiency of the Robust Screening Design to the D-optimal design for the main effects model

Number of Factors:	6	7	8	9	10	11	12
D-Efficiency (%)	85.5	83.3	88.8	86.8	87.3	89.1	89.8

Compared to the D-optimal design for six to twelve factors, the robust screening design gives up some efficiency in the estimation of the main effects. In compensation the robust screening design allows for the estimation of pure quadratic effects of each factor and the main effects are unbiased by any two-factor interaction. In general, a D-optimal design for the main effects model has main effects that can experience substantial bias from active two-factor interactions.

We note that for the proposed class of designs, the correlation between quadratic effects of a factor q and a factor r for a design involving m factors, denoted $r_{qq,rr}(m)$ is:

$$r_{qq,rr}(m) = \frac{1}{3} - \frac{1}{m-1} \quad m > 1 \tag{1}$$

This correlation is increasing in m and approaches $+1/3$ as $m \rightarrow \infty$. Values for two through 20 factors are shown in Table 5. We note that for $m = 4$ factors, the

correlation is zero. It turns out that our algorithm produces a graeco-latin square in this case.

Table 5: Correlations between quadratic effects for two through 20 factors

Number of Factors (m)	Correlation Between Quadratic Effects ($r_{qq,rr}(m)$)
2	-0.6667
3	-0.1667
4	0.0000
5	0.0833
6	0.1333
7	0.1667
8	0.1905
9	0.2083
10	0.2222
11	0.2333
12	0.2424
13	0.2500
14	0.2564
15	0.2619
16	0.2667
17	0.2708
18	0.2745
19	0.2778
20	0.2807

Table 6: Correlations between quadratic effects and two-factor interactions for four through 20 factors

Number of Factors (m)	Correlation Between Quadratic Effects ($r_{qq,rz}(m)$)
4	0.7071
6	0.4655
8	0.3673
10	0.3118
12	0.2752
14	0.2489
16	0.2289
18	0.2129
20	0.1999
22	0.1890
24	0.1797
26	0.1716
28	0.1645
30	0.1582
32	0.1526
34	0.1476
36	0.1430
38	0.1388
40	0.1350

Characterizing correlations between quadratic effects and interactions is more complex. If the number of factors is even and the two-factor interaction columns are balanced, the correlations between quadratic effects and two-factor interactions have a simple closed form. Although we did observe balanced interaction columns in the examples presented, we cannot guarantee that this condition will always be met in globally optimal designs for the class. However, when the interaction columns are balanced and the number of factors is even, correlation assumes one of three values, depending on whether or not the two terms have a factor in common. For a common factor, $r_{qq,qr} = 0$, otherwise:

$$r_{qq,rz} = \frac{\pm 2}{\sqrt{\frac{12m^2 - 36m + 24}{2m + 1}}} \quad m > 3, \text{ and } m \text{ even} \tag{2}$$

The absolute value of this correlation decreases in m , approaching zero as $m \rightarrow \infty$. Absolute values of the correlation for even numbers of factors ranging from $m = 2$ through $m = 40$ are shown in Table 6. We have not been able to develop general closed-form expressions for the correlation if either m is odd, or if the optimal design does not produce balanced interaction columns.

Note that quadratic effects of any factor are uncorrelated with any two-factor interaction involving that factor. Assuming that models exhibit effect heredity, this is another beneficial property of this design class.

5 Suggestions for Analysis

The analysis of these designs is straightforward if only main effects or main and pure quadratic effects are active. Then a multiple regression model containing the main effects only or a saturated model containing both main and pure quadratic effects will produce coefficients that are BLUE assuming no third order effects.

The analysis becomes more challenging if both two-factor interactions and pure quadratic effects are active because these may have substantial correlations. Figure 1 shows the column correlations for the design shown in Table 1. Columns AA through FF represent the pure quadratic terms. Columns AB through EF show the correlations for the two-factor interaction columns. Note that the main effects are uncorrelated with each other and all second-order effects.

The properties of the design do not depend on the response y , but in order to illustrate how the analysis might proceed, we generated the column for y in Table 1 using the formula, $y = 20 + 4A + 3B + -2C + -D + 5BC + 6A^2 + \epsilon$, where the errors, ϵ , are independently normally distributed with mean zero and variance 1. We did an all subsets regression for all models up to 7 terms including all terms up through second order for consideration in the “full model”. The model with the minimum corrected Akaike’s Information (AIC_c) criterion, 70.63, included terms A , B , C , BC and A^2 . The next best model added the main effect of D yielding the true model, for which $AIC_c = 71.25$.

We recommend all subset regression for models with up to $m + 1$ terms, where the full set of model terms is comprised of all first- and second-order effects.

Note that if there are multiple active pure quadratic and two-factor interaction terms, there may be model confounding. That is, two or more models may yield identical predicted y -vectors. In such cases, the all subsets regression will identify the confounded models and additional runs will be necessary to resolve the confounding.

6 Summary

We have introduced a class of screening designs for quantitative factors that are robust to active second order effects. We have also provided an algorithm for generating these designs for any number of factors. Our designs have the minimum possible number of runs for estimating both the main and pure quadratic effects of the factors. For each factor pair they project to a 3×3 factorial design which is D-optimal for the full quadratic model.




Fig. 1: Absolute values of column correlations of terms through 2nd order for the six factor robust screening design

References

- Box, G. E. P. and J. S. Hunter (2008). The 2^{k-p} fractional factorial designs. *Technometrics* 3, 449–458.
- Cheng, S. W. and C. F. J. Wu (2001). Factor screening and response surface exploration (with discussion). *Statistica Sinica* 11, 553–604.
- Tsai, P. W., Gilmour, S. G., and R. Mead (2000). Projective three-level main effects designs robust to model uncertainty. *Biometrika* 87, 467–475.
- Jones, B., and C. J. Nachtsheim (2009). D-efficient designs with minimal aliasing. Technical report, University of Minnesota. 467–475.
- Meyer, R. K. and C. J. Nachtsheim (1995). The coordinate-exchange algorithm for constructing exact optimal experimental designs. *Technometrics* 37, 60–69.

D-Optimal Design for a Five-Parameter Logistic Model

Zorayr Manukyan and William F. Rosenberger

Abstract We explore the D -optimal design for a five-parameter logistic model, which includes a shape parameter to handle asymmetries, and two threshold parameters to account for situations where the asymptotes are not at 0 and 1. The optimal design is five points, including points at $-\infty$ and ∞ representing the thresholds. We compare the efficiencies of the optimal designs arising from the two- and five-parameter models. We find a significant loss of efficiency when the two-parameter model is used on data generated from the five-parameter model.

1 Introduction

The standard two-parameter logistic model is often used to characterize dose-response relationships. Let $d \in \Omega_d$ be a dose-level and y be a binary response. For example, y could be toxicity or no toxicity, death or no death, cure or no cure, etc. For a location parameter α and a scale parameter $\beta > 0$, we have

$$P(Y = 1|d) = F\left(\frac{d - \alpha}{\beta}\right), \quad (1)$$

where

$$F(x) = \frac{e^x}{(1 + e^x)}. \quad (2)$$

Zorayr Manukyan

The EMMES Corporation, 401 N. Washington St., Rockville, MD 20850 USA e-mail: zmanukyan@emmes.com

William F. Rosenberger

Department of Statistics, George Mason University, 4400 University Dr., MS 4A7, Fairfax, VA 22030-4444 USA e-mail: wrosenbe@gmu.edu

When d is the log of a dose level, Ω_d is some subset of the real line. In some cases, Ω_d may be restricted due to ethical or other constraints, but we do not deal with restricted design spaces in this paper (see Biedermann, Dette, and Pepelyshev 2006 for a careful treatment of restricted designs spaces). It is well-known that the D -optimal design for the model in (1) is a two-point design, putting equal weights at the 17.6th and 82.4th percentiles (Kalish and Rosenberger 1978). Since these percentiles depend on the unknown parameters $\boldsymbol{\theta}_2 = (\alpha, \beta)$, they must be evaluated at specified or *local* values of the parameter, leading to a *locally optimal design*. Such designs are not useful in practice, but can be used to gauge the efficiency of alternate designs, and can give insight into the relationship of the parameters in designing an experiment.

The two-parameter logistic model is not rich enough to incorporate asymmetry in the dose-response curves, or to account for dose-response problems that do not have an asymptote at 0 or 1. For example, there may be no dose at which the cure rate is 0 (e.g., consider the placebo effect), and often doses do not exist that are 100 percent effective or 100 percent toxic. (In rare cases, the dose-response curve may not be monotonically non-increasing. We do not deal with that problem.) In this paper, we examine a five-parameter logistic model which has a shape parameter that can handle asymmetries and threshold parameters for when the asymptotes are not 0 or 1. The model is described as follows:

$$P(Y = 1|d) = (c_{\max} - c_{\min}) \left[F \left(\frac{d - \alpha^*}{\beta^*} \right)^\gamma \right] + c_{\min}, \quad (3)$$

where F is defined by (2). The γ parameter influences the shape and captures asymmetries. The parameters c_{\min} and c_{\max} determine the asymptotes of the dose-response curve at $d = -\infty$ and $d = \infty$, respectively.

In this paper, we focus on the D -optimality criterion, which maximizes the log determinant of the information matrix generated by the likelihood for model (3), and consequently minimizes the volume of the confidence ellipsoid for the joint estimation of $\boldsymbol{\theta}_5 = (c_{\max}, c_{\min}, \gamma, \alpha^*, \beta^*)$. We compare the relative efficiency of the D -optimal design for the two- and five-parameter models. In an attempt to determine the effects of model misspecification, we also generate data from the five-parameter model and fit the two-parameter model to these data, obtaining the maximum likelihood (ML) estimators of α and β . We then find the optimal design for the two-parameter model at the local values of the ML estimates and compare its efficiency to that of the true five-parameter model.

2 Methods

Let $\boldsymbol{\xi}$ be a design measure; i.e., dose levels d_1, \dots, d_K with associated nonnegative weights w_1, \dots, w_K , where $\sum_{i=1}^K w_i = 1$, where the dimension K is a priori unknown. The information matrix is defined as $\mathbf{M}(\boldsymbol{\xi}, \boldsymbol{\theta}) = \sum_{i=1}^K w_i \mathbf{I}(d_i, \boldsymbol{\theta})$, where

$I(d_i, \theta)$ is the information contributed by a single dose level d_i . For the two-parameter model in (1), we use M_2, I_2 , and θ_2 to refer to M, I , and θ , respectively. Similarly, for the five-parameter model in (3), we use M_5, I_5 and θ_5 , respectively. The form of $I_5(d, \theta_5)$, which is a 5×5 matrix, is given in the Appendix. In general, the D -optimal design will be the design measure ξ^* which maximizes $\Psi(\xi, \theta) = \log |\det(I(\xi, \theta))|$. For the standard two-parameter logistic model in (1), the information matrix is given by $M_2(\xi, \theta_2) = \sum_{i=1}^K w_i I_2(d_i, \theta_2)$, where

$$I_2(d, \theta_2) = \frac{e^z}{\beta^2(1 + e^z)^2} \begin{bmatrix} 1 & z \\ z & z^2 \end{bmatrix},$$

and $z = (d - \alpha)/\beta$. The D -optimal design is ξ_2^* , which maximizes $\Psi_2(\xi, \theta_2) = \log |\det(I_2(\xi, \theta_2))|$ for all K .

Define a design $\bar{\xi}$ to put unit mass at dose d . Given a design ξ , consider a perturbation towards the point d :

$$\xi' = (1 - \lambda)\xi + \lambda\bar{\xi}.$$

Due to the additive structure of the information matrix M , we have

$$M(\xi') = (1 - \lambda)M(\xi) + \lambda M(\bar{\xi}).$$

Then the directional derivative in the direction of $\bar{\xi}$ is given by:

$$\phi(d, \xi) = \lim_{\lambda \rightarrow 0^+} \frac{1}{\lambda} \left((1 - \lambda)M(\xi) + \lambda M(\bar{\xi}) \right) - \Psi \left(M(\bar{\xi}) \right).$$

We employ the graphical technique of plotting directional derivatives that derives from Kiefer and Wolfowitz (1960) general equivalence theorem. We begin by finding an arbitrary K -point ($n \geq 5$) optimal design and then perturb the design toward the design point with a large directional derivative value, and continue until the directional derivatives satisfy the general equivalence criterion for D -optimality, given in Atkinson, Donev, and Tobias (2007, p. 122): under mild assumptions, the most important of which are the compactness of the design space and the convexity and differentiability of Ψ , the following three conditions are equivalent:

- the design ξ^* maximizes Ψ ;
- the maximum of $\phi(d, \xi)$ does not exceed 0;
- the directional derivative $\phi(d, \xi)$ achieves its maximum at the design points.

MATLAB6R12 was used to compute the directional derivatives and optimize the criterion for a fixed number of design points. Programs are available from the first author upon request.

Let ξ_5^* be the D -optimal design for the five-parameter model. Let $\xi_{2,ML}^*$ be the D -optimal design evaluated at local ML estimates computed from data generated from the five-parameter model. We define the efficiency of the five-parameter model at a parameter vector θ_5 to be $\Psi_5(\xi^*, \theta_5)$. The relative efficiency is computed for

Table 1: Parameterizations of the numerical example in Figure 1

Parameterization	Model	c_{\min}	c_{\max}	α	β	γ
a	5-parameter	0.1	0.9	843.0	222.0	2.0
	2-parameter MLE			1052.8	341.8	
	2-parameter			843.0	222.0	
b	5-parameter	0.1	0.9	843.0	222.0	3.0
	2-parameter MLE			1209.3	392.5	
	2-parameter			843.0	222.0	
c	5-parameter	0.2	0.7	843.0	222.0	0.6
	2-parameter MLE			775.7	622.3	
	2-parameter			843.0	222.0	

the five-parameter model evaluated at the optimal design obtained from the two-parameter model, relative to the D -optimal design for the five-parameter model; i.e.,

$$\frac{\Psi_5(\xi, \theta_5)}{\Psi_5(\xi_5^*, \theta_5^*)}, \tag{4}$$

where $\xi = \xi_2^*$ and $\xi = \xi_{2,ML}^*$. Note that the numerator cannot be computed directly since the information matrix is singular at a two-point design. We therefore established a five-point design with negligible weights for three of the points.

Our numerical examples are based on an actual dose-response clinical trial of a chemotherapeutic agent in leukemia, conducted at the University of Maryland Medical School (see Haines, Perevozskaya, and Rosenberger 2003). Here $\Omega_d = (100, 300, 600, 1200)$ (in mg). From the data on 34 patients, the two-parameter logistic model yielded an MLE of α of 843.0 and an MLE of β of 222.0. We generated 7000 data points from the five-parameter model and plugged these data into the two-parameter model to yield $\hat{\alpha}$ and $\hat{\beta}$. Parameter values for each of our numerical examples are given in Table 1.

3 Results

The left side of Figure 1 shows the three dose-response curves for Parameterizations a, b, and c from Table 1. In Parameterization c, we have established a particularly unlikely model, but it illustrates what happens when the thresholds are in the interior, and γ is less than 1. The right side gives the directional derivative plots for the three dose-response curves. In each case, the optimal design points occur when the directional derivatives equal zero. Note that the directional derivatives are less than zero at all other points, hence the general equivalence criterion is satisfied.

Table 2 gives the optimal design for each of the three parameterizations in Table 1. (Note that we did not take the logarithm of dose in these examples.) One can see that the D -optimal design in each case puts points at $\pm\infty$ and three points in

Table 2: *D*-Optimal designs for the two- and five-parameter logistic models, corresponding to Figure 1

Parameterization Model		Design Points	Weights
a	5-parameter	$(-\infty, 606, 963, 1415, \infty)$	$(0.2, 0.2, 0.2, 0.2, 0.2)$
	2-parameter MLE	$(500, 1600)$	$(0.5, 0.5)$
	2-parameter	$(500, 1185)$	$(0.5, 0.5)$
b	5-parameter	$(-\infty, 744, 1051, 1486, \infty)$	$(0.2, 0.2, 0.2, 0.2, 0.2)$
	2-parameter MLE	$(630, 1800)$	$(0.5, 0.5)$
	2-parameter	$(500, 1185.6)$	$(0.5, 0.5)$
c	5-parameter	$(-\infty, 50, 660, 1135, \infty)$	$(0.2, 0.2, 0.2, 0.2, 0.2)$
	2-parameter MLE	$(-100, 1800)$	$(0.5, 0.5)$
	2-parameter	$(500, 1185)$	$(0.5, 0.5)$

Table 3: Relative efficiencies of designs arising from the two- and five-parameter models

Parameterization Model		$\Psi_5(\xi, \theta_5)$	Relative Efficiency
a	5-parameter	-33.94	
	2-parameter MLE	-41.47	0.82
	2-parameter	-41.46	0.82
b	5-parameter	-35.21	
	2-parameter MLE	-42.92	0.82
	2-parameter	-42.88	0.82
c	5-parameter	-35.92	
	2-parameter MLE	-43.46	0.83
	2-parameter	-43.31	0.83

the middle with equal weights. The resulting design is always five points. As γ increases, the tendency is for the three middle points to be shifted to the right. The three middle points are always bounded by the optimal design from the 2-parameter MLE.

Finally, in Table 3, we compute the efficiencies and relative efficiencies, as defined in (4). It is clear that there is a significant loss of efficiency by using a two-parameter model rather than the five-parameter model, in the range of 17 – 18 percent.

4 Discussion

In this brief report, we have shown the effects of model misspecification when the asymptotes of the logistic model are not 0 and 1. The two-parameter model ignores the points $\pm\infty$, which are intuitively necessary to pick up the information on the upper and lower threshold asymptotes. The effects of γ are not as clear and merit more inspection.

Bayesian optimal designs can be obtained by establishing a prior distribution on $\boldsymbol{\theta} \in \Theta$, $g(\boldsymbol{\theta})$, and integrating over the prior distribution to optimize the average; i.e., find $\boldsymbol{\xi}$ that maximizes

$$\int_{\Theta} \log \det |\mathbf{I}(\boldsymbol{\xi}, \boldsymbol{\theta})| g(\boldsymbol{\theta}) d\boldsymbol{\theta}.$$

This requires five-dimensional integration. In the past, this was computationally infeasible, but we were able to use Markov chain Monte Carlo methods quite efficiently. Determining an appropriate prior distribution for the problem is nontrivial. We typically used independent normal and beta distributions.

Of greater interest are adaptive or sequential designs in which the posterior is updated after each response is recorded and a new design point selected based on optimizing the updated criterion. This has application in clinical trials and extends the work of Haines, Perevozskaya, and Rosenberger (2003) from the two-parameter logistic distribution to the five-parameter logistic distribution. After n responses, let $g(\boldsymbol{\theta} | \mathcal{D}_n)$ be the posterior distribution incorporating the data \mathcal{D}_n observed thus far and let $\boldsymbol{\xi}_n$ be the design measure. Then for the $(n + 1)$ th dose assignment, we find the dose d that maximizes

$$\int_{\Theta} \log \det |n\mathbf{I}(\boldsymbol{\xi}_n, \boldsymbol{\theta}) + \mathbf{I}(d, \boldsymbol{\theta})| g(\boldsymbol{\theta} | \mathcal{D}_n) d\boldsymbol{\theta}.$$

Some preliminary results can be found in Manukyan (2009) and is being prepared for publication elsewhere.

5 Appendix: Information Matrix

The entries of the information matrix $\mathbf{I}_5(d, \boldsymbol{\theta}_5)$ for the parameter $\boldsymbol{\theta}_5 = (c_{\min}, c_{\max}, \gamma, \alpha^*, \beta^*)$ at the dose level d , under the model framework introduced in (3), are given by:

$$\begin{aligned} I_{1,1}(\boldsymbol{\theta}) &= [1 - F(z)^\gamma]^2 \frac{1}{p(1-p)}, \\ I_{2,2}(\boldsymbol{\theta}) &= F(z)^{2\gamma} \frac{1}{p(1-p)}, \\ I_{3,3}(\boldsymbol{\theta}) &= \left[\Delta c F(z)^\gamma \ln F(z) \right]^2 \frac{1}{p(1-p)}, \\ I_{4,4}(\boldsymbol{\theta}) &= \left[\frac{\Delta c}{\beta^*} \gamma F(z)^{\gamma-1} F'(z) \right]^2 \frac{1}{p(1-p)}, \\ I_{5,5}(\boldsymbol{\theta}) &= z^2 \cdot I_{4,4}, \end{aligned}$$


$$\begin{aligned}
I_{1,2}(\boldsymbol{\theta}) &= I_{2,1}(\boldsymbol{\theta}) = \left[1 - F(z)^\gamma\right] F(z)^\gamma \frac{1}{p(1-p)}, \\
I_{1,3}(\boldsymbol{\theta}) &= I_{3,1}(\boldsymbol{\theta}) = \Delta c \ln F(z) I_{1,2}, \\
I_{1,4}(\boldsymbol{\theta}) &= I_{4,1}(\boldsymbol{\theta}) = -\Delta c \frac{\gamma}{\beta^*} F(z)^{\gamma-1} F'(z) [1 - F(z)^\gamma] \frac{1}{p(1-p)}, \\
I_{1,5}(\boldsymbol{\theta}) &= I_{5,1}(\boldsymbol{\theta}) = z I_{1,4}(\boldsymbol{\theta}), \\
I_{2,3}(\boldsymbol{\theta}) &= I_{3,2}(\boldsymbol{\theta}) = \Delta c F(z)^{2\gamma} \ln F(z) \frac{1}{p(1-p)}, \\
I_{2,4}(\boldsymbol{\theta}) &= I_{4,2}(\boldsymbol{\theta}) = -\frac{\Delta c}{\beta^*} \gamma F(z)^{2\gamma-1} F'(z) \frac{1}{p(1-p)}, \\
I_{2,5}(\boldsymbol{\theta}) &= I_{5,2}(\boldsymbol{\theta}) = z I_{2,4}(\boldsymbol{\theta}), \\
I_{3,4}(\boldsymbol{\theta}) &= I_{4,3}(\boldsymbol{\theta}) = -\frac{\Delta c^2}{\beta^*} \gamma F(z)^{2\gamma-1} F'(z) \ln F(z) \frac{1}{p(1-p)}, \\
I_{3,5}(\boldsymbol{\theta}) &= I_{5,3}(\boldsymbol{\theta}) = I_{3,4}(\boldsymbol{\theta}) z, \\
I_{4,5}(\boldsymbol{\theta}) &= I_{5,4}(\boldsymbol{\theta}) = \left[\frac{\Delta c}{\beta^*} \gamma F(z)^{\gamma-1} F'(z)\right]^2 z \frac{1}{p(1-p)},
\end{aligned}$$

where $z = (d - \alpha^*)/\beta^*$, $\Delta c := c_{max} - c_{min}$ and $p = P(Y = 1|d)$.


Acknowledgements The authors thank two excellent referees for their detailed comments. One of the referees found a major mistake, and the authors are deeply appreciative.

References


- Atkinson, A. C., A. N. Donev, and R. D. Tobias (2007). *Optimum Experimental Designs, with SAS*. Oxford: Oxford University Press.
- Biedermann, S., H. Dette, and A. Pepelyshev (2006). Some robust design strategies for percentile estimation in binary response models. *Canadian Journal of Statistics* 34, 603–622.
- Haines, L. M., I. Perevozskaya, and W. F. Rosenberger (2003). Bayesian optimal designs for Phase I clinical trials. *Biometrics* 59, 591–600.
- Kalish, L. A. and J. L. Rosenberger (1978). Optimal designs for the estimation of the logistic function. Technical Report 33, Pennsylvania State University.
- Kiefer, J. and J. Wolfowitz (1960). The equivalence of two extremum problems. *Canadian Journal of Mathematics* 12, 363–366.
- Manukyan, Z. (2009). *Sequential Designs for Estimating Toxicity and Efficacy in a Dose-Response Setting*. Ph. D. thesis, George Mason University.




(a) Parameterization a




(b) Directional Derivatives




(c) Parameterization b



(d) Directional Derivatives



(e) Parameterization c



(f) Directional Derivatives

Fig. 1: Dose-response curves and directional derivatives for the parameterizations in Table 1

Sequential Barycentric Interpolation

Hugo Maruri-Aguilar and Paula Camelia Trandafir

Abstract Polynomial interpolators may exhibit oscillating behaviour which often makes them inadequate for modelling functions. A well-known correction to this problem is to use Chebyshev design points. However, in a sequential strategy it is not very clear how to add points, while still improving polynomial interpolation. We present a sequential design alternative by allocating an extra observation where the difference between consecutive interpolators is largest. Our proposal is independent of the response and does not require distributional assumptions. In simulated examples, we show the good interpolation performance of our proposal and its asymptotical convergence to the Chebyshev distribution.

1 Introduction

In classical optimal design theory, a connection can be established between a certain optimality criterion, a linear polynomial model and a design constructed with zeros of T_n , i.e. Chebyshev points. This connection was first noted by Studden (1968). Subsequent research led to various articles and books which also exhibited designs whose points are zeros of Jacobi, Laguerre, or Hermite polynomials, among others. Among some of the better known examples of the connection between polynomial models and Chebyshev points, we may cite, classified by optimality criterion, Pukelsheim and Torsney (1991) for A -optimal designs, Fedorov (1972, pp.85), Pázman (1986, pp.178), Pukelsheim (1993, pp.214–216) and Karlin and Studden (1966) for D -optimality, Dette (1993b) for D_s -optimality; and Dette (1993a) and Heiligers (1998) for E -optimality.

In the analysis of computer experiments, usually there is no random error associated with the response and models interpolate the observed response values. A variety of models are available for the analysis of such experiments. Spline models can be used, but also models based on radial bases and kriging have become widely used, see Müller (2001), O’Hagan (2006) and Fedorov and Müller (2007).

Hugo Maruri-Aguilar

School of Mathematical Sciences, Queen Mary, University of London, Mile End Road, London E1 4NS UK, e-mail: H.Maruri-Aguilar@qmul.ac.uk

Paula Camelia Trandafir

Departamento de Estadística e Investigación Operativa, Universidad de Valladolid, Prado de la Magdalena s/n, 47005, Valladolid, Spain, e-mail: camelia@eio.uva.es

We are concerned with sequential polynomial interpolation. Polynomials are simple and potentially effective models, often with a straightforward interpretation. However, they have a tendency to oscillate between design points. Those oscillations are called the *Runge phenomenon*, which in some cases can only get worse as the number of data points increases. It is a well-known classical result that the oscillation caused by Runge's phenomenon can be minimized by interpolating at the Chebyshev nodes. In Epperson (1987), additional conditions for mitigating the phenomenon are studied. In the literature of approximation theory there are several proposals which use Chebyshev points and may be used to interpolate. In Boyd and Ong (2009) and Boyd and Xu (2009) the authors use a subsample of the uniform distribution to generate "mock Chebyshev" (i.e. approximately Chebyshev) points, while in Platte and Driscoll (2005), interpolation points are selected using more general polynomial approximation techniques. The methods produce good interpolation results, although the strategies are non-sequential.

The aim of the present paper is to introduce a sequential design strategy for univariate polynomial interpolation. Our proposal is based on an equivalent form of the Lagrange interpolator called the barycentric interpolator. In Section 2 we review the Lagrange and barycentric interpolators. In Section 3 we present our sequential problem, introduce a design algorithm and prove that it does not depend on response values. Section 4 presents simulated examples to evaluate the performance of our algorithm. Conclusions, future work and a conjecture are presented in Section 5.

2 Barycentric Lagrange Interpolation

Let $D_n = \{d_1, \dots, d_n\}$ be a design of n distinct univariate points. The values f_1, \dots, f_n are observations, one for every design point. Those values are assumed to be evaluations of a deterministic (but unknown) function which is to be interpolated. A classic solution is the Lagrange interpolator $g_n(x) = \sum_{j=1}^n f_j \prod_{i=1, i \neq j}^n \frac{x-d_i}{d_j-d_i}$. Lagrange interpolation exhibits numerical and computational drawbacks and an alternative form of it is available, known as the barycentric interpolator:

$$g_n(x) = m_n(x) \sum_{j=1}^n \frac{w_{n,j} f_j}{x - d_j}, \quad (1)$$

with barycentric weights defined by $w_{n,j} = \{\prod_{i=1, i \neq j}^n (d_j - d_i)\}^{-1}$ and $m_n(x) = \prod_{i=1}^n (x - d_i)$. The first subindex in $w_{n,j}$ denotes the number of design points used for computing it, while the second subindex relates the weight to a design point.

3 Sequential Interpolation

Barycentric formulæ allow sequential interpolation of data, that is, adding an extra observation to an existing data set and updating the barycentric interpolator. Sequen-

tial updating can be made part of an adaptive procedure, in which using information from the interpolation process helps in selecting a new design point.

3.1 Response-based Update

Consider two interpolators, one of which is considered to be more accurate than the other, as it is built with one extra observation. We postulate that the difference between them can be used as a guide to future experimentation. In other words, the more accurate interpolator may be used to validate the less accurate fit and to insert another design point where this difference is largest over an arbitrary design region. We set the design region to $[0, 1]$ but it can be adapted to other design region $[a, b]$.

The starting point is $g_n(x)$, the barycentric interpolator as defined in Equation (1). Consider an extra design point d_{n+1}^* and its corresponding observation. We term d_{n+1}^* a *dummy point* and only require it to be different from existing design points. Denote by $G_{n+1}(x)$ the interpolator constructed with the temporary design consisting of the original design plus the dummy point, $D_n \cup d_{n+1}^*$. A new design point is selected according to

$$d_{n+1} = \arg \max_{x \in [0,1]} |G_{n+1}(x) - g_n(x)|. \quad (2)$$

After the search, the dummy point d_{n+1}^* is discarded and the original design D_n is augmented to $D_{n+1} = D_n \cup d_{n+1}$. The search problem is well posed, i.e. maximisation of a bounded function over a closed compact set.

Example 1. Consider the function $f(x) = 1/\{1 + 25(2x - 1)^2\}$, which is to be interpolated using the design $D_{10} = \{0, \frac{1}{5}, \dots, 1\}$. Let $g_{10}(x)$ be the interpolator function constructed with observations of $f(x)$ at D_{10} . The dummy point $d_{11}^* = \frac{1}{2}$ is added to build an updated interpolator $G_{11}(x)$. The next point is selected where the absolute difference between the interpolators $G_{11}(x)$ and $g_{10}(x)$ is largest over $[0, 1]$; this occurs at the points 0.0325 and 0.9675. Any of these two points can be added to D_{10} and d_{11}^* is discarded.

3.2 Sequential Design Algorithm

The sequential design procedure described above simplifies to a response-independent alternate maximization and update of $m_n(x)$. We now describe the algorithm.

Input An initial design D_n of n distinct points $d_1, \dots, d_n \subset [0, 1]$; a number k of extra design points required.

Output A set of additional runs $d_{n+1}, \dots, d_{n+k} \subset [0, 1]$.

Initialization Set $m_n(x) := \prod_{i=1}^n (x - d_i)$; set $j := 0$.

- Step 1** Maximize $|m_{n+j}(x)|$ with respect to x , in the interval $[0, 1]$, i.e. let $d_{n+j+1} := \arg \max_{x \in [0,1]} |m_{n+j}(x)|$.
- Step 2** Update $m_{n+j+1}(x) := m_{n+j}(x)(x - d_{n+j+1})$ and $j := j + 1$. If $j < k$, repeat from Step 1.

The algorithm does not depend on actual response values observed, but only on the design points. Additionally, it does not depend on the actual location of the dummy point. These two characteristics are implied by Theorem 1, which is proven in the Appendix.

Theorem 1. For $n > 0$, let $g_n(x)$ and $G_{n+1}(x)$ for $n > 0$, be two barycentric interpolators, where $g_n(x)$ is defined as in Equation (1); and $G_{n+1}(x)$ is constructed with an additional dummy design point d_{n+1}^* . Then

$$G_{n+1}(x) - g_n(x) = m_n(x) \sum_{j=1}^{n+1} w_{n+1,j} f_j. \quad (3)$$

The right hand side of Equation (3) is the product of $m_n(x)$, which depends on x and on D_n (but not on the dummy d_{n+1}^*), and a second quantity $\sum_{j=1}^{n+1} w_{n+1,j} f_j$ that depends on design points and responses (including dummy data), but not on x and thus it can be ignored when searching for the new design point. Theorem 1 makes the search for a new design point independent of the response, indeed it makes Equation (2) equivalent to $d_{n+1} = \arg \max_{x \in [0,1]} |m_n(x)|$.

Example 2. Consider again the design of Example 1. The sequential algorithm is applied for $k = 10$ extra runs and points d_{11} to d_{20} are sequentially obtained: 0.0325, 0.9684, 0.9335, 0.0662, 0.8306, 0.1677, 0.4999, 0.0099, 0.9902 and 0.2813.

A special condition arises from Equation (3), when $\sum_{j=1}^{n+1} w_{n+1,j} f_j = 0$ holds. This implies that $G_{n+1}(x) - g_n(x) \equiv 0$, in other words, that the dummy point d_{n+1}^* does not update the interpolator and consequently, this step does not yield information for the next design point d_{n+1} . This condition appears, for instance, when all f_j values are equal. This could occur when sampling a constant function or a periodic function at the same point in every period. A different instance appears when response data truly come from a polynomial of degree at most $n - 1$. In any of the above situations, any point in the interval $[0, 1]$ could be selected as the new design point. However, in all the examples we tried, none of them occurred and we suggest that they should not be a cause of concern.

4 Performance and Large Sample Properties

In this section we first evaluate the performance of our sequential design strategy for interpolation. We then study the large sample properties of our sequential designs.

4.1 Interpolating Performance

The accuracy of polynomial interpolators with our sequential design algorithm was assessed in a simulation study. The following four functions with domain $[0, 1]$ were used as test functions: $s_1(x)$ is the function of Example 1; $s_2(x) = \frac{1}{20} \exp(u^{1/3}) \sin(u/2) \phi(u)$ with $\phi(u)$ the Heaviside unit step function and $u = 60x - 30$; $s_3(x) = \sin(10x)$ and $s_4(x) = 2 \frac{1 - \cos(v)}{v^2}$ with $v = 35x - 15$. The functions were selected to exhibit features which are not easy for modeling with polynomials, such as flat regions followed by regions with sharp change, or periodic behaviour.




Fig. 1: Maximum distance between simulators and barycentric interpolators, plotted against number of extra points k added.

From a uniform design of size ten, fifteen points were sequentially added, totalling 25 points. At every step, barycentric interpolators were fitted independently for each function. The maximum distance between the true function $s_1(x), \dots, s_4(x)$ and its barycentric interpolator, over the design region, was recorded. Decreasing values of this distance show good approximation, while increasing values point to the presence of the Runge phenomenon.

The results are plotted in Figure 1, where a decreasing trend is evident, thus showing good approximation to simulators for all cases. Convergence to the true function s_3 was faster than the other cases, while convergence was slowest for s_2 .

4.2 Large Sample Properties

The points generated with our algorithm cluster in the borders of the design region. We studied whether the points converge asymptotically to a known distribution.

To study large sample behaviour, points were sequentially added to each of the following eight initial designs of size n : uniform designs $0, \frac{1}{n-1}, \dots, 1$ (termed UI) and $\frac{1}{n+1}, \dots, \frac{n}{n+1}$ (termed UII); first n points of Sobol’s space filling sequence (termed S), see Bratley and Fox (1988); Chebyshev type I and II points (CI and CII, respectively), see Berrut and Trefethen (2004); the designs labelled TI and TII

were generated by transforming UI and UII to the symmetric triangular distribution with mode in $\frac{1}{2}$; and a design with random points (termed R). We used initial design sizes $n = 5, 10, 35, 50, 100, 150$, in each case sequentially adding points with our algorithm up to one thousand. Two statistics were computed: a) Quantile-Quantile (QQ) plot and b) goodness of fit Kolmogorov-Smirnov (KS) statistic. The Beta distribution $\beta(\frac{1}{2}, \frac{1}{2})$ (also known as the Chebyshev or arcsine distribution) was used in computations. This choice was suggested by the literature on polynomial interpolation convergence (Berrut and Trefethen 2004).




Fig. 2: QQ plot for design UI.

We show results for $n = 50$, which are representative of the results for other initial sizes. Figure 2 shows the QQ plot for UI, which converges to the Chebyshev distribution with 20 extra points. QQ plots for other initial designs exhibit similar pattern.




Fig. 3: KS statistic vs. design size. The left-hand panel is a close up view of the right-hand panel.

The KS statistic “grouped” designs according to their initial performance, with best (i.e. low) values obtained by Chebyshev points (CI, CII). In second place were uniform designs (UI, UII, S), followed by the random design R. The worst values were observed for designs with points clustered in the centre of the design region (TI, TII). Figure 3 shows the evolution of the KS statistic for one design for each of the observed “groups”: CII, R, S and TII. After adding about twice as many points

as the initial design size, the designs behave similarly, showing a non-monotonic decreasing linear trend (in the log-log scale) for the KS statistic. Simulation results suggest a value for the slope of the linear trend between -0.8 and -1 , see also Figure 3.

5 Discussion and Future Work

We introduced a univariate sequential adaptive design algorithm. In the examples we tried, the algorithm produced good points for polynomial interpolation, which converged rapidly to the Chebyshev distribution and lead to the following claim:

Conjecture 1. For any initial design in $[0, 1]$, as the number of extra points k tends to infinity, the algorithm of Section 3.2 produces samples from the distribution $\beta(\frac{1}{2}, \frac{1}{2})$; and the KS statistic is of order $O(k^\alpha)$, with α a suitable constant.

Acknowledgements The first author was supported by UK Research Council grant EP/D048893/1 (MUCM), and the second author by Spain's *Ministerio de Ciencias e Innovación* (ECO2008-02358/ECON) and *Conserjería de Educación de la Junta de Castilla y León* (VA004B08).

Appendix A: Proof of Theorem 1

The barycentric interpolator $G_{n+1}(x)$ is

$$G_{n+1}(x) = m_n(x)(x - d_{n+1}^*) \left(\sum_{j=1}^n \frac{w_{n+1,j} f_j}{x - d_j} + \frac{w_{n+1,n+1} f_{n+1}}{x - d_{n+1}^*} \right),$$

where barycentric weights $w_{n+1,j}$ are computed using design and dummy points. We have that $G_{n+1}(x) - g_n(x) = m_n(x)A$, where

$$A = (x - d_{n+1}^*) \left(\sum_{j=1}^n \frac{w_{n+1,j} f_j}{x - d_j} + \frac{w_{n+1,n+1} f_{n+1}}{x - d_{n+1}^*} \right) - \sum_{j=1}^n \frac{w_{n,j} f_j}{x - d_j}.$$

We now show that A does not depend on x . After simplifying and using the updating formula of barycentric weights $w_{n+1,j} = w_{n,j}(d_j - d_{n+1}^*)^{-1}$, we have

$$\begin{aligned} A &= \sum_{j=1}^n \frac{f_j}{x - d_j} \left((x - d_{n+1}^*) \frac{w_{n,j}}{(d_j - d_{n+1}^*)} - w_{n,j} \right) + w_{n+1,n+1} f_{n+1} \\ &= \sum_{j=1}^n \frac{w_{n,j} f_j}{(d_j - d_{n+1}^*)} + w_{n+1,n+1} f_{n+1} = \sum_{j=1}^{n+1} w_{n+1,j} f_j. \end{aligned}$$

References

- Berrut, J.-P. and L. N. Trefethen (2004). Barycentric Lagrange interpolation. *SIAM Review* 46, 501–517.
- Boyd, J. P. and J. R. Ong (2009). Exponentially-convergent strategies for defeating the Runge phenomenon for the approximation of non-periodic functions. I. Single-interval schemes. *Communications in Computational Physics* 5, 484–497.
- Boyd, J. P. and F. Xu (2009). Divergence (Runge phenomenon) for least-squares polynomial approximation on an equispaced grid and Mock-Chebyshev subset interpolation. *Applied Mathematics and Computation* 210, 158–168.
- Bratley, P. and B. L. Fox (1988). ALGORITHM 659 Implementing Sobol’s quasirandom sequence generator. *ACM Transactions on Mathematical Software* 14, 88–100.
- Dette, H. (1993a). A note on E-optimal designs for weighted polynomial regression. *Annals of Statistics* 21, 767–771.
- Dette, H. (1993b). On a mixture of the D- and D_1 -optimality criterion in polynomial regression. *Journal of Statistical Planning and Inference* 35, 233–249.
- Epperson, J. F. (1987). On the Runge example. *American Mathematical Monthly* 94, 329–341.
- Fedorov, V. V. (1972). *Theory of Optimal Experiments*. New York: Academic Press.
- Fedorov, V. V. and W. G. Müller (2007). Optimum design for correlated fields via covariance kernel expansions. In J. Lopez-Fidalgo, J. M. Rodriguez-Diaz, and B. Torsney (Eds.), *mODA 8—Advances in Model-oriented Design and Analysis*, pp. 57–66. Heidelberg: Physica-Verlag.
- Heiligers, B. (1998). E-optimal designs for polynomial spline regression. *Journal of Statistical Planning and Inference* 75, 159–172.
- Karlin, S. and W. J. Studden (1966). Optimal experimental designs. *Annals of Mathematical Statistics* 37, 783–815.
- Müller, W. G. (2001). *Collecting Spatial Data, rev. edn*. Heidelberg: Physica-Verlag.
- O’Hagan, A. (2006). Bayesian analysis of computer code outputs: a tutorial. *Reliability Engineering & System Safety* 37, 1290–1300.
- Pázman, A. (1986). *Foundations of Optimum Experimental Design*. Dordrecht: Reidel.
- Platte, R. B. and T. A. Driscoll (2005). Polynomials and potential theory for Gaussian radial basis function interpolation. *SIAM Journal on Numerical Analysis* 43, 750–766.
- Pukelsheim, F. (1993). *Optimal Design of Experiments*. New York: Wiley.
- Pukelsheim, F. and B. Torsney (1991). Optimal weights for experimental designs on linearly independent support points. *Annals of Statistics* 19, 1614–1625.
- Studden, W. J. (1968). Optimal designs on Tchebycheff points. *Annals of Mathematical Statistics* 39, 1435–1447.

Some Considerations on the Fisher Information in Nonlinear Mixed Effects Models

Tobias Mielke and Rainer Schwabe

Abstract The inverse of the Fisher Information Matrix is a lower bound for the covariance matrix of any unbiased estimator of the parameter vector and, given this, it is important for the construction of optimal designs. For normally distributed observation vectors with known variance, the Fisher Information can be easily constructed. For nonlinear mixed effects models, the problem of the missing closed-form solution of the likelihood function carries forward to the calculation of the Fisher Information matrix. The often used approximation of the Fisher Information by linearizing the model-function in the fixed effects case is generally not reliable, as will be shown in this article.

1 Introduction

In population pharmacokinetic studies, the observations of different individuals are often assumed to follow one common function with small differences, which are generated by random individual parameters. One main interest in these studies lies in the estimation of the population parameters. Usually maximum likelihood estimation is desirable as fewer observations per individual are needed to estimate the population parameters, than for a two-stage procedure. The occurring models are nonlinear in the random parameters and with this the likelihood generally cannot be described in an explicit form. Numerical procedures, such as described by Davidian and Giltinan (1995) or Pinheiro and Bates (2000), are used to approximately solve the maximum likelihood problem. Knowledge of the Fisher Information is of interest for designing the experiment. A well known approach to approximating the Fisher Information is to linearize the regression function and to assume the linearized model to be normally distributed.

In this article we outline some problems occurring, when using this approximated Fisher Information. The second section briefly describes estimation in nonlinear models and asymptotic distributions of estimators. Continuing from the ideas of nonlinear regression and of linear mixed effects models, we describe in section 3

T. Mielke , R. Schwabe

Otto-von-Guericke University Magdeburg, e-mail: tobias.mielke@ovgu.de

the first-order linearization, which is used to approximate the Fisher Information. In section 4 a problem with this approximation is illustrated through a simple example.

2 Non-linear Models

In our considered model, the observation $Y(x_i)$, taken at known experimental settings x_i in a design region X is modeled by

$$Y(x_i) = \eta(x_i, \beta) + \varepsilon_i, \text{ with } E(\varepsilon_i) = 0 \text{ and } \text{Var}(\varepsilon_i) = \sigma^2.$$

The real valued regression function η is assumed to be nonlinear in the unknown parameter vector $\beta \in \mathbb{R}^p$. To avoid difficulties we assume that η is continuous in x_i and differentiable in β . For an unknown error-distribution, a standard approach for the estimation of the vector β would be the use of least squares techniques.

Let $\xi = (x_1, \dots, x_k)$ be a concrete design and denote

$$F_\beta(\xi) := \left(\frac{\partial \eta(x_1, \beta)}{\partial \beta}, \dots, \frac{\partial \eta(x_k, \beta)}{\partial \beta} \right)^T.$$

For a vector β_0 near to the true parameter vector β , the nonlinear model can be approximated by a linear model:

$$Y \approx \eta(\xi, \beta_0) + F_{\beta_0}(\xi)(\beta - \beta_0) + \varepsilon,$$

with vectors $Y = (Y(x_1), \dots, Y(x_k))^T$, $\eta(\xi, \beta_0) = (\eta(x_1, \beta_0), \dots, \eta(x_k, \beta_0))^T$ and $\varepsilon = (\varepsilon_1, \dots, \varepsilon_k)^T$. Under the assumption of a negligible linearization error, estimation of β in the approximated model

$$Y_{\beta_0} = F_{\beta_0}(\xi)\beta + \varepsilon, \text{ where } Y_{\beta_0} := Y - \eta(\xi, \beta_0) + F_{\beta_0}(\xi)\beta_0,$$

leads to an updated guess for the true parameter vector β . For β_0 close enough to β , this procedure leads to an estimate approximately fulfilling the estimating equation

$$F_\beta(\xi)^T (y - \eta(\xi, \beta)) = 0$$

which is fulfilled by the ordinary least squares estimator $\hat{\beta}_{OLS}$. For homoscedastic errors ε_i , n replications of the design ξ and under appropriate regularity conditions, the ordinary least squares estimator $\hat{\beta}_{OLS}$ is asymptotically normally distributed:

$$\sqrt{n}(\hat{\beta}_{OLS} - \beta) \rightarrow N(0, \sigma^2(F_\beta(\xi)^T F_\beta(\xi))^{-1}) \text{ as } n \rightarrow \infty$$

and has for normally distributed homoscedastic errors ε_i the nice property of coinciding with the *ML*-estimator.

Since in ordinary least squares estimation each deviation receives equal weight, this method may be inefficient for heteroscedastic observation errors. Weighted least squares estimators for known variance structures and generalized least squares methods for unknown variance structures take the heteroscedasticities into account. With a variance matrix $V_{\beta}(\xi) = \text{diag}(\sigma^2(x_1, \beta), \dots, \sigma^2(x_k, \beta))$ depending on the experimental settings and the parameter vector β , one might transform the original model into a homoscedastic error model. Starting from a prior guess β_0 of the parameter vector β and a prior guess $V_{\beta_0}(\xi)$ of $V_{\beta}(\xi)$, new iterates for estimating β and $V_{\beta}(\xi)$ in the transformed model can be deduced. For variance functions $\sigma^2(x, \beta)$, which are known up to the vector β , independent errors and with some regularity conditions, the *GLS* estimator $\hat{\beta}_{GLS}$ is asymptotically normally distributed:

$$\sqrt{n}(\hat{\beta}_{GLS} - \beta) \rightarrow N(0, (F_{\beta}(\xi)^T V_{\beta}(\xi)^{-1} F_{\beta}(\xi))^{-1}) \text{ as } n \rightarrow \infty.$$

Contrary to the homoscedastic case with normally distributed errors, the *ML*-estimator and *GLS*-estimator generally do not coincide for heteroscedastic errors. For nonlinear regression functions with normally distributed heteroscedastic errors and a known variance matrix $V_{\beta}(\xi)$, additional information can be drawn from the variance of the observations, such that the Fisher information for β results in

$$\begin{aligned} M_{\beta}(\xi) &= E\left(\frac{\partial \ln(f_Y(y, \beta))}{\partial \beta} \frac{\partial \ln(f_Y(y, \beta))}{\partial \beta}\right) \\ &= F_{\beta}^T(\xi) V_{\beta}(\xi)^{-1} F_{\beta}(\xi) + \frac{1}{2} \tilde{S}(\xi), \end{aligned}$$

where $f_Y(y, \beta)$ is the likelihood function for the model of Y and $\tilde{S}(\xi)$ is a matrix with

$$\tilde{S}(\xi)_{ij} = Tr\left(\frac{\partial V_{\beta}(\xi)}{\partial \beta_i} V_{\beta}(\xi)^{-1} \frac{\partial V_{\beta}(\xi)}{\partial \beta_j} V_{\beta}(\xi)^{-1}\right), \text{ } i \text{ and } j = 1, \dots, p.$$

The *ML*-estimator in the normal model with heteroscedastic errors is known to be asymptotically normally distributed:

$$\sqrt{n}(\hat{\beta}_{ML} - \beta) \rightarrow N(0, M_{\beta}(\xi)^{-1}) \text{ as } n \rightarrow \infty$$

and with this it follows that for normally distributed heteroscedastic errors the *ML*-estimator is asymptotically more efficient than the *GLS*-estimator.

As Davidian and Giltinan (1995) point out, the *ML*-estimator under assumed normality loses the advantage of efficiency very quickly in the case of nonnormal data and is highly sensitive to outlying observations, while the *GLS*-estimator is more robust. Moreover, misspecified variance functions $\sigma^2(x, \beta)$ may result in biased *ML*-estimates, whereas *GLS*-estimates are not so sensitive to variance function misspecification.

Nonlinear models with heteroscedastic and homoscedastic errors have in common that the variance of the estimator $\hat{\beta}$ depends on the parameter vector β itself, so that optimal designs for these models are in general just locally optimal.

3 Mixed-Effects Models

Consider that the j -th observation of individual i with experimental settings $x_{ij} \in X$ is described by

$$Y(x_{ij}) = \eta(x_{ij}, \beta_i) + \varepsilon_{ij}.$$

The individual parameter vector β_i is assumed to be a random vector with mean β and some covariance matrix D . The observation error is assumed to have zero mean and a known constant variance σ^2 . Observation errors and individual parameter vector are considered to be independent of each other.

For linear regression functions, the assumption of normally distributed individual parameter vectors and observation errors carries forward to the marginal distribution of the observation vector. Let $Y = (Y_1^T, \dots, Y_N^T)$ describe the vector of all observations, where $Y_i = (Y(x_{i1}), \dots, Y(x_{im_i}))^T$ is the observation vector of the i -th individual with a concrete design ξ_i . As the regression function is linear in the parameter vector β , the design matrix $F(\xi_i) = F_\beta(\xi_i)$ defined in the previous section does not depend on the parameter vector. For the observation vector Y_i we obtain

$$Y_i \sim N(F(\xi_i)\beta, (F(\xi_i)DF(\xi_i)^T + \sigma^2 I_{m_i})).$$

Often primary interest lies in estimating the mean parameter vector β . With

$$\begin{aligned} F &:= (F(\xi_1)^T, \dots, F(\xi_N)^T)^T, \\ G &:= \text{diag}(F(\xi_1), \dots, F(\xi_N)), \\ V(\xi_i) &:= (F(\xi_i)DF(\xi_i)^T + \sigma^2 I_{m_i}) \text{ and} \\ V &:= \text{diag}(V(\xi_1), \dots, V(\xi_N)), \end{aligned}$$

the model of all observations is described by

$$Y = F\beta + Gb + \varepsilon, \text{ where } b = ((\beta_1 - \beta)^T, \dots, (\beta_N - \beta)^T)^T \text{ and } \varepsilon = (\varepsilon_1^T, \dots, \varepsilon_N^T)^T.$$

It readily follows that $Y \sim N(F\beta, V)$ and that the *ML*- and *GLS*-estimators, in the case of a known matrix D and σ^2 , coincide:

$$\hat{\beta}_{ML} = \hat{\beta}_{GLS} = (F^T V^{-1} F)^{-1} F^T V^{-1} Y \text{ and } \text{cov}(\hat{\beta}_{ML}) = (F^T V^{-1} F)^{-1}.$$

Note that for a variance matrix V_β depending on the parameter vector β , *ML*- and *GLS*-estimator generally do not coincide, as the Fisher Information is then of a

similar form as for nonlinear regression functions with normally distributed heteroscedastic errors (Atkinson and Cook 1995). For nonlinear mixed effects models, these results usually cannot be observed, as the observations will generally not be normally distributed. For estimating the population parameters in nonlinear mixed effects models, two-stage procedures might be helpful. In a first step individual parameter vectors could be estimated and based on these estimates the population parameter vector might be estimated. However, often reliable individual estimates cannot be obtained for the subjects. Maximum likelihood estimation on the marginal model of the observations would then be an alternative approach to obtain reliable estimates of the population parameter vector. Due to the nonlinearity of the regression function in the random parameters, a closed-form description of the likelihood of the observations y is in general nonexistent. Different numerical approaches are used to make the optimization of the likelihood a tractable problem.

If we have prior knowledge in the form of a first guess β_0 of the true population mean β , then linearization of the model around β_0 leads on the individual level by

$$\begin{aligned} Y_i &= \eta(\xi_i, \beta_i) + \varepsilon_i \\ &\approx \eta(\xi_i, \beta_0) + F_{\beta_0}(\xi_i)(\beta_i - \beta_0) + \varepsilon_i \\ &= \eta(\xi_i, \beta_0) + F_{\beta_0}(\xi_i)(\beta - \beta_0) + F_{\beta_0}(\xi_i)(\beta_i - \beta) + \varepsilon_i \end{aligned}$$

to a linear mixed effects model. With the earlier assumptions of the normal distribution for the parameter and the error vector and with the assumption that the approximating model is almost exact, one obtains

$$Y_{i,\beta_0} = F_{\beta_0}(\xi_i)\beta + F_{\beta_0}(\xi_i)(\beta_i - \beta) + \varepsilon_i, \text{ with } Y_{i,\beta_0} := Y_i - \eta(\xi_i, \beta_0) + F_{\beta_0}(\xi_i)\beta_0.$$

As a consequence it is assumed

$$Y_{i,\beta_0} \sim N(F_{\beta_0}(\xi_i)\beta, V_{\beta_0}(\xi_i)), \text{ where } V_{\beta_0}(\xi_i) := F_{\beta_0}(\xi_i)DF_{\beta_0}(\xi_i)^T + \sigma^2I_m$$

and with this

$$\hat{\beta} = \left(\sum_{i=1}^N F_{\beta_0}(\xi_i)^T V_{\beta_0}(\xi_i)^{-1} F_{\beta_0}(\xi_i) \right)^{-1} \sum_{i=1}^N F_{\beta_0}(\xi_i) V_{\beta_0}(\xi_i)^{-1} Y_{i,\beta_0}$$

is the *ML*-Estimator for β in the linearized model around β_0 and might be used as starting point for a next iteration.

A second possible approach would be the linearization of the function η around the unknown expected value of β_i as described by Davidian and Giltinan (1995):

$$\begin{aligned} Y_i &= \eta(\xi_i, \beta_i) + \varepsilon_i \\ &\approx \eta(\xi_i, \beta) + F_{\beta}(\xi_i)(\beta_i - \beta) + \varepsilon_i. \end{aligned}$$

Assuming the linearization error as negligible and with a covariance matrix $V_\beta(\xi_i) = F_\beta(\xi_i)DF_\beta(\xi_i)^T + \sigma^2 I_{m_i}$ depending on β , the marginal model results in

$$Y_{i,\beta} \sim N(\eta(\xi_i, \beta), V_\beta(\xi_i)).$$

For the estimation of β one might in this case resort to parameter estimation techniques in nonlinear heteroscedastic models with normal observation errors.

This first-order linearization around the expected value of the individual parameter vector is often used to approximate the true nonlinear mixed effects model. In their description of the first-order linearization in nonlinear mixed effects models, Davidian and Giltinan (1995) point out that the observation vector Y is taken in this method as approximately normally distributed with the moments

$$E(Y_i) \approx \eta(\xi_i, \beta) \text{ and } cov(Y_i) \approx V_\beta(\xi_i).$$

This has the drawback that if the inter-individual variation is substantial, then the linearized model may lead to biased imprecise estimation of the fixed parameters. In fact, the linearization around the population parameter vector might misleadingly suggest generating some information, as can be seen in the following example.

4 Example

Assume the observations of an experiment to follow some quadratic model. The measurements in the experimental settings $x \in X$ are considered to be exact:

$$Y_i(x) = \beta_{1,i} + \beta_{2,i}x + \beta_{3,i}x^2 =: f(x)^T \beta_i, \varepsilon_i = 0$$

and the individual parameter vector β_i is assumed to be normally distributed with mean vector β and a positive definite covariance matrix D . For simplicity assume that each individual is observed under 3 different experimental settings $x_{ij} \in X$. With these assumptions it follows that

$$Y_i \sim N(F_i \beta, V_i), \text{ where } F_i = (f(x_{i1}), f(x_{i2}), f(x_{i3}))^T \text{ and } V_i = F_i D F_i^T.$$

For the population model with N individuals follows:

$$Y \sim N(F \beta, V), \text{ where } F = (F_1^T, \dots, F_N^T)^T \text{ and } V = \text{diag}(V_1, \dots, V_N).$$

In this model the *ML*-estimator and least squares estimator coincide:

$$\hat{\beta} = (F^T V^{-1} F)^{-1} F^T V^{-1} Y \text{ with the covariance } cov(\hat{\beta}) = (F^T V^{-1} F)^{-1} = \frac{1}{N} D.$$

This is obvious, since we obtain for each individual the true parameter vector and with this:

$$\hat{\beta} \sim N(\beta, \frac{1}{N}D).$$

In a next step consider the lognormal model:

$$Y_i(x) = \eta(x, \beta_i) = \exp(\beta_{1,i} + \beta_{2,i}x + \beta_{3,i}x^2)$$

with the same assumptions as in the former example. Notice that the regression function is no longer linear in the parameters. For the *ML*-estimate under the implied lognormal model, with the design matrix F_i and variance V_i as before, we obtain

$$\begin{aligned} \hat{\beta}_{ML} &= \left(\sum_{i=1}^N F_i^T V_i^{-1} F_i \right)^{-1} \sum_{i=1}^N F_i^T V_i^{-1} \ln(Y_i) \\ &= (ND)^{-1} D \sum_{i=1}^N F_i^{-1} \ln(Y_i) = \frac{1}{N} \sum_{i=1}^N F_i^{-1} F_i \beta_i \\ &= \frac{1}{N} \sum_{i=1}^N \beta_i \sim N(\beta, \frac{1}{N}D). \end{aligned}$$

Ignoring the obvious distribution of Y_i and considering the linearization of the model around some vector β_0 , we obtain for the linearized model:

$$Y_i(x) \approx \eta(x, \beta_0) + f_{\beta_0}(x)^T (\beta_i - \beta) + f_{\beta_0}(x)^T (\beta - \beta_0)$$

with

$$f_{\beta}(x) := \left(\frac{\partial \eta(x, \beta)}{\partial \beta_1}, \frac{\partial \eta(x, \beta)}{\partial \beta_2}, \frac{\partial \eta(x, \beta)}{\partial \beta_3} \right)^T \text{ and } f_{\beta_0}(x) := f_{\beta}(x)|_{\beta=\beta_0}.$$

As the linearized design matrix $F_{i,\beta_0} = (f_{\beta_0}(x_{i1}), f_{\beta_0}(x_{i2}), f_{\beta_0}(x_{i3}))^T$ is for $\beta_0 \neq 0$ and 3 different experimental settings x_{ij} in X regular, it follows that the individual information is

$$F_{i,\beta_0}^T V_{i,\beta_0}^{-1} F_{i,\beta_0} = F_{i,\beta_0}^T (F_{i,\beta_0} D F_{i,\beta_0}^T)^{-1} F_{i,\beta_0} = D^{-1}.$$

The resulting observation vector $\tilde{Y}_{i,\beta_0} = Y_i - \eta(\xi_i, \beta_0) + F_{i,\beta_0} \beta_0$ and the assumption $\tilde{Y}_{i,\beta_0} \sim N(F_{i,\beta_0} \beta, V_{i,\beta_0})$ yield

$$\begin{aligned} \hat{\beta}_{GLS} &= \left(\sum_{i=1}^N F_{i,\beta_0}^T V_{i,\beta_0}^{-1} F_{i,\beta_0} \right)^{-1} \sum_{i=1}^N F_{i,\beta_0}^T V_{i,\beta_0}^{-1} \tilde{Y}_{i,\beta_0} \\ &= (ND)^{-1} D \sum_{i=1}^N F_{i,\beta_0}^{-1} \tilde{Y}_{i,\beta_0} \sim N(\beta, \frac{1}{N}D). \end{aligned}$$

The linearized model around the expectation of the individual effects follows as described in the previous section:

$$Y_{i,\beta} \sim N(\eta(\xi_i, \beta), V_{i,\beta}) \text{ with } V_{i,\beta} = F_{i,\beta} D F_{i,\beta}^T.$$

The asymptotic variance of the *ML*-estimator in the heteroscedastic normal model leads according to section 2 to

$$\text{cov}(\hat{\beta}) = M_{\beta}^{-1} = \left(\sum_{i=1}^N F_{i,\beta}^T V_{i,\beta}^{-1} F_{i,\beta} + \frac{1}{2} \tilde{\mathcal{S}}_{i,\beta} \right)^{-1} = (N D^{-1} + \frac{1}{2} \sum_{i=1}^N \tilde{\mathcal{S}}_{i,\beta})^{-1} < \frac{1}{N} D.$$

These differences in the information matrices are generally not negligible. Consider $D = I_3$ and $X = [-1, 1]$; the determinant of the *D*-optimal information in the linearized model is more than 20000 times the determinant of the information of the *ML*-estimator in the lognormal model.

A consequence of these results for the linear mixed model would be, that the information might be improved by simply “nonlinearizing” the model and afterwards applying a Taylor expansion conditional on the unknown population parameter vector β and assuming that this approximation is exact. This would mean that information is generated by systematically misspecifying the model.

5 Discussion

The above example is a simple illustration of one main problem of the derivation of the Fisher Information using the first-order linearization. Misspecifications of the model may lead to wrong approximations of the Fisher Information and with this to wrong optimal designs. For nonlinear mixed effects models, the vector of observations will generally not be normally distributed, to such an extent that the approximation of the Fisher Information by linearization around the population mean is not reliable.

Acknowledgements This work was supported by the BMBF grant SKAVOE 03SCPAB3

References

- Atkinson, A. C. and R. Cook (1995). D-optimum designs for heteroscedastic linear models. *Journal of the American Statistical Association* 90, 204–212.
- Davidian, M. and D. Giltinan (1995). *Nonlinear Models for Repeated Measurement Data*. London: Chapman & Hall.
- Pinheiro, J. and D. Bates (2000). *Mixed-Effects Models in S and S-Plus*. New York: Springer-Verlag.

Designs with High Breakdown Point in Nonlinear Models

Christine H. Müller and Christina Schäfer

Abstract Least trimmed squares estimators are outlier robust since they have a high breakdown point because of trimming large residuals. But the breakdown point depends also on the design. In generalized linear models and nonlinear models, the connection between breakdown point and design is given by the fullness parameter defined by Vandev and Neykov (1998). As Müller and Neykov (2003) have shown, this fullness parameter is given in generalized linear models by the largest subdesign where the parameter of interest is not identifiable. In this paper, we show that this connection does not hold for all nonlinear models. This means that the identifiability at subdesigns cannot be used for finding designs which provide high breakdown points. Instead of this, the fullness parameter itself must be determined. For some nonlinear models with two parameters, the fullness parameter is derived here. It is shown that the fullness parameter and thus a lower bound for the breakdown point depends heavily on the design and the parameter space.

1 Introduction

We assume a nonlinear model given by

$$Y_n = g(t_n, \theta) + Z_n,$$

where Y_1, \dots, Y_N are independent observations, Z_1, \dots, Z_N are independent errors, $\theta \in \Theta \subset \mathcal{R}^r$ is an unknown parameter, $t_1, \dots, t_N \in \mathcal{T} \subset \mathcal{R}^q$ are nonrandom exper-

Christine H. Müller

Fakultät Statistik, Technische Universität Dortmund, D-44221 Dortmund, Germany, e-mail: cmueller@statistik.tu-dortmund.de

Christina Schäfer

Institut für Mathematik, Carl von Ossietzky Universität Oldenburg, D-26111 Oldenburg i.O., Germany, e-mail: christina-schaefer@hotmail.de

imental conditions, and $g : \mathcal{T} \times \Theta \rightarrow \mathfrak{R}$ is a known function which is nonlinear in θ . Set $Y = (Y_1, \dots, Y_N)^\top$ with realization $y = (y_1, \dots, y_N)^\top$ and let $D = (t_1, \dots, t_N)^\top$ be the design. The density of Y_n is given by

$$f(y_n, t_n, \theta) = h((y_n - g(t_n, \theta))^2) \quad (1)$$

where h is a monotone decreasing function which is known. Then the negative log-likelihood function given by $l_n(y, D, \theta) = -\log(f(y_n, \theta))$ is a monotone increasing function of $|y_n - g(t_n, \theta)|$.

Here we will consider the breakdown point behaviour of estimators $\hat{\theta} : \mathfrak{R}^N \rightarrow \Theta$. The breakdown point of an estimator $\hat{\theta}$ is defined according to Rousseeuw and Leroy (1987) by

$$\varepsilon^*(\hat{\theta}, y) := \frac{1}{N} \min \{M;$$

there exists no compact set $\Theta_0 \subset \Theta$ with $\{\hat{\theta}(\bar{y}); \bar{y} \in \mathcal{B}_M(y)\} \subset \Theta_0\}$,

where $\mathcal{B}_M(y) := \{\bar{y} \in \mathfrak{R}^N; \text{card}\{n; y_n \neq \bar{y}_n\} \leq M\}$ is the set of contaminated samples corrupted by at most M observations. Often the condition $\Theta_0 \subset \Theta$ is replaced by

$$\Theta_0 \subset \text{int}(\Theta) \quad (2)$$

to include also the implosion point for restricted parameter spaces. To facilitate the task here, we will consider only $\Theta_0 \subset \Theta$ which means that the breakdown point is only an explosion point.

There are several approaches to high breakdown point estimators for nonlinear models. See e.g. Stromberg and Ruppert (1992), Vandev (1993), Sakata and White (1995), Vandev and Neykov (1998). High breakdown point estimators are in particular obtained by trimming large residuals. See e.g. the least trimmed squares estimators in Rousseeuw and Leroy (1987), Procházka (1988), or Jurečková and Procházka (1994). However, in generalized linear models or nonlinear models it is more appropriate to trim the smallest likelihood functions or the largest negative loglikelihood functions as Vandev (1993) and Hadi and Luceño (1997) proposed.

Trimming the least likely observations, i.e. the observations with the largest $l_n(y, \theta)$, leads to trimmed likelihoods. Maximizing the trimmed likelihood provides the trimmed likelihood estimator $TL_h(y)$ given by

$$TL_h(y) := \arg \min_{\theta} \sum_{n=1}^h l_{(n)}(y, D, \theta),$$

where $N - h$ observations are trimmed and $l_{(1)}(y, D, \theta) \leq \dots \leq l_{(N)}(y, D, \theta)$. Vandev (1993) and Vandev and Neykov (1998) studied the breakdown point behavior of trimmed likelihood estimators and showed a relation between the breakdown point and the fullness parameter d of $\{l_n(y, D, \cdot); n = 1, \dots, N\}$. They defined the d -fullness as follows.

Definition 1. A finite set $\Psi = \{\psi_n : \Theta \rightarrow \mathfrak{R}; n = 1, \dots, N\}$ of functions is called d -full if for every $\{n_1, \dots, n_d\} \subset \{1, \dots, N\}$ the function ψ given by $\psi(\theta) := \max\{\psi_{n_k}(\theta); k = 1, \dots, d\}$ is sub-compact. If $\psi(\theta) := \max\{\psi_n(\theta); n = 1, \dots, N\}$ is not subcompact, then the fullness parameter of $\Psi = \{\psi_n : \Theta \rightarrow \mathfrak{R}; n = 1, \dots, N\}$ is defined as $N + 1$.

Thereby a function $\psi : \Theta \rightarrow \mathfrak{R}$ is called sub-compact if the set $\{\theta \in \Theta; \psi(\theta) \leq C\}$ is contained in a compact set $\Theta_C \subset \Theta$ for all $C \in \mathfrak{R}$.

Again we use here for simplicity $\Theta_C \subset \Theta$ instead of $\Theta_C \subset \text{int}(\Theta)$ in the original definition of Vandev and Neykov.

The relation between breakdown point and d -fullness was derived in more detail by Müller and Neykov (2003). In particular they showed the following theorem.

Theorem 1. Assume that $\{l_n(y, D, \cdot); n = 1, \dots, N\}$ is d -full and $\lfloor \frac{N+d}{2} \rfloor \leq h \leq \lfloor \frac{N+d+1}{2} \rfloor$. Then the breakdown point of a trimmed likelihood estimator TL_h satisfies

$$\varepsilon^*(TL_h, y) \geq \frac{1}{N} \left\lfloor \frac{N-d+2}{2} \right\rfloor.$$

Theorem 1 means in particular that the fullness parameter d should be as small as possible to achieve a high breakdown point. Müller and Neykov (2003) also proved that the fullness parameter of $\{l_n(y, \cdot); n = 1, \dots, N\}$ in linear models and in many generalized linear models satisfies $d = \mathcal{N}(D) + 1$ where the so called identifiability parameter

$$\mathcal{N}(D) := \max \left\{ \sum_{n=1}^N 1_{\mathcal{D}}(t_n); \mathcal{D} \subset \{t_1, \dots, t_N\} \text{ where } \theta \text{ is not identifiable at } \mathcal{D} \right\}$$

was introduced by Müller (1995). In linear models and generalized linear models, where $g(t_n, \theta) = x(t_n)^\top \theta$ is satisfied, we have that

$$\mathcal{N}(D) = \max_{\theta \neq 0 \in \mathfrak{R}^p} \text{card} \left\{ n \in \{1, \dots, N\}; x(t_n)^\top \theta = 0 \right\},$$

so that $\mathcal{N}(D)$ provides the maximum number of explanatory variables lying in a subspace. This means in particular for linear and generalized linear models that $d = \mathcal{N}(D) + 1$ is the smallest number so that every subset of the design with this number of points provides identifiability of θ .

Although Theorem 1 holds also for nonlinear models and identifiability can be defined also for nonlinear models, there is no simple relation between identifiability and the fullness parameter d which holds for all nonlinear models. This is shown in Section 2. Section 3 and Section 4 treat the determination of the fullness parameter for two special nonlinear models with two parameters. In particular, nonlinear models with unrestricted parameter spaces are considered in Section 3, and nonlinear models with restricted parameter spaces are studied in Section 4. It is shown that

the fullness parameter and thus the lower bound for the breakdown point depends heavily on the design and the parameter space. In Section 5, extensions of the results are discussed.

2 Identifiability and d Fullness

If the density satisfies (1), then the monotony of h and the logarithm implies with the triangle inequality

$$\begin{aligned} & \max\{l_{n_k}(y, D, \theta); k = 1, \dots, d\} \\ & = \max\{-\log(h((y_{n_k} - g(t_{n_k}, \theta))^2)); k = 1, \dots, d\} \leq C \\ & \iff \max\{|g(t_{n_k}, \theta)|; k = 1, \dots, d\} \leq C_2, \end{aligned}$$

where the constants C , C_1 , and C_2 are independent of θ , but depend on y . Hence the following theorem holds.

Theorem 2.

$$\{l_n(y, D, \cdot); n = 1, \dots, N\} \text{ is } d\text{-full} \iff \{|g(t_n, \theta)|; n = 1, \dots, N\} \text{ is } d\text{-full}.$$

Identifiability in nonlinear models is defined as follows.

Definition 2. θ is identifiable at D with respect to g if and only if

$$g(t_n, \theta) = g(t_n, \tilde{\theta}) \text{ for all } n = 1, \dots, N \implies \theta = \tilde{\theta}$$

for all $\theta, \tilde{\theta} \in \Theta$.

Identifiability in nonlinear models with more than two unknown parameters is often difficult to verify. Therefore, only a simple nonlinear model is regarded, namely $g(t, \theta) = \alpha \cdot \exp(\beta t)$. Then the following result holds.

Theorem 3. If $g(t, \theta) = \alpha \cdot \exp(\beta t)$ with $\theta = (\alpha, \beta)^\top \in \Theta = [a, \infty) \times [b, \infty)$ and $0 < a < \frac{1}{\exp(bt)}$ and $D = t$ with $t > 0$ then we have

- θ is not identifiable at D ,
- $|g(t, \cdot)|$ is subcompact and thus $\{l_1(y, D, \cdot)\}$ is 1-full.

Theorem 3 means for all designs $D = (t_1, \dots, t_N) \in \mathfrak{R}^N$ with $t_n \neq 0$ for $n = 1, \dots, N$ that the fullness parameter d of $\{l_n(y, D, \cdot); n = 1, \dots, N\}$ is 1 while the identifiability parameter satisfies $\mathcal{N}(D) \geq 1$. Hence the relationship $d = \mathcal{N}(D) + 1$, which holds in linear and many generalized linear models, is not satisfied.

Proof of Theorem 3. Since $a < \frac{1}{\exp(bt)}$, there exists $\alpha, \tilde{\alpha}$ with $a < \alpha < \tilde{\alpha} < \frac{1}{\exp(bt)}$. Set $\beta = \ln\left(\frac{1}{\alpha}\right) \frac{1}{t}$ and $\tilde{\beta} = \ln\left(\frac{1}{\tilde{\alpha}}\right) \frac{1}{t}$. Then we have $\beta, \tilde{\beta} > \ln(\exp(bt)) \frac{1}{t} = b$ and

$g(t, (\alpha, \beta)) = \alpha \exp\left(\ln\left(\frac{1}{\alpha}\right) \frac{1}{t} t\right) = 1 = \tilde{\alpha} \exp\left(\ln\left(\frac{1}{\tilde{\alpha}}\right) \frac{1}{t} t\right) = g(t, (\tilde{\alpha}, \tilde{\beta}))$, so that $\theta = (\alpha, \beta)^\top$ is not identifiable at $D = t$. Furthermore, $|g(t, (\alpha, \beta))| = \alpha \cdot \exp(\beta t) \leq C$ implies $a \leq \alpha \leq \frac{C}{\exp(\beta t)} \leq \frac{C}{\exp(bt)}$ and $\exp(\beta t) \leq \frac{C}{\alpha} \leq \frac{C}{a}$ so that $\theta = (\alpha, \beta)^\top \in \left[a, \frac{C}{\exp(bt)}\right] \times \left[b, \frac{1}{t} \cdot \ln\left(\frac{C}{a}\right)\right]$ which is a compact set. Hence $\{I_1(y, D, \cdot)\}$ is 1-full. \square

A restricted parameter space like $\Theta = [a, \infty) \times [b, \infty)$ used in Theorem 3 is typical for nonlinear models based on exponential functions with high breakdown point. This is discussed in more detail in the following two sections.

3 Nonlinear Models with Unrestricted Parameter Space

In this section, nonlinear models based on the exponential function with two parameters are studied. If the support of the design consists of one negative and one positive value then no restriction of the parameter space is necessary.

Theorem 4. *If $t_1 < 0 < t_2$ and $g(t, \theta) = \alpha + \exp(\beta t)$ or $g(t, \theta) = \alpha t + \exp(\beta t)$ with $\theta = (\alpha, \beta)^\top \in \Theta = \mathfrak{R} \times \mathfrak{R}$ then*

$$\max\{|g(t_1, \cdot)|, |g(t_2, \cdot)|\}$$

is subcompact.

Proof. Consider at first $g(t, \theta) = \alpha + \exp(\beta t)$ and let be $C \geq 0$ arbitrary. Then $\max\{|g(t_1, \theta)|, |g(t_2, \theta)|\} \leq C$

implies $-C \leq \alpha + \exp(\beta t_i) \leq C$ for $i = 1, 2$ so that $-C - \exp(\beta t_i) \leq \alpha \leq C - \exp(\beta t_i) \leq C$ for $i = 1, 2$. Since $t_1 < 0 < t_2$, it holds $\beta t_1 \leq 0$ or $\beta t_2 \leq 0$ for any β so that $\alpha \geq -C - \exp(0) = -C - 1$. Hence $\alpha \in [-C - 1, C]$.

Moreover, $\exp(\beta t_i) \leq C - \alpha \leq 2C + 1$ for $i = 1, 2$ so that $\beta t_i \leq \ln(2C + 1)$ for $i = 1, 2$ which implies $\beta \in \left[\frac{\ln(2C+1)}{t_1}, \frac{\ln(2C+1)}{t_2}\right]$. Hence $\max\{|g(t_1, \cdot)|, |g(t_2, \cdot)|\}$ is subcompact for $g(t, \theta) = \alpha + \exp(\beta t)$.

Now consider $g(t, \theta) = \alpha t + \exp(\beta t)$. Again, let be $C \geq 0$ arbitrary. Then

$$\alpha t_i + e^{\beta t_i} \leq C \text{ for } i = 1, 2 \tag{3}$$

implies $\alpha t_i \leq C - e^{\beta t_i} \leq C$ for $i = 1, 2$ so that $\alpha \geq \frac{C}{t_1}$, $\alpha \leq \frac{C}{t_2}$. Hence there exists $k \geq 0$ with $-k \leq \alpha \leq k$. With this k we obtain $-kt_1 \geq \alpha t_1 \geq kt_1$, and $-kt_2 \leq \alpha t_2 \leq kt_2$ so that

$$kt_1 \leq -\alpha t_1 \leq -kt_1, \quad kt_2 \geq -\alpha t_2 \geq -kt_2. \tag{4}$$

Inequality (3) also implies $e^{\beta t_i} \leq C - \alpha t_i$ for $i = 1, 2$. With (4) we obtain $e^{\beta t_1} \leq C - \alpha t_1 \leq C - kt_1$ and $e^{\beta t_2} \leq C - \alpha t_2 \leq C + kt_2$ so that $\beta t_1 \leq \ln(C - kt_1)$, $\beta t_2 \leq \ln(C + kt_2)$ and $\beta \geq \frac{\ln(C - kt_1)}{t_1}$, $\beta \leq \frac{\ln(C + kt_2)}{t_2}$. Hence, there exists $k' \geq 0$ with $-k' \leq$

$\beta \leq k'$ so that $(\alpha, \beta)^\top \in [-k, k] \times [-k', k']$. This means that $\max\{|g(t_1, \cdot)|, |g(t_2, \cdot)|\}$ is subcompact for $g(t, \theta) = \alpha t + \exp(\beta t)$ as well. \square

If all experimental conditions are either negative or positive, then no subcompactness is possible. Without loss of generality, we can consider only the case where all experimental conditions are positive.

Theorem 5. *If $0 \leq t_1 \leq t_2 \leq \dots \leq t_N$ and $g(t, \theta) = \alpha + \exp(\beta t)$ or $g(t, \theta) = \alpha t + \exp(\beta t)$ with $\theta = (\alpha, \beta)^\top \in \Theta = \mathfrak{R} \times \mathfrak{R}$ then*

$$\max\{|g(t_n, \cdot)|; n = 1, \dots, N\}$$

is not subcompact. In particular, the fullness parameter of $\{l_n(y, D, \cdot); n = 1, \dots, N\}$ is $N + 1$.

Proof. Set $\alpha = 0$. Then $\exp(\beta t_n) \leq C$ for all $n = 1, \dots, N$ is satisfied by $\beta \leq \frac{\ln(C)}{t_N}$. Hence

$$\{0\} \times \left(-\infty, \frac{\ln(C)}{t_N}\right) \subset \{\theta; \max\{|g(t_n, \cdot)|; n = 1, \dots, N\} \leq C\}$$

so that $\max\{|g(t_n, \cdot)|; n = 1, \dots, N\}$ is not subcompact. \square

Now define for any design $D = (t_1, \dots, t_N) \in \mathfrak{R}^N$

$$N^+(D) := \text{card}\{t_n; t_n > 0\} \text{ and } N^-(D) := \text{card}\{t_n; t_n < 0\}.$$

Corollary 1. *If $g(t, \theta) = \alpha + \exp(\beta t)$ or $g(t, \theta) = \alpha t + \exp(\beta t)$ with $\theta = (\alpha, \beta)^\top \in \Theta = \mathfrak{R} \times \mathfrak{R}$ and $\min\{N^+(D), N^-(D)\} > 0$ then the fullness parameter of $\{l_n(y, D, \cdot); n = 1, \dots, N\}$ is given by*

$$\max\{N - N^+(D) + 1, N - N^-(D) + 1\}.$$

Since the fullness parameter should be as small as possible to maximize the lower bound for the breakdown point according to Theorem 1, a breakdown point maximizing design for the setup of Corollary 1 is a design with $N^+(D) = N^-(D) = \frac{N}{2}$. In this case, the lower bound for the breakdown point is approximately $\frac{1}{4}$.

However, in most applications, a nonnegative design region is assumed for a model like $g(t, \theta) = \alpha + \exp(\beta t)$ or $g(t, \theta) = \alpha t + \exp(\beta t)$. Then a fullness parameter less than $N + 1$ and thus a lower bound for the breakdown point greater than 0 is only achieved if the parameter space is restricted. This situation is studied in the next section.

4 Nonlinear Models with Restricted Parameter Space

Considering $g(t, \theta) = \alpha + \exp(\beta t)$ or $g(t, \theta) = \alpha t + \exp(\beta t)$, it is enough to restrict the parameter space of β .

Theorem 6. *If $0 \leq t_1 < t_2$ and $g(t, \theta) = \alpha + \exp(\beta t)$ or $g(t, \theta) = \alpha t + \exp(\beta t)$ with $\theta = (\alpha, \beta)^\top \in \Theta = \mathfrak{X} \times [b, \infty)$ and $b \geq 0$ then*

$$\max\{|g(t_1, \cdot)|, |g(t_2, \cdot)|\}$$

is subcompact.

Proof. Consider at first $g(t, \theta) = \alpha t + \exp(\beta t)$ and let be $C \in [0, \infty)$ arbitrary. Then

$$-C \leq \alpha t_i + e^{\beta t_i} \leq C \text{ for } i = 1, 2 \tag{5}$$

implies $\alpha t_i \leq C - e^{\beta t_i} \leq C$, $\alpha t_i \geq -C - e^{\beta t_i}$ so that $\alpha \leq \frac{C}{t_i}$ and

$$\alpha \geq \frac{1}{t_i}(-C - e^{\beta t_i}). \tag{6}$$

(6) means $-\alpha \leq \frac{1}{t_i}(C + e^{\beta t_i})$ so that with (5) we obtain $e^{\beta t_j} \leq C - \alpha t_j \leq C + \frac{1}{t_i}(C + e^{\beta t_i})t_j = C\left(1 + \frac{1}{t_i}\right) + \frac{t_j}{t_i}e^{\beta t_i}$. Dividing by $e^{\beta t_i} \geq 1$ ($t_i \geq 0, \beta \geq 0$) yields $e^{\beta(t_j - t_i)} \leq \frac{C}{e^{\beta t_i}}\left(1 + \frac{1}{t_i}\right) + \frac{t_j}{t_i} \leq C\left(1 + \frac{1}{t_i}\right) + \frac{t_j}{t_i}$ so that $\beta(t_j - t_i) \leq \ln\left(C\left(1 + \frac{1}{t_i}\right) + \frac{t_j}{t_i}\right)$. With $t_j = t_2, t_i = t_1$ we obtain $\beta \leq \frac{1}{t_2 - t_1} \ln\left(C\left(1 + \frac{1}{t_1}\right) + \frac{t_2}{t_1}\right) =: K_1$ because of $t_2 - t_1 > 0$. Inequality (6) provides then $\alpha \geq \frac{1}{t_i}(-C - e^{\beta t_i}) \geq \frac{1}{t_i}(-C - e^{K_1 t_i})$. Hence there exists $K_2 \geq 0$ such that $(\alpha, \beta)^\top \in [-K_2, K_2] \times [b, K_1]$.

The assertion for $g(t, \theta) = \alpha + \exp(\beta t)$ follows similarly. \square

Theorem 6 means that $\{l_n(y, D, \cdot); n = 1, \dots, N\}$ is 2-full if $0 \leq t_1 < t_2 < \dots < t_N$. In this case the lower bound for the breakdown point is approximately $\frac{N}{2}$ which is the maximum possible value for the lower bound. It is also obvious that repeated observation at the same experimental condition would reduce the breakdown point as Müller (1995) showed for linear models.

5 Discussion

Extensions to nonlinear models with more than two parameters are possible and will be published elsewhere. However, all these results concern only the explosion point and not the implosion point, where condition (2) would be necessary in the definition of the breakdown point. In particular for restricted parameter spaces, the implosion point is of interest, in particular when the bound is 0. However, using

condition (2) in the definition of subcompactness, as Vandev (1993) and Vandev and Neykov (1998) did, would not help. It seems that the d -fullness criterion is only useful for the explosion point.

References

- Hadi, A. S. and A. Luceño (1997). Maximum trimmed likelihood estimators: a unified approach, examples, and algorithms. *Computational Statistics and Data Analysis* 25, 251–272.
- Jurečková, J. and B. Procházka (1994). Regression quantiles and trimmed least squares estimator in nonlinear regression model. *Journal of Nonparametric Statistics* 3, 201–222.
- Müller, Ch. H. and N. Neykov (2003). Breakdown points of trimmed likelihood estimators and related estimators in generalized linear models. *Journal of Statistical Planning and Inference* 116, 503–519.
- Müller, Ch. H. (1995). Breakdown points for designed experiments. *Journal of Statistical Planning and Inference* 45, 413–427.
- Procházka, B. (1988). Regression quantiles and trimmed least squares estimator in the nonlinear regression model. *Computational Statistics and Data Analysis* 6, 385–391.
- Rousseeuw, P. J. and A. M. Leroy (1987). *Robust Regression and Outlier Detection*. New York: Wiley.
- Sakata, S. and H. White (1995). An alternative definition of finite-sample breakdown point with applications to regression model estimators. *Journal of the American Statistical Association* 90, 1099–1106.
- Stromberg, A. J. and D. Ruppert (1992). Breakdown in nonlinear regression. *Journal of the American Statistical Association* 87, 991–997.
- Vandev, D. (1993). A note on the breakdown point of the least median of squares and least trimmed squares estimators. *Statistics and Probability Letters* 16, 117–119.
- Vandev, D. and N. Neykov (1998). About regression estimators with high breakdown point. *Statistics* 32, 111–129.

A Note on the Relationship between Two Approaches to Optimal Design under Correlation

Andrej Pázman and Werner G. Müller

Abstract The note demonstrates the relationship between two recently developed methods for characterizing optimal designs, when the errors/observations in the experiments are correlated according to a given correlation structure. The understanding of this relationship can help to improve the applicability of the methods by providing new frameworks for their tuning parameters.

1 Introduction

In two previous mODa proceedings, Pázman and Müller (1998) and Fedorov and Müller (2007) put forward two different approaches to finding optimal designs, when the observations from an experiment are correlated according to a prespecified correlation function. The former approach was based on an interpretation of design measures that was very different from the classical one and it was eventually further developed into Pázman and Müller (2001) and Müller and Pázman (2003). The latter utilizes an expansion of the covariance kernel into independent components, an idea which was first formulated in Fedorov (1996).

The setup we are interested in will be the linear random field observed at n distinct points x_1, \dots, x_n yielding

$$y(x_i) = f^T(x_i)\beta + \varepsilon(x_i), \quad (1)$$

where $f^T(x)\beta$ is a linear(ized) response function at $x \in \mathfrak{X}$ containing q unknown parameters $\beta = (\beta_1, \dots, \beta_q)^T \in \mathbb{R}^q$ and \mathfrak{X} denotes the design space, corresponding to a finite set of potential trials. Note that we are using the notation of Fedorov and Müller (2007) here and in the following.

Prof. A. Pázman

Comenius University, Bratislava, Slovakia, e-mail: pazman@fmph.uniba.sk

Prof. W.G. Müller

Dept. of Applied Statistics, Johannes-Kepler-University, Linz, Austria.

Let us further assume that the random noise $\varepsilon(x)$ consists of two independent components

$$\varepsilon(x) = u(x) + e(x),$$

such that both $E[u] = E[e] = 0$, thus $E[\varepsilon] = 0$, and that $\text{Cov}[e(x), e(x')] = \sigma^2(x)\delta_{x,x'}$, and $\text{Cov}[u(x), u(x')] = k(x, x')$, where $x, x' \in \mathcal{X}$, respectively. The latter - the so-called covariance kernel - is a known function and $\delta_{x,x'}$ denotes the Kronecker-symbol. Component $u(x)$ is a random field that describes the deviation on a particular instance from the local average and can thus not be replicated. Component $e(x)$ can be either viewed as an observational error (as in Fedorov and Müller 2007), or as a regulatory device without physical meaning (as in Pázman and Müller 1998). There, it was named ‘virtual noise’ and allows for the use of design measures in this setup (cf. Müller and Pázman 2003). For that purpose its variance $\sigma^2(x)$ needs to have specific forms, some of which are presented in Pázman and Laca (2008), a particularly suitable form being

$$\sigma^2(x) = \rho \frac{(\xi(x) - 1/n)^2}{\xi(x)}, \quad (2)$$

with a small tuning parameter ρ and n the required number of points in the design. Here $\xi(x)$ denotes a design measure at x .

The resulting covariance structure is then given by $\text{Cov}[\varepsilon(x), \varepsilon(x')] = \sigma^2(x)\delta_{x,x'} + k(x, x')$, which results in a setup frequently utilized in computer simulation experiments (cf. eg. Santner, Williams, and Notz 2003) or spatial sampling (cf. eg. Müller 2007).

2 Information Matrices

In the above setting Fedorov and Müller (2007) proceed by showing that the random field (1) can be approximated by

$$y_i = \beta^T f(x_i) + \sum_{l=1}^p \gamma_l \varphi_l(x_i) + e_i,$$

with independent errors e , where the $\varphi_l(x)$ and λ_l are the eigenfunctions and eigenvalues, respectively, of the covariance kernel $k(x, x')$ up to some degree p . Here, the γ_l are random and independent, with $\text{Var}(\gamma_l) = \lambda_l$. The information matrix corresponding to best linear unbiased estimation of the trend parameters β , which are the focus of interest, is then given by

$$FWF^T - FW\Phi^T (\Phi W \Phi^T + \Lambda^{-1})^{-1} \Phi W F^T, \quad (3)$$

with $F = \{f(x_1), \dots, f(x_n)\}$, $\Phi = \{\varphi(x_1), \dots, \varphi(x_n)\}$, $\Lambda_{l,l'} = \text{Cov}[\gamma_l, \gamma_{l'}] = \lambda_l \delta_{l,l'}$, and

$$W_{i,i'} = \frac{n}{\sigma^2} \delta_{i,i'} \xi(x_i). \quad (4)$$

On the other hand, the typical form of the information matrices used in the approach of Pázman and Müller (1998) is

$$F[W^{-1} + K]^{-1}F^T,$$

with $K_{ij} = k(x_i, x_j)$, which they term ‘approximate information matrices’, however with $\sigma^2(x_i)$ from (2) instead of σ^2/ξ in formula (4). Since the eigenfunction decomposition means $k(x_i, x_j) = \sum_l \phi_l(x_i)\lambda_l\phi_l(x_j)$, we can write

$$F[W^{-1} + K]^{-1}F^T = F[W^{-1} + \Phi^T \Lambda \Phi]^{-1}F^T. \quad (5)$$

Although the Appendix of Fedorov and Müller (2007) explores various transformations, the relationship between the two types of information matrices was overlooked and we can indeed formulate the following

Theorem. The information matrices (3) and (5) are algebraically equivalent.

We first have for the inner part of (5):

$$[W^{-1} + \Phi^T \Lambda \Phi]^{-1} = [I + W\Phi^T \Lambda \Phi]^{-1}W.$$

Then we expand

$$\begin{aligned} [I + \underbrace{W\Phi^T}_A \underbrace{\Lambda\Phi}_B]^{-1} &= I - \underbrace{W\Phi^T}_A [I + \underbrace{\Lambda\Phi}_B \underbrace{W\Phi^T}_A]^{-1} \underbrace{\Lambda\Phi}_B \\ &= I - W\Phi^T [\Lambda^{-1} + \Phi W\Phi^T]^{-1} \Phi, \end{aligned}$$

which leads to

$$\begin{aligned} F[W^{-1} + \Phi^T \Lambda \Phi]^{-1}F^T &= F[I - W\Phi^T [\Lambda^{-1} + \Phi W\Phi^T]^{-1} \Phi]WF^T \\ &= FW F^T - FW\Phi^T [\Phi W\Phi^T + \Lambda^{-1}]^{-1} \Phi W F^T, \end{aligned}$$

which corresponds to (3). \square

Here we have used a standard matrix equivalence, which is for instance given in Problem 2.9 in Rao (2001).

3 Conclusions

We have thus shown in this brief note that the two seemingly different approaches share a common framework. The remaining, but still important, difference lies in the role of the matrix W reflecting the variances of the independent noise. In Fedorov and Müller (2007) this variance needs to remain nonzero (albeit potentially small) at all design points, whereas in Pázman and Müller (1998) it vanishes at the points of the optimal design (zero virtual noise) and is very large at all other points, despite

the small tuning parameter ρ . Note that both approaches allow for tuning, the former via the choice of the order of the eigenvector expansion, the latter through choice and parametrization of $\sigma^2(x)$, eg. specific choices of ρ . In particular situations the relationships established above may aid in proper tuning.

Note also that the latter method can also be expanded to the case of unknown correlation functions as proposed in Pázman and Laca (2008), which is the focus of much current research (cf. Müller and Pronzato 2009).

Acknowledgements We are grateful to two anonymous referees, whose comments helped to improve our paper, particularly our English.

References

- Fedorov, V. (1996). Design of spatial experiments: Model fitting and prediction. In S. Gosh and C. R. Rao (Eds.), *Handbook of Statistics 13*, 515–553. Amsterdam: Elsevier.
- Fedorov, V. and W. Müller (2007). Optimum design for correlated fields via covariance kernel expansions. In J. López-Fidalgo, J. M. Rodríguez-Díaz, and B. Torsney (Eds.), *mODa 8 - Advances in Model-Oriented Design*, 57–66. Heidelberg: Physica-Verlag.
- Müller, W. G. (2007). *Collecting Spatial Data: Optimum Design of Experiments for Random Fields, 3rd ed.* Heidelberg: Springer-Verlag.
- Müller, W. G. and A. Pázman (2003). Measures for designs in experiments with correlated errors. *Biometrika* 90, 423–434.
- Müller, W. G. and L. Pronzato (2009). Towards an optimal design equivalence theorem for random fields? (Submitted).
- Pázman, A. and J. Laca (2008). Modifications and extensions of the method of virtual noise in correlated experimental design. *Tatra Mountains Mathematical Publications* 39, 125–133.
- Pázman, A. and W. G. Müller (1998). A new interpretation of design measures. In A. C. Atkinson, L. Pronzato, and H. P. Wynn (Eds.), *MODA 5 - Advances in Model-Oriented Data Analysis and Experimental Design*, 239–246. Heidelberg: Physica-Verlag.
- Pázman, A. and W. G. Müller (2001). Optimal design of experiments subject to correlated errors. *Statistics & Probability Letters* 52, 29–34.
- Rao, C. R. (2001). *Linear Statistical Inference and its Applications, 2nd ed.* New York: Wiley.
- Santner, T. J., B. J. Williams, and W. Notz (2003). *The Design and Analysis of Computer Experiments.* New York: Springer-Verlag.

The Role of the Nugget Term in the Gaussian Process Method

Andrey Pepelyshev

Abstract The maximum likelihood estimate of the correlation parameter of a Gaussian process with and without a nugget term is studied in the case of the analysis of deterministic models.

1 Introduction

The Gaussian process method is an elegant way to analyze the results of experiments in many areas of science including machine learning (Rasmussen and Williams 2006), spatial statistics (Matheron 1973, Ripley 1981, Cressie 1993, Müller 2007), and the Bayesian analysis of computer experiments (Sacks, Welch, Mitchell, and Wynn 1989, Kennedy and O'Hagan 2001, Santner, Williams, and Notz 2003). Each area has its own specific ways of employing and interpreting the Gaussian processes. The purpose of this paper is not to give a full overview, which can be found in the above references, but to discuss some issues concerning the nugget term for the analysis of computer experiments.

The concept of the nugget term was first introduced in geostatistics by Matheron (1962). Roughly speaking, the variogram and covariance often show a discontinuity at the origin, termed the nugget effect. The nugget effect is considered as a random noise and may represent a measurement error or short scale variability. The nugget term is a well-explored object in spatial statistics (Pitard 1993).

Another area of the application of Gaussian processes is the Bayesian approach developed for the analysis of computer experiments. In this approach, a so-called emulator is introduced for making probabilistic judgments on the true output of the given computer model, which is called a simulator. A Gaussian process is used for a full probabilistic specification of the emulator. Thus, the emulator is utilized to measure uncertainty of different kinds (see Kennedy and O'Hagan 2001).

Dr A. Pepelyshev
University of Sheffield, Sheffield, S3 7RH, UK, e-mail: a.pepelyshev@sheffield.ac.uk

Formally, there is no nugget term in the Gaussian process method for the analysis of deterministic models, but the nugget term can be introduced artificially, for example, for the regularization of the inversion of a covariance matrix, see Neal (1997) for details. Gramacy and Lee (2009) reported on the usefulness of the nugget term in their research on supercomputer experiments.

The presence of the nugget term in the Gaussian process method is natural for the analysis of stochastic and simulation models. The nugget effect may represent a measurement error or an effect of random values used inside computer models (Kleijnen and van Beers 2005, Kleijnen 2008).

The influence of the nugget term for optimal designs of experiments for a number of cases has been studied in Zhu and Stein (2005) and in Stehlík et al. (2008).

The present paper focuses on the Gaussian process method applied for the analysis of deterministic models. It is shown that the nugget term has a great impact on the likelihood and on the estimate of the correlation parameter.

2 The Likelihood for a Gaussian Process Without the Nugget Term

In this section, it is shown that the likelihood of a Gaussian process has an unexpected behaviour in the analysis of non-stochastic models. More precisely, for a deterministic model of observations, the maximum likelihood estimate of the correlation parameter may tend to the infinity as the number of points increases. This means that a deterministic model is approximated by a Gaussian process with the correlation function $r(x) \approx 1$ for any x .

Indeed, let $y_i = \eta(x_i)$ be the output of the model $\eta(x)$ at the point $x_i \in [0, 1]$, $i = 1, \dots, n$. Note that for a deterministic model, the replication of an observation at some point gives the same output. Without loss of generality, let $x_1 < \dots < x_n$. The likelihood for a Gaussian process with constant mean β , variance σ^2 and correlation function $r(x, \tilde{x}) = e^{-|x-\tilde{x}|/\psi}$ has the form

$$p(y|\beta, \sigma, \psi) = \frac{|R|^{-1/2}}{(2\pi\sigma^2)^{n/2}} e^{-\frac{1}{2\sigma^2}(y-H\beta)^T R^{-1}(y-H\beta)}$$

where $y = (y_1, \dots, y_n)^T$ is the vector of output values, $R = (r(x_i, x_j|\psi))_{i,j=1}^n$ is the correlation matrix, $H = (h(x_1), \dots, h(x_n))$, and $h(x) \equiv 1$.

The maximum likelihood (ML) estimates of β and σ have the following explicit forms

$$\hat{\beta}_{ML} = (H^T R^{-1} H)^{-1} H R^{-1} y$$

and

$$\hat{\sigma}_{ML}^2 = \frac{1}{n} (y - H \hat{\beta}_{ML})^T R^{-1} (y - H \hat{\beta}_{ML}).$$

The ML estimate of ψ can be found only numerically in the following way

$$\hat{\psi}_{ML} = \arg \max_{\psi \in (0, \infty)} p(y | \hat{\beta}_{ML}, \hat{\sigma}_{ML}, \psi).$$

After substituting and simplifying, we obtain that the estimate $\hat{\psi}_{ML}$ maximizes

$$L(\psi) = \ln \left[|R|^{-1/2} \right] - \frac{n}{2} \ln \left[(y - H\hat{\beta}_{ML})^T R^{-1} (y - H\hat{\beta}_{ML}) \right].$$

For the exponential correlation function, the inverse of matrix R admits the explicit representation $R^{-1} = V^T V$ where the matrix V is defined by

$$V = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & 0 \\ -\frac{\mu_2}{\sqrt{1-\mu_2^2}} & \frac{1}{\sqrt{1-\mu_2^2}} & 0 & \cdots & 0 & 0 \\ 0 & -\frac{\mu_3}{\sqrt{1-\mu_3^2}} & \frac{1}{\sqrt{1-\mu_3^2}} & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & -\frac{\mu_n}{\sqrt{1-\mu_n^2}} & \frac{1}{\sqrt{1-\mu_n^2}} \end{bmatrix},$$

$\mu_i = e^{-(x_i - x_{i-1})/\psi}$. For n equidistant points $x_i = (i - 1)/(n - 1), i = 1, \dots, n$, straight-forward calculation shows that

$$yR^{-1}y = \frac{y_1^2 + y_n^2}{1 - \lambda^2} + \sum_{i=2}^{n-1} y_i^2 \frac{1 + \lambda^2}{1 - \lambda^2} - 2 \sum_{i=1}^{n-1} y_i y_{i+1} \frac{\lambda}{1 - \lambda^2}$$

where $\lambda = e^{-\frac{1}{(n-1)\psi}}$, and

$$|R|^{-1/2} = \frac{1}{(1 - \lambda^2)^{(n-1)/2}}.$$

For the model $\eta(x) = x - 1/2$, we obtain that $\hat{\beta}_{ML} = 0$ and

$$yR^{-1}y = \frac{1}{2} \frac{1}{1 - \lambda^2} + \frac{n^2 - 5n + 6}{12(n - 1)} \cdot \frac{1 + \lambda^2}{1 - \lambda^2} - \frac{n^2 - 2n - 3}{6(n - 1)} \cdot \frac{\lambda}{1 - \lambda^2}.$$

The estimate $\hat{\psi}_{ML}$ can be found explicitly in Maple and is not presented since it is a very large expression. Applying a power series expansion, we have

$$e^{-\frac{1}{(n-1)\hat{\psi}_{ML}}} = 1 - \frac{2}{n^2} - \frac{20}{3n^2} + O\left(\frac{1}{n^3}\right) \text{ and } \hat{\psi}_{ML} = \frac{n}{2} - \frac{7}{6} - \frac{7}{18n} - \frac{17}{54n^2} + O\left(\frac{1}{n^3}\right).$$

The dependence of $\hat{\psi}_{ML}$ on n is given in Figure 1 for the model $\eta(x) = x - 1/2$ in the left-hand panel and for the model $\eta(x) = \sin(2\pi x)$ in the right-hand panel. We observe that the estimate $\hat{\psi}_{ML}$ increases almost linearly as n increases for both models.

The maximum likelihood estimate of ψ for the Gaussian correlation function $r(x, \tilde{x}) = e^{-(x - \tilde{x})^2/\psi}$ is given in Figure 2. For the model $\eta(x) = x - 1/2$ we have that $\hat{\psi}_{ML} = \infty$ for any n . Note that for the model $\eta(x) = \sin(2\pi x)$, the condition




Fig. 1: The maximum likelihood estimate of ψ for the Gaussian process with the exponential correlation function and n equidistant points on the interval $[0, 1]$ for the model $\eta(x) = x - 1/2$ (left panel) and for the model $\eta(x) = \sin(2\pi x)$ (right panel) for $n = 6, \dots, 20$.

number of the correlation matrix $R(\hat{\psi}_{ML})$ is of order 10^7 , 10^{14} , 10^{22} , 10^{30} , and 10^{38} for $n = 8, 11, 14, 17,$ and 20 , respectively. These calculations were done in Maple with 45 digits precision. However, the computer representation of floating numbers typically has only 17 digits. Thus, it is impossible to find the maximum likelihood estimate for large n using the ordinary floating representation in a computer. In particular, Ababou, Bagtzoglou, and Wood (1994) have shown that the condition number grows linearly for the exponential correlation function and grows exponentially for Gaussian correlation functions.




Fig. 2: At left: the maximum likelihood estimate of ψ for the Gaussian process with the Gaussian correlation function and n equidistant points on the interval $[0, 1]$ for the model $\eta(x) = \sin(2\pi x)$ for $n = 6, \dots, 20$. At right: the likelihood function of ψ for $n = 7, 14, 20$.

In more general situations for other correlation functions and other models, the dependence of the maximum likelihood estimate and the restricted maximum likelihood estimate of ψ on n remains typically the same and can be verified numerically (Pepelyshev 2009).

Thus, roughly speaking, the estimate of the parameters of a Gaussian process is associated with the given data set and is not associated with the deterministic model. This estimation is not simple and is not well-defined. It is easy to observe that if one divides an input space into several regions, one may get quite different estimates of parameters for different regions. However, if one is looking for one Gaussian process over the full space, one has difficulty in finding the single estimate.

3 The Likelihood for a Gaussian Process With a Nugget Term

3.1 MLE for a Gaussian Process

In this section, the likelihood with the presence of the nugget term is investigated. For this case, the correlation matrix R in the formulae from Section 2 should be replaced by the correlation matrix

$$R_v = \{ (1 - v)r(x_i - x_j) + v\delta_{i,j} \}_{i,j}$$

where v is the nugget term.




Fig. 3: At left: the maximum likelihood estimate of ψ for the Gaussian process with the Gaussian correlation function and the nugget term $v = 0.02, 0.05$ for measurements of the model $\eta(x) = \sin(2\pi x)$ at n equidistant points on the interval $[0, 1]$, $n = 6, \dots, 20$. At right: the likelihood function of ψ for the nugget term $v = 0.02$ and for $n = 7, 14, 20$.

The likelihood function and the maximum likelihood estimate for fixed values of the nugget term are presented in Figure 3. One can observe that the nugget term essentially changes the maximum likelihood estimate of ψ (and also of σ). The estimate $\hat{\psi}_{ML}$ does not increase to infinity as n increases, since the Gaussian process is fitted to a band around the deterministic function. It should also be noted that the condition number of the correlation matrix R_α is of order 10^2 and is increasing very slowly as n increases. Moreover, the estimate $\hat{\psi}_{ML}$ is smaller with the presence of the nugget term that also reduces the condition number of the correlation matrix.

Ababou, Bagtzoglou, and Wood (1994) have shown that the condition number of the correlation matrix for the Gaussian process models increases to a finite limit in the presence of the nugget term.

Note one undesired effect of the nugget term. The likelihood may have a second mode for large values of the correlation parameter (see Figure 4). The second mode strongly depends on the value of the nugget term and can be considered as a false mode. For some data, the likelihood function at the second mode may have a larger value than at the first mode.




Fig. 4: The likelihood function of ψ for the Gaussian process with the Gaussian correlation function and the nugget term $v = 0, 0.01, 0.001, 0.0001$ for measurements of the model $\eta(x) = \sin(2\pi x)$ at 7 equidistant points on the interval $[0, 1]$.

Note that in the presence of the nugget term, the meta-model

$$m_v(x) = H\beta + t^T(x)R_v^{-1}(y - H\beta),$$

where $t(x) = (r(x, x_1), \dots, r(x, x_n))^T$, does not possess the interpolation property. Nevertheless, the deviations $\varepsilon_i = y_i - m_v(x_i)$ are very small. One may construct a meta-model, that interpolates the dataset $\{(x_i, \varepsilon_i)\}_{i=1}^n$, by a method given in Cressie (1993, Sect. 5.9). It is not necessary for the deviations ε_i to use the Kriging approach without the nugget term. One may use the inverse distance weighted interpolation (Cressie 1993, p. 371, Lu and Wong 2008) and define the meta-model in the following form

$$m(x) = m_v(x) + \frac{\sum_{i=1}^n \varepsilon_i \|x - x_i\|_2^{-2}}{\sum_{i=1}^n \|x - x_i\|_2^{-2}}.$$

3.2 MLE for Stationary Processes

Let us perform a small simulation study. Assume that the results of experiments satisfy

$$y(x_i) = \beta + \sigma^2 \varepsilon^{(1)}(x_i) + \tau^2 \varepsilon^{(2)}(x_i),$$

where x_1, \dots, x_n are points of measurement, $\varepsilon^{(1)}(x)$ denotes a stationary Gaussian process with correlation function $r(x) = e^{-x^2/\psi}$ and $\varepsilon^{(2)}(x)$ is white noise. Let $\mathbf{E}\varepsilon^{(j)}(x) = 0$, $\mathbf{D}\varepsilon^{(j)}(x) = 1$, with processes $\varepsilon^{(1)}(x)$ and $\varepsilon^{(2)}(x)$ independent. The values $\beta + \sigma^2 \varepsilon^{(1)}(x_i)$ may be conceived as true values of a physical process. The values $\tau^2 \varepsilon^{(2)}(x_i)$ may be interpreted as a measurement error or a rough rounding of measured values. We compute the maximum likelihood estimators for 1000 realizations obtained for $n = 8$, $x_i = (i - 1)/7$, $i = 1, \dots, 8$, $\beta = 2$, $\psi = 1.5$, $\sigma = 1$, $\tau = 0$ or $\tau = 0.01$. Results of maximum likelihood estimation for different values of the nugget term are presented in Table 1.

Table 1: The mean of maximum likelihood estimators of parameters using different values of the nugget term. Standard deviations are given in brackets.

v	$\tau = 0$			$\tau = 0.01$		
	0	0.01	0.02	0	0.01	0.02
$\hat{\beta}_{ML}$	2.03(0.68)	2.01(0.85)	2.02(0.86)	2.02(0.92)	2.04(0.85)	2.04(0.86)
$\hat{\sigma}_{ML}$	0.83(0.40)	0.29(0.17)	0.27(0.16)	0.33(0.23)	0.30(0.17)	0.28(0.16)
$\hat{\psi}_{ML}$	1.44(0.37)	0.54(0.25)	0.47(0.20)	0.14(0.06)	0.58(0.29)	0.49(0.23)

One can observe that the maximum likelihood estimators with a nonzero nugget term does not depend on small perturbations $\{\tau^2 \varepsilon^{(2)}(x_i)\}_i$ of the data $\{\beta + \sigma^2 \varepsilon^{(1)}(x_i)\}_i$. In contrast, for $v = 0$, the maximum likelihood estimators of σ and ψ are significantly changed due to adding small perturbations. In all cases, the accuracy of $\hat{\beta}_{ML}$ is approximately the same. Thus, as can be seen, the nugget term yields a regularization effect on the maximum likelihood estimators.

4 Conclusions

In the analysis of deterministic models the presence of a nugget term has a significant impact on the likelihood of a Gaussian process. The maximum likelihood estimate of the correlation parameter with a nonzero nugget term is more reliable and the condition number of the correlation matrix is moderate. Even if a deterministic model does not have any internal computational errors or other perturbations, the artificial introduction of the nugget term can be recommended.

Acknowledgements Andrey Pepelyshev thanks two referees for their valuable comments and suggestions, and acknowledges the financial support provided by the MUCM project (EPSRC grant EP/D048893/1, <http://mucm.group.shef.ac.uk>).

References

- Ababou, R., A. Bagtzoglou, and E. Wood (1994). On the condition number of covariance matrices in kriging, estimation, and simulation of random fields. *Mathematical Geology* 26, 99–133.
- Cressie, N. A. C. (1993). *Statistics for Spatial Data, revised reprint*. New York: Wiley.
- Gramacy, R. and H. Lee (2009). Adaptive design and analysis of supercomputer experiments. *Technometrics* 51, 130–145.
- Kennedy, M. C. and A. O’Hagan (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society, B* 63, 425–450.
- Kleijnen, J. P. C. (2008). *Design and Analysis of Simulation Experiments*. New York: Springer-Verlag.
- Kleijnen, J. P. C. and W. C. M. van Beers (2005). Robustness of Kriging when interpolating in random simulation with heterogeneous variances: some experiments. *European Journal of Operational Research* 165, 826–834.
- Lu, G. and D. Wong (2008). An adaptive inverse-distance weighting spatial interpolation technique. *Computers and Geosciences* 34, 1044–1055.
- Matheron, G. (1962). *Traite de Géostatistique Appliquée*. Paris: Bureau de Recherches Géologiques et Minières.
- Matheron, G. (1973). The intrinsic random functions and their applications. *Advances in Applied Probability* 5, 439–468.
- Müller, W. G. (2007). *Collecting Spatial Data, revised edn*. Heidelberg: Physica-Verlag.
- Neal, R. (1997). Monte-Carlo implementation of Gaussian process models for Bayesian classification and regression. Technical Report 9702.
- Pepelyshev, A. (2009). Fixed-domain asymptotics of maximum likelihood estimators for observations of deterministic models. Submitted to *AISTAT2010*.
- Pitard, F. (1993). *Exploration of the Nugget Effect*. Dordrecht: Kluwer.
- Rasmussen, C. E. and C. K. I. Williams (2006). *Gaussian Processes for Machine Learning*. Cambridge: MIT Press.
- Ripley, B. D. (1981). *Spatial Statistics*. New York: Wiley.
- Sacks, J., W. J. Welch, T. J. Mitchell, and H. P. Wynn (1989). Design and analysis of computer experiments (with discussion). *Statist. Sci.* 4, 409–435.
- Santner, T. J., B. J. Williams, and W. I. Notz (2003). *The Design and Analysis of Computer Experiments*. New York: Springer-Verlag.
- Stehlík, M., J. M. Rodríguez-Díaz, W. G. Müller, and J. López-Fidalgo (2008). Optimal allocation of bioassays in the case of parametrized covariance functions: an application to lung’s retention of radioactive particles. *TEST* 17, 56–68.
- Zhu, Z. and M. L. Stein (2005). Spatial sampling design for parameter estimation of the covariance function. *Journal of Statistical Planning and Inference* 134, 583–603.

A Bonferroni-Adjusted Trend Testing Method for Excess over Highest Single Agent

John J. Peterson

Abstract Combination drug therapy offers much promise for discovering pharmaceutical treatments that are efficacious and safe. A key efficacy criterion for a combination of two compounds is that the combination is superior to both of its component compounds used alone. This article proposes a simultaneous testing procedure, based upon step-down trend tests, that identifies dose combinations for pairs of compounds that produce efficacy results with excess over highest single agent (i.e. the combination is superior to both of the component compounds). This testing procedure is applied to data from experiments for a pilot high-throughput screening study for pairs of compounds evaluated at nine dose levels using 9×9 factorial experiments. This procedure is easily automated for high-throughput screening and can be computed using the SAS[®] MULTTEST procedure.

1 Testing for Excess Over Highest Single Agent (EOHSA)

1.1 Model and Testing for EOHSA

In dose-response high-throughput screening experiments, increasing doses of a compound are robotically dispensed onto laboratory plates with many very small wells. Typically, optical plate readers are used to measure (perhaps indirectly) the result of a specific biological phenomenon, such as compounds binding to cell receptor targets or cell death. Monotone dose response profiles are the expected, with noticeable increases (or decreases) for potent compounds. In this paper it is assumed that “larger is better” for dose response. For experiments where pairs of chemical agents are being screened in the search for some type of “synergy” or enhanced

John J. Peterson

GlaxoSmithKline Pharmaceuticals, 1250 So. Collegeville Rd, PO Box 5089, Collegeville, PA, 19426 e-mail: john.peterson@gsk.com

effects between the compounds, (square) factorial experiments are typically used. Robots are programmed to dispense different doses of the two compounds to produce $(K + 1) \times (K + 1)$ factorial experiments, where there are K (positive) doses of compound A and K (positive) doses of compound B . In total there are K^2 (positive) dose combinations, K doses of compound A alone, K doses of compound B alone, and possibly a control group (i.e. zero doses of compound). This experiment may be replicated r ($r \geq 1$) times. But typically in high-throughput screening r is a small number (2 or 3 say) so that resources can be used to screen as many pairs of compounds as possible. In some situations, such an experiment would be ripe for modelling by way of a sophisticated combination-drug dose-response surface model. However, for high-throughput screening, it is desirable to have a statistical inference procedure that one can easily automate. This paper presents such a procedure that appears to have better power than similar tests for “excess of highest single agent” (EOHSA) that can be easily automated.

The statistical model proposed here is a saturated ANOVA (cell means) model. The response (or a transformed version of it) is denoted by Y_{ijk} . Here, Y_{ijk} is the k^{th} replicate of the response at the i^{th} dose of compound A and the j^{th} dose of compound B . The ANOVA model is,

$$Y_{ijk} = \mu_{ij} + e_{ijk}, \quad (1)$$

where e_{ijk} are iid $N(0, \sigma^2)$ errors. If $i = 0$, then Y_{ijk} is the k^{th} replicate response associated with the j^{th} dose of compound B alone. Likewise, if $j = 0$, then Y_{ijk} is the k^{th} replicate response associated with the i^{th} dose of compound A alone. If $(i, j) = (0, 0)$, then Y_{ijk} is the k^{th} replicate associated with the control group. For a single compound combination, the most common test for EOHSA is the “min” test, denoted here as MIN. (See Laska and Meisner (1989) for details.) Define $\theta_{ij} = \min\{\delta_{ij}^A, \delta_{ij}^B\}$, where $\delta_{ij}^A = \mu_{ij} - \mu_{0j}$ and $\delta_{ij}^B = \mu_{ij} - \mu_{i0}$ for $i = 1, \dots, K$ and $j = 1, \dots, K$. Here, μ_{ij} is the true mean response for the $(i, j)^{\text{th}}$ dose group, while μ_{i0} and μ_{0j} denote the true mean responses for the single compound dose groups at doses i and j , respectively. Note that δ_{ij}^A represents a difference of means along a compound A row (or column) and at the j^{th} level of compound B , while δ_{ij}^B represents a difference of means along a compound B row (or column) at the i^{th} level of compound A . A MIN test for the (i, j) compound pairs tests $H_0^{(i,j)} : \theta_{ij} \leq 0$ vs. $H_1^{(i,j)} : \theta_{ij} > 0$. A p -value for the MIN test can be obtained by taking the maximum of the p -values for separate tests of $H_{0A}^{(i,j)} : \delta_{ij}^A \leq 0$ vs. $H_{1A}^{(i,j)} : \delta_{ij}^A > 0$ and $H_{0B}^{(i,j)} : \delta_{ij}^B \leq 0$ vs. $H_{1B}^{(i,j)} : \delta_{ij}^B > 0$ (Hung 2000).

For a factorial design, Hung et al. (1993) proposed two tests (the AVE test and the MAX test) to test for the existence of at least one $\theta_{ij} > 0$. As with the MIN test, when all of the δ_{ij}^A and δ_{ij}^B are zero (indicating no compound effects), the probability of rejecting for the AVE or MAX test becomes much less than α (Hung et al. 1993).

Hung (2000) proposes an alternative MAX test based upon using the Hochberg (1988) p -value adjustment of all of the MIN test p -values for each combination. This approach also provides an adjusted p -value for each of the compound combinations. Westfall et al. (2001) have also proposed adjusted p -value approaches for

the testing EOHSA for various clinical endpoints or fixed dose combinations. They study three types “union-intersection intersection-union” (UIIU) tests, one based upon Bonferroni -Holm (BHUIIU), one based upon Simes-Hommel (SHUIIU), and one based upon parametric resampling (RBUIIU). They find that the SHUIIU and RBUIIU tests are more powerful, with no overall domination between the two. They recommend the SHUIIU due to it being easier to compute as it is simply the Hommel (1988) step-up procedure applied to the MIN test p-values for each combination.

1.2 Approaches Based Upon Trend Tests

As a compromise approach, one possibility is to compute a (one-sided) trend test for compound *A* at each (positive) dose level of compound *B*. Likewise, one can also compute a (one-sided) trend test for compound *B* at each (positive) dose level of compound *A*. For a fixed (positive) dose level of compound *B*, trend tests for compound *A* can include all dose levels from 0 to the highest or from 0 to some intermediate dose level by appropriate choice of contrast coefficients. An analogous set of trend tests can be computed for compound *B* at fixed levels of compound *A*. If a trend is statistically significant from 0 to the i^{th} dose level of some drug, then it follows that the mean response for i^{th} dose of that drug must exceed the mean response of that drug at the zero dose.

A trend testing approach for identifying EOHSA (and other related inferences) has been proposed by Hellmich and Lehmacher (2005). However, this procedure is a closed testing procedure that is difficult to implement unless the factorial designs are rather small. As such, they suggest seeking approximate testing procedures which have good power and are easier to compute. Such a procedure is proposed in this article. Trend tests can be computed using standard ANOVA trend contrasts (e.g. arithmetic, log, or ordinal contrasts). If two trends are (simultaneously) statistically significant and intersect at a compound combination, then that combination exhibits EOHSA in a statistically significant fashion (provided that proper adjustment has been done to prevent excess Type I errors). A trend test uses all of the $(K + 1)$ levels in a one-way ANOVA model to gain power to detect differences between an (i, j) combination level and its corresponding $(i, 0)$, or $(0, j)$, single-agent dose level.

A process for testing all $K \times K$ dose combinations using trend tests is proposed as follows. Form sets of one-way ANOVA trend tests utilizing the K row trends for compound *A* (at each positive dose level of compound *B*) and the K column trends for compound *B* (at each positive dose level of compound *A*). Assemble these K row trends and K column trends to form $2K$ trends of length $(K + 1)$. These $2K$ trends are then tested in a step-down fashion. The first set of $2K$ trend tests involves dose levels from 0 to the highest level. Next, a second set of $2K$ trend tests is done which involves dose levels from 0 to the ‘next-to-highest’ dose level. This continues until the K^{th} set of $2K$ trend tests is done which involves only the two dose levels 0 and the lowest (positive) dose. In each case, however, the mean squared error for each

trend test uses all of the data (in a specific row or column of the factorial design) pooled across the $(K + 1)$ relevant dose levels.

The first set of $2K$ trend tests (using dose levels from 0 to the highest level) must be adjusted for multiple testing. This can be done, for example, using the Bonferroni method to assign a significance level of $\alpha / (2K)$ to each trend test. By applying the Tukey step-down trend testing process (Tukey et al. 1985) to each of the $2K$ trends one can obtain trend tests where the highest doses tested intersect so that eventually all of the K^2 (i, j) -dose combinations are tested. So whenever a statistically significant compound A trend is found at level i (for the j^{th} level compound B) and a statistically significant compound B trend is found at level j (for the i^{th} level of compound A) it follows that the $(i, j)^{\text{th}}$ dose combination exhibits EOHSAs in a statistically significant fashion. For each of the $2K$ trends, the Tukey step-down trend test does not require adjustment of the significance level as the step-down tests are conducted. This is because for each of the $2K$ trends, statistical significance for a specific trend stops as soon as a trend is found not to be statistically significant.

1.3 Multiplicity Adjusted p -Values

This procedure for testing for compound combinations that have EOHSAs can also be used to determine p -value measures for each combination. Let $p_{ij}(A)$ be the (raw) p -value associated with a trend test of compound A for dose levels $(0 - i)$ for the j^{th} dose level of compound B . Likewise, let $p_{ij}(B)$ be the (raw) p -value associated with a trend test of compound B for dose levels $(0 - i)$ for the j^{th} dose level of compound A . For $i = 1, \dots, K$, let

$$p_{i1}^m(A), \dots, p_{iK}^m(A), p_{i1}^m(B), \dots, p_{iK}^m(B) \quad (2)$$

denote $2K$ multiplicity-adjusted p -values adjusted with respect to the $2K$ raw p -values: $p_{i1}(A), \dots, p_{iK}(A), p_{i1}(B), \dots, p_{iK}(B)$. Next, let $p_{i1}^{sm}(A), \dots, p_{iK}^{sm}(A), p_{i1}^{sm}(B), \dots, p_{iK}^{sm}(B)$ denote the step-down adjusted p -values corresponding to the p -values in (2). Then one can take each $p_{ij}^{\max} = \max \{ p_{ij}^{sm}(A), p_{ji}^{sm}(B) \}$ as a measure of evidence against the null hypothesis, $H_0^{(i,j)} : \theta_{ij} \leq 0$. The smallest p_{ij}^{\max} can be taken as evidence of the existence of at least one compound combination with EOHSAs. The following five step procedure summarizes the multiplicity-adjusted (step-down) trend testing procedure described above.

Multiplicity-Adjusted (Step-Down) Trend Test Procedure

Step 1. Set $k = K$.

Step 2. Compute the $2K$ multiplicity adjusted (one-sided) trend test p -values:

$$p_{k1}^m(A), \dots, p_{kK}^m(A), p_{k1}^m(B), \dots, p_{kK}^m(B).$$

Step 3. If $k = 1$ then go to Step 4. Otherwise, set k equal to $(k - 1)$ and go back to Step 2.

Step 4. For $i = 1, \dots, (K - 1)$ and $j = 1, \dots, (K - 1)$ compute the step-down p -value adjustments, $p_{ij}^{sm}(A) = \max \left\{ p_{ij}^m(A), p_{i+1,j}^{sm}(A) \right\}$ and $p_{ij}^{sm}(B) = \max \left\{ p_{ij}^m(B), p_{i+1,j}^{sm}(B) \right\}$. For $i = K, j = 1, \dots, K$ set: $p_{Kj}^{sm}(A) = p_{Kj}^m(A)$ and $p_{Kj}^{sm}(B) = p_{Kj}^m(B)$.

Step 5. For each (i, j) , compute $p_{ij}^{\max} = \max \left\{ p_{ij}^{sm}(A), p_{ij}^{sm}(B) \right\}$. If $p_{ij}^{\max} < \alpha$, then label the $(i, j)^{th}$ combination as having EOHSAs in a statistically significant way.

If the Bonferroni procedure is used in Step 2, then the above five-step procedure strongly controls the Type I error rate of declaring any combination to have EOHSAs when in fact the corresponding true means do not possess EOHSAs. This is stated more formally in the theorem below.

Theorem. *Given the ANOVA model described in (1) above, suppose that the multiplicity adjustment used for the $2K$ tests in Step 2 is the Bonferroni procedure with an overall error Type I rate of α . Then the proposed procedure (Steps 1-5) strongly controls the overall error rate of falsely detecting at least one (i, j) combination as having EOHSAs to at be most α .*

Proof: Let A_{ij} be the event that $p_{ij}^{sm}(A) < \alpha$ and let B_{ij} be the event that $p_{ij}^{sm}(B) < \alpha$. It follows directly that the $(i, j)^{th}$ combination is indicated as having EOHSAs in a statistically significant way if the event $(A_{ij} \cap B_{ji})$ occurs.

Let C be the index set of all (i, j) combinations such that $C = \{(i, j) : \theta_{ij} \leq 0\}$. To establish the proof one must show that

$$\Pr \left(\bigcup_{(i,j) \in C} \{A_{ij} \cap B_{ji}\} \right) \leq \alpha \tag{3}$$

for any non-empty C . Now define

$$S_{ij} = \begin{cases} A_{ij} & \text{if } \delta_{ij}^A \leq 0 \\ B_{ji} & \text{if } \delta_{ij}^A > 0 \text{ and } \delta_{ij}^B \leq 0 \end{cases} \quad \text{for all } i = 1, \dots, K \text{ and } j = 1, \dots, K.$$

Note that the definition of S_{ij} is such that the null space corresponding to $H_0^{(i,j)}$ is divided into two mutually exclusive and exhaustive parts. Next define N as

$$N = \bigcup_{(i,j) \in C} S_{ij}.$$

It follows then that

$$\Pr \left(\bigcup_{(i,j) \in C} \{A_{ij} \cap B_{ji}\} \right) \leq \Pr(N) \tag{4}$$

because $\{A_{ij} \cap B_{ji}\} \subseteq S_{ij}$. Now, due to the Tukey step-down process, for each trend, $A_{ij} \Rightarrow A_{i'j}$ where $i' > i$ and $B_{ji} \Rightarrow B_{j'i}$ where $j' > j$. It follows then that for all (i, j) , (i', j) , (j, i) and (j', i) in C , $A_{ij} \cup A_{i'j} = A_{i'j}$ for $i' > i$ and $B_{ji} \cup B_{j'i} = B_{j'i}$ for $j' > j$. So it further follows that N can be expressed as a union of A_{ij} and B_{ji} null events where (the first indices) i and j are, respectively, as large as possible in C . But this implies that such an N can be expressed as the union of at most $2K$ events, each of which corresponds to a different null trend. Denote the corresponding index set of (i, j) pairs that include only the largest possible first indices as C_{max} . As such, C_{max} has only $2K$ (i, j) elements. It then follows that

$$\Pr(N) = \Pr\left(\bigcup_{(i,j) \in C_{max}} S_{ij}\right) \leq \sum_{(i,j) \in C_{max}} \Pr(S_{ij}). \tag{5}$$

Note that for any $(\delta_{ij}^A, \delta_{ji}^B)$ in the null space corresponding to $H_0^{(i,j)}$, $\Pr(S_{ij}) \leq \alpha/2K$. Since the Bonferroni procedures used in Step 2 are all based upon $2K$ tests, and the Bonferroni adjustment is independent of the *relative* values of the raw p -values, it follows that

$$\sum_{(i,j) \in C_{max}} \Pr(S_{ij}) \leq 2K \left(\frac{\alpha}{2K}\right) = \alpha. \tag{6}$$

Thus (3) follows directly from (4), (5) and (6).

2 An Example

This section summarizes the results of 16 experiments, conducted for a compound-pair screening pilot study. Here, the experimenters are looking for pairs of compounds exhibiting EOHSA with respect to their ability to kill cancer cells.

The Bonferroni-adjusted trend testing procedure of this article is applied along with four other procedures: AVE, SHUIIU, a version of the SHUIIU test employing trends, and a straightforward ANOVA application of PROC MULTTEST using only (multiplicity adjusted) pairwise contrasts (denoted as MAPC). The multiplicity adjusted trend testing approach proposed in this paper that is based upon the Bonferroni adjustment is denoted as the MATBON procedure. Hung’s alternative MAX test based upon the Hochberg procedure was not used because the Simes-Hommel procedure (SHUIIU) is always more powerful, and will provide a more stringent comparison with the MATBON procedure.

The MIN test concept can also be applied to trend contrast test statistics. This can be done by simply replacing the p -value in (3) with $\max\{p(A_{ij}), p(B_{ji})\}$, where $p_{ij}(A)$ and $p_{ji}(B)$ are the raw trend test p -values as defined previously. The SHUIIU procedure can then be applied to the resulting MIN test p -values to gain more power. Note, however, the MIN test p -values will not be independent as some trends will

crisscross each other in the factorial design. However, SHUIIU procedure will typically not exceed a family-wise error rate for positive dependence among the p -values (Westfall et al. 2001). Some positive dependence among the MIN test (trend) p -values is expected due to the fact that intersecting trends are positively correlated.

The factorial experimental designs were 9×9 layouts. The responses recorded for data analysis were the percent reductions as a measure of viable cancer cells relative to a positive control group (zero doses of both compounds A and B). There were two replications of the 9×9 factorial design (except for the positive control group which had 16 replications). Since the responses for data analysis were defined as “percent increase from the control mean”, the control group was not used directly as a response in the factorial design, leaving 80 treatment groups for analysis.

One further point to make is that the competing AVE, SHUIIU, SHUIUT, and MAPC tests in this paper all use the pooled mean-squared-error (MSE) from the entire ANOVA experiment, resulting in 80 df for error. However, the MATBON test as shown in this paper only used the MSE from the trend test with nine groups, resulting in 9 df for error. These two tests could be easily modified to use the pooled error, but separate MSE’s for the trend tests was considered more flexible with regard to a model to use for automated screening. Nonetheless, this test still appears to be more powerful than the other four testing procedures due to the adjustment for $2K$ multiplicities rather than K^2 multiplicities.

A brief summary of all six testing procedures for the 16 pilot screening experiments is given as follows. (Further details are available in Peterson 2006.) The proposed MATBON testing procedure found a compound combination with statistically significant EOHSAs, in 8 out of 16 cases, and it found more such combinations than any of the other procedures capable of identifying individual combinations as statistically significant. In 7 of the 16 cases it tied with the SHUIUT method for finding the largest number of compound combinations with statistically significant EOHSAs. Hung’s AVE test behaved erratically. In four cases, the AVE test achieved a very low p -value, but in other cases it did poorly relative to other tests. This is because, for some experiments, there were many compound combinations with small positive or negative $\hat{\theta}_{ij}$ values. As such, the resulting test statistic, which is the average of the $\hat{\theta}_{ij}$ values, was low. In each of the four cases where the AVE test detected the existence of EOHSAs, the MATBON test identified several statistically significant compound combinations. When average rank of the p -values (ranked across test types and averaged over the 16 experimental cases) was computed, the MATBON test had the lowest average rank.

One of the reasons for the superiority of the MATBON test over the SHUIUT, SHUIIU, and MAPC tests is that for the MATBON test the p -values are adjusted in a more parsimonious fashion. The MATBON p -values are adjusted over only $2K$ samples, and then adjusted in a prespecified, step-down fashion. However, the SHUIUT and SHUIIU tests adjust the MIN test p -values over K^2 combinations. The MAPC uses resampling-based adjustment, but the power of the trend test is not used, and adjustment is done over $2K^2$ raw p -values to make inferences about EOHSAs combinations. Further work needs to be done to see how low K can be to

remain competitive with other procedures. In the literature, K as low as five or six has been reported, although K as large as ten has also been used (Borisy et al. 2003).

References

- Borisy, A. A., P. Elliott, N. Hurst, M. Lee, J. Lehar, E. Price, G. Serbedzija, G. Zimmerman, M. Foley, B. Stockwell, and C. Keith (2003). Systematic discovery of multicomponent therapeutics. *Proceedings of the National Academy of Sciences* 100, 7977–7982.
- Hellmich, M. and W. Lehmacher (2005). Closure procedures for monotone bifactorial dose-response designs. *Biometrics* 61, 269–276.
- Hochberg, Y. (1988). A sharper Bonferroni procedure for multiple tests of significance. *Biometrika* 75, 800–802.
- Hommel, G. (1988). A stagewise rejective multiple test procedure based on a modified Bonferroni test. *Biometrika* 75, 383–386.
- Hung, H. (2000). Evaluation of a combination drug with multiple doses in unbalanced factorial design clinical trials. *Statistics in Medicine* 19, 2079–2087.
- Hung, H., G. Chi, and R. Lipicky (1993). Testing for the existence of a desirable dose combination. *Biometrics* 49, 85–94.
- Laska, E. and M. Meisner (1989). Testing whether an identified treatment is best. *Biometrics* 45, 1139–1151.
- Peterson, J. (2006). Testing for excess over highest single agent with an application to high-throughput screening of pairs of compounds. In *BDS Technical Report 2006-04*, http://biometrics.com/wp-content/uploads/2009/06/tr_2006_04.pdf.
- Tukey, J., J. Ciminera, and J. Heyse (1985). Testing the statistical certainty of a response to increasing doses of a drug. *Biometrics* 41, 295–301.
- Westfall, P. H., S.-Y. Ho, and B. Prillaman (2001). Properties of multiple intersection-union tests for multiple endpoints in combination therapy trials. *Journal of Biopharmaceutical Statistics* 11, 125–138.

Asymptotic Properties of Adaptive Penalized Optimal Designs over a Finite Space

Luc Pronzato

Abstract Adaptive optimal design with a cost constraint is considered, both for LS estimation in nonlinear regression and ML estimation in Bernoulli-type experiments, with possible applications in clinical trials. We obtain the strong consistency of the estimators for designs over a finite space, both when the cost level is fixed (and the adaptive design converges to an optimum constrained design) and when the objective is to minimize the cost. Moreover, the asymptotic normality of the estimators is obtained in the first situation, with an asymptotic covariance matrix given by the inverse of the usual information matrix, calculated as if the design were not constructed sequentially.

1 Introduction

Let \mathcal{X} , a compact subset of \mathbb{R}^d , denote the admissible domain for the experimental variables x (design points) and $\theta \in \Theta$, a compact subset of \mathbb{R}^p , denote the p -dimensional vector of parameters, all of interest, in a parametric model with independent observations $Y_i(x_i)$ conditionally on the x_i , $i = 1, 2, \dots$. The information matrix for parameters θ and design measure ξ (a probability measure on \mathcal{X}) is denoted by $\mathbf{M}(\xi, \theta) = \int_{\mathcal{X}} \mu(x, \theta) \xi(dx)$, with $\mu(x, \theta)$ the contribution of the design point x . We only consider the case of scalar observations, so that $\mu(x, \theta)$ is a rank-one matrix, which we denote $\mu(x, \theta) = \mathbf{f}_\theta(x) \mathbf{f}_\theta^\top(x)$ with $\mathbf{f}_\theta(x)$ a p -dimensional vector. We shall suppose that $\mathbf{f}_\theta(x)$ is continuously differentiable with respect to θ in the interior of Θ for all $x \in \mathcal{X}$. In a nonlinear situation, $\mathbf{M}(\xi, \theta)$ depends on θ and locally optimal design maximizes a concave function $\Psi(\cdot)$ of $\mathbf{M}(\xi, \theta)$ for some nominal value of θ . Here we shall only consider D -optimal design, i.e. $\Psi(\mathbf{M}) = \log \det(\mathbf{M})$,

Laboratoire I3S, CNRS/Université de Nice–Sophia Antipolis, France, e-mail: pronzato@i3s.unice.fr

but the extension to other global optimality criteria, such as $[\text{trace}(\mathbf{M}^{-1})]^{-1}$ for instance, could be obtained by following a similar route. A rather common approach to overcome the difficulty caused by the dependence of a locally optimal design on the unknown value of the model parameters is to design the experiment sequentially.

In fully-adaptive D -optimal design, the next design point x_{n+1} after n observations is taken as

$$x_{n+1} = \arg \max_{x \in \mathcal{X}} \text{trace}[\mu(x, \hat{\theta}^n) \mathbf{M}^{-1}(\xi_n, \hat{\theta}^n)], \quad (1)$$

where $\hat{\theta}^n \in \Theta$ is the current estimated value for θ , based on x_1, \dots, x_n and the associated observations Y_1, \dots, Y_n , and $\xi_n = (1/n) \sum_{i=1}^n \delta_{x_i}$ is the current empirical design measure. We leave aside initialisation issues and simply assume that x_1, \dots, x_p are such that $\mathbf{M}(\xi_p, \theta)$ is nonsingular for any $\theta \in \Theta$.

When $\hat{\theta}^n$ is frozen to a fixed value θ , the iteration (1) corresponds to one step of a steepest-ascent vertex-direction algorithm with step-length $1/n$ and convergence to a D -optimal design measure is proved in Wynn (1970). The fact that $\hat{\theta}^n$ is estimated in adaptive design creates dependency among observations and makes the investigation of the asymptotic behaviour of the design and estimator a much more complicated issue for which few results are available: Ford and Silvey (1980), Wu (1985) and Müller and Pötscher (1992) focus on a particular example with LS estimation; Hu (1998) is specific to Bayesian estimation by posterior mean and does not use a fully sequential design of the form (1); Lai (1994) and Chaudhuri and Mykland (1995) require the introduction of a subsequence of non-adaptive design points to ensure consistency of the estimator and Chaudhuri and Mykland (1993) require that the size of the initial non-adaptive experiment grows with the increase in size of the total experiment. Notice that the situation is different in clinical trials for comparing treatments: the designs considered are typically such that the number of allocations of each treatment goes to infinity a.s., which then yields the strong consistency of the ML estimators, see for instance the ML design in Antognini and Giovagnoli (2005). It is shown in Pronzato (2009b) that the situation becomes much simpler when \mathcal{X} is a finite set and that, under reasonable assumptions, (1) yields the a.s. convergence and asymptotic normality of the estimator $\hat{\theta}^n$. Using the results in Pronzato (2009a), we show here that similar asymptotic properties are obtained for adaptive penalized D -optimal design. We shall always assume that

$$\mathcal{X} = \{x^{(1)}, x^{(2)}, \dots, x^{(K)}\}, \quad K < \infty.$$

2 Asymptotic Properties of Estimators when \mathcal{X} is Finite

Consider a nonlinear regression model with observations

$$Y_i = Y(x_i) = \eta(x_i, \bar{\theta}) + \varepsilon_i, \quad (2)$$

where the ε_i are i.i.d. with zero mean and finite variance (which we take equal to one without any loss of generality) and $\eta(x, \theta)$ is a known function of θ and x . We suppose that $\bar{\theta}$, the unknown ‘true’ value of θ , is in the interior of Θ . We have $\mu(x, \theta) = \mathbf{f}_\theta(x) \mathbf{f}_\theta^\top(x)$ with $\mathbf{f}_\theta(x) = \partial \eta(x, \theta) / \partial \theta$. The LS estimator $\hat{\theta}_{LS}^n$ minimizes $S_n(\theta) = \sum_{k=1}^n [Y(x_k) - \eta(x_k, \theta)]^2$ and we define

$$D_n(\theta, \theta') = \sum_{i=1}^n [\eta(x_i, \theta) - \eta(x_i, \theta')]^2. \quad (3)$$

The following properties, see Pronzato (2009a), will be used in §3.2 and 3.3.

Theorem 1. *Suppose that \mathcal{X} is finite. If $D_n(\theta, \bar{\theta})$ given by (3) satisfies for all $\delta > 0$, $[\inf_{\|\theta - \bar{\theta}\| \geq \delta / \tau_n} D_n(\theta, \bar{\theta})] / (\log \log n) \xrightarrow{\text{a.s.}} \infty$ ($n \rightarrow \infty$), with $\{\tau_n\}$ a nondecreasing sequence of positive deterministic constants, then $\tau_n \|\hat{\theta}_{LS}^n - \bar{\theta}\| \xrightarrow{\text{a.s.}} 0$ as $n \rightarrow \infty$.*

Theorem 2. *Suppose that \mathcal{X} is finite and that there exist non-random symmetric positive definite $p \times p$ matrices \mathbf{C}_n such that $\mathbf{C}_n^{-1} \mathbf{M}^{1/2}(\xi_n, \bar{\theta}) \xrightarrow{p} \mathbf{I}$, with \mathbf{I} the p -dimensional identity matrix. If $c_n = \lambda_{\min}(\mathbf{C}_n)$ and $D_n(\theta, \bar{\theta})$ satisfy $n^{1/4} c_n \rightarrow \infty$ and for all $\delta > 0$, $\inf_{\|\theta - \bar{\theta}\| \geq c_n^2 \delta} D_n(\theta, \bar{\theta}) / (\log \log n) \xrightarrow{\text{a.s.}} \infty$ ($n \rightarrow \infty$), then $\hat{\theta}_{LS}^n$ satisfies $\sqrt{n} \mathbf{M}^{1/2}(\xi_n, \hat{\theta}_{LS}^n) (\hat{\theta}_{LS}^n - \bar{\theta}) \xrightarrow{d} \omega \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ as $n \rightarrow \infty$.*

Consider now the case of dose-response experiments with

$$Y \in \{0, 1\}, \quad \text{with } \Pr\{Y = 1 | x_i, \theta\} = \eta(x_i, \theta). \quad (4)$$

Suppose that the ‘true’ value of θ that generates the observations lies in the interior of Θ , that $\eta(x, \theta) \in (0, 1)$ for any $\theta \in \Theta$ and $x \in \mathcal{X}$, and that when n observations Y_1, \dots, Y_n are performed at the design points x_1, \dots, x_n , the Y_i are independent conditionally on the x_i . Also suppose that x_i is a non-random function of Y_1, \dots, Y_{i-1} , x_1, \dots, x_{i-1} for all i . Theorems 1 and 2 are then also valid for the ML estimator $\hat{\theta}_{ML}^n$ in this model, see Pronzato (2009a), and, in the rest of the paper, $\bar{\theta}^n$ will denote indifferently $\hat{\theta}_{LS}^n$ in the model (2) or $\hat{\theta}_{ML}^n$ in (4).

3 Adaptive Penalized D-optimal Design

Consider constrained locally D -optimal design that maximizes $\log \det[\mathbf{M}(\xi, \theta)]$ under a constraint $\Phi(\xi, \theta) \leq C$ on the average cost $\Phi(\xi, \theta) = \int_{\mathcal{X}} \phi(x, \theta) \xi(dx)$. We suppose that $\phi(x, \theta)$, the cost induced by a single observation at x , is a positive continuous function of θ for all $x \in \mathcal{X}$. The extension to nonlinear or multiple constraints is considered, e.g., in Cook and Fedorov (1995) and Fedorov and Hackl (1997, Chap. 4). A necessary and sufficient condition for the optimality of ξ_C^* satisfying $\Phi(\xi_C^*, \theta) \leq C$ is the existence of a Lagrange coefficient $\lambda^* = \lambda^*(\theta) \geq 0$ satisfying $\lambda^* [C - \Phi(\xi_C^*, \theta)] = 0$ and $\forall x \in \mathcal{X}$, $\text{trace}[\mu(x, \theta) \mathbf{M}^{-1}(\xi_C^*, \theta)] \leq p + \lambda^* [\phi(x, \theta) - \Phi(\xi_C^*, \theta)]$. In practice, ξ_C^* can be determined by maximizing

$$H_{\theta}(\xi, \lambda_i) = \log \det[\mathbf{M}(\xi, \theta)] - \lambda_i \Phi(\xi, \theta) \quad (5)$$

for an increasing sequence $\{\lambda_i\}$ of coefficients, starting at $\lambda_0 = 0$ and stopping at the first λ_i such that the associated optimal design ξ^* satisfies $\Phi(\xi^*, \theta) \leq C$, see, e.g., Mikulecká (1983) (notice that for C large enough the unconstrained D -optimal design is optimal for the constrained problem). The coefficient λ in $H_{\theta}(\xi, \lambda)$ can thus be considered as a penalty coefficient that penalizes costly experiments and sets the tradeoff between the maximization of $\log \det[\mathbf{M}(\xi, \theta)]$ and the minimization of $\Phi(\xi, \theta)$. One may refer to Cook and Wong (1994) for the equivalence between constrained and compound optimal designs.

In adaptive constrained D -optimal design, we take x_{n+1} that gives the steepest ascent direction for $H_{\hat{\theta}^n}(\xi_n, \lambda_n)$,

$$x_{n+1} = \arg \max_{x \in \mathcal{X}} \{ \text{trace}[\mu(x, \hat{\theta}^n) \mathbf{M}^{-1}(\xi_n, \hat{\theta}^n)] - \lambda_n \phi(x, \hat{\theta}^n) \}, \quad (6)$$

where different choices for λ_n are discussed below. Since (1) can be considered as a special case of (6), the results to be presented also cover the case of classical (unconstrained) adaptive D -optimal design (1) treated in Pronzato (2009b) (they therefore also cover the case of the adaptive penalized designs considered in Dragalin and Fedorov (2006), Dragalin, Fedorov, and Wu (2008), where the constrained problem is formulated as a standard D -optimal design problem). One may notice the similarity between (6) and the construction used in Pronzato (2000) to optimize a parametric function, the parameters of which being estimated by least-squares in a linear regression model.

Two situations will be considered concerning the choice of the sequence $\{\lambda_n\}$ in (6), respectively in §3.2 and 3.3. In the first one, the objective is to obtain an optimal design with a specified cost: we adapt λ_n to $\hat{\theta}^n$ and take $\lambda_n = \lambda^*(\hat{\theta}^n)$, the optimal Lagrange coefficient for the constrained D -optimal design problem with parameters $\hat{\theta}^n$. The second situation corresponds to the case where $\{\lambda_n\}$ forms an increasing sequence, which gives more and more importance to the constraint in the construction of the design. When $\phi(x, \theta)$ has a single minimum, by letting the Lagrange coefficient λ_n increase with n one may hope to be able to force the design to concentrate at the minimizer of ϕ associated with the true value of θ . In clinical trials, when $\phi(x, \theta)$ is related to the probability of success of treatment x , this means that we can focuss more and more on individual ethics by allocating treatments with increasing efficacy, see Pronzato (2010).

3.1 A bound on the sampling rate of nonsingular designs

The key idea used below for investigating the asymptotic properties of an estimator for a design generated by (6) is to suppose first that $\{\hat{\theta}^n\}$ is an arbitrary sequence in Θ . We shall use the following assumptions on the design space \mathcal{X} , the vectors $\mathbf{f}_{\theta}(x)$ and the Lagrange coefficients λ_n .

$$\mathbf{H}_{\mathcal{X}}\text{-}(i): \inf_{\theta \in \Theta} \lambda_{\min} \left[\sum_{i=1}^K \mathbf{f}_{\theta}(x^{(i)}) \mathbf{f}_{\theta}^{\top}(x^{(i)}) \right] > \gamma > 0;$$

$$\mathbf{H}_{\lambda}\text{-}(i): 0 \leq \lambda_n < \bar{\lambda} < \infty, \quad \forall n;$$

$$\mathbf{H}_{\lambda}\text{-}(ii): \{\lambda_n\} \text{ is a non-decreasing positive sequence and } \lim_{n \rightarrow \infty} \lambda_n = \infty.$$

Theorem 3. *Let $\{\hat{\theta}^n\}$ be an arbitrary sequence in Θ used to generate design points according to (6) in a finite design space satisfying $H_{\mathcal{X}}\text{-}(i)$, with an initialisation such that $\mathbf{M}(\xi_n, \theta)$ is non-singular for all θ in Θ and all $n \geq p$. Let $r_{n,i} = r_n(x^{(i)})$ denote the number of times $x^{(i)}$ appears in the sequence x_1, \dots, x_n , $i = 1, \dots, K$, and consider the associated order statistics $r_{n,1:K} \geq r_{n,2:K} \geq \dots \geq r_{n,K:K}$. Define*

$$q^* = \max\{j : \text{there exists } \alpha > 0 \text{ such that } \liminf_{n \rightarrow \infty} r_{n,j:K}/n > \alpha\},$$

$$q^{**} = \max\{j : \text{there exists } \alpha > 0 \text{ such that } \liminf_{n \rightarrow \infty} \lambda_n r_{n,j:K}/n > \alpha\}.$$

Then $H_{\lambda}\text{-}(i)$ implies $q^ \geq p$ and $H_{\lambda}\text{-}(ii)$ implies $q^{**} \geq p$. When the sequence $\{\hat{\theta}^n\}$ is random, the statement holds with probability one.*

The proof is similar to that of Lemma 2 in Pronzato (2009b). \mathcal{X} finite implies that q^* and $q^{**} > 1$. Supposing that $p \geq 2$, we show that assuming q^* or $q^{**} < p$ leads to a contradiction under $H_{\lambda}\text{-}(i)$ or $H_{\lambda}\text{-}(ii)$ respectively.

3.2 λ_n is bounded in (6)

When λ_n is bounded, for any sequence $\{\hat{\theta}^n\}$ used in (6), the conditions of Th. 3 ensure the existence of n_1 and $\alpha > 0$ such that $r_{n,j:K} > \alpha n$ for all $n > n_1$ and all $j = 1, \dots, p$. Under the additional assumption

$\mathbf{H}_{\mathcal{X}}\text{-}(ii)$: For all $\delta > 0$ there exists $\varepsilon(\delta) > 0$ such that for any subset $\{i_1, \dots, i_p\}$ of distinct elements of $\{1, \dots, K\}$, $\inf_{\|\theta - \bar{\theta}\| \geq \delta} \sum_{j=1}^p [\eta(x^{(i_j)}, \theta) - \eta(x^{(i_j)}, \bar{\theta})]^2 > \varepsilon(\delta)$;

we thus obtain that $D_n(\theta, \bar{\theta})$ given by (3) satisfies $\inf_{\|\theta - \bar{\theta}\| \geq \delta} D_n(\theta, \bar{\theta}) > \alpha n \varepsilon(\delta)$, $n > n_1$. Therefore, $\tilde{\theta}^n \xrightarrow{\text{a.s.}} \bar{\theta}$ ($n \rightarrow \infty$) from Th. 1, with $\tilde{\theta}^n = \hat{\theta}_{LS}^n$ in (2) or $\hat{\theta}_{ML}^n$ in (4). Since this holds for any sequence $\{\hat{\theta}^n\}$ in Θ , it is true in particular when $\tilde{\theta}^n$ is substituted for $\hat{\theta}^n$ in (6). One can take in particular $\lambda_n = \lambda^*(\tilde{\theta}^n)$, with $\lambda^*(\theta)$ the optimal Lagrange coefficient for the constrained D -optimal design problem with parameters θ . The following condition then guarantees that $H_{\lambda}\text{-}(i)$ is satisfied so that Th. 3 applies and $\tilde{\theta}^n$ is strongly consistent from Th. 1.

$\mathbf{H}_{\lambda}\text{-}(i')$: There exists $C' < C$ such that $\forall \theta \in \Theta, \exists \hat{\xi}(\theta) \in \Xi$ with $\Phi[\hat{\xi}(\theta), \theta] \leq C'$ and $\mathbf{M}[\hat{\xi}(\theta), \theta]$ has full rank.

Making the following additional assumption on \mathcal{X}

$\mathbf{H}_{\mathcal{X}}\text{-}(iii)$: $\lambda_{\min} \left[\sum_{j=1}^p \mathbf{f}_{\bar{\theta}}(x^{(i_j)}) \mathbf{f}_{\bar{\theta}}^{\top}(x^{(i_j)}) \right] \geq \bar{\gamma} > 0$ for any subset $\{i_1, \dots, i_p\}$ of distinct elements of $\{1, \dots, K\}$,

we then obtain the following concerning the convergence of $\mathbf{M}(\xi_n, \tilde{\theta}^n)$.

Theorem 4. *Suppose that the design points for $n > p$ are generated sequentially according to (6) with $\lambda_n = \lambda^*(\tilde{\theta}^n)$ and $\hat{\theta}^n = \tilde{\theta}^n$, the LS-estimator $\hat{\theta}_{LS}^n$ in (2) or the ML-estimator $\hat{\theta}_{ML}^n$ in (4). Suppose, moreover, that the first p design points are such that the information matrix is nonsingular for any $\theta \in \Theta$. Then, under $H_{\mathcal{X}}$ -(i-iii) and H_{λ} -(i') we have $\tilde{\theta}^n \xrightarrow{\text{a.s.}} \bar{\theta}$ and $\mathbf{M}(\xi_n, \tilde{\theta}^n) \xrightarrow{\text{a.s.}} \mathbf{M}[\xi^*(\bar{\theta}), \bar{\theta}]$, $n \rightarrow \infty$, with $\xi^*(\bar{\theta})$ a constrained D-optimal design for $\bar{\theta}$.*

From Th. 4 we can take $\mathbf{C}_n = \mathbf{M}^{1/2}[\xi^*(\bar{\theta}), \bar{\theta}]$ in Th. 2 and obtain the usual asymptotic normality of $\tilde{\theta}^n$ for the adaptive design (6) (although the sequential construction of the design implies that $\mathbf{M}(\xi_n, \theta)$ is not the information matrix for parameters θ).

3.3 λ_n tends to infinity in (6)

For any sequence $\{\hat{\theta}^n\}$ used in (6), the conditions of Th. 3 ensure the existence of n_1 and $\alpha > 0$ such that $r_{n,j;K} > \alpha n / \lambda_n$ for all $n > n_1$ and all $j = 1, \dots, p$. $H_{\mathcal{X}}$ -(ii) then implies that $D_n(\theta, \bar{\theta})$ given by (3) satisfies $[\inf_{\|\theta - \bar{\theta}\| \geq \delta} D_n(\theta, \bar{\theta})] / (\log \log n) > \alpha n \varepsilon(\delta) / [\lambda_n (\log \log n)]$ for $n > n_1$. Therefore, if $\lambda_n (\log \log n) / n \rightarrow 0$ as $n \rightarrow \infty$, $\tilde{\theta}^n \xrightarrow{\text{a.s.}} \bar{\theta}$ from Th. 1. Since this holds for any sequence $\{\hat{\theta}^n\}$ in Θ , it is true in particular when $\tilde{\theta}^n$ is substituted for $\hat{\theta}^n$ in (6). (One may notice that Th. 1 provides some indication about the rate of convergence of $\tilde{\theta}^n$ towards $\bar{\theta}$: for $\|\theta - \bar{\theta}\| = \delta$ small enough, $D_n(\theta, \bar{\theta}) / n \approx (\theta - \bar{\theta})^\top \mathbf{M}(\xi_n, \bar{\theta})(\theta - \bar{\theta})$, which is larger than $\alpha \bar{\gamma} \delta^2 / \lambda_n$ from $H_{\mathcal{X}}$ -(iii); therefore, $\|\tilde{\theta}^n - \bar{\theta}\| = O(\sqrt{\lambda_n (\log \log n) / \sqrt{n}})$ a.s.). The next theorem indicates that when the following is satisfied in addition to H_{λ} -(ii):

H_{λ} -(iii): λ_n / n is non-increasing and $\lambda_n (\log \log n) / n \rightarrow 0$, $n \rightarrow \infty$;

\mathbf{H}_{ϕ} : $\phi(x, \bar{\theta})$ has a unique global minimizer in \mathcal{X} : $\phi(x^{(i^*)}, \bar{\theta}) = \min_{x \in \mathcal{X}} \phi(x, \bar{\theta}) < \phi(x^{(i)}, \bar{\theta})$, $\forall i \in \{1, \dots, K\}$, $i \neq i^*$;

then $\{x_n\}$ tends to accumulate at the point of minimum cost for $\bar{\theta}$.

Theorem 5. *Suppose that the design points for $n > p$ are generated sequentially according to (6), where λ_n satisfies H_{λ} -(ii) and H_{λ} -(iii). Suppose, moreover, that the first p design points are such that the information matrix is nonsingular for any $\theta \in \Theta$. Then, under $H_{\mathcal{X}}$ -(i-iii) we have $\tilde{\theta}^n \xrightarrow{\text{a.s.}} \bar{\theta}$ and*

$$\Phi(\xi_n, \bar{\theta}) \xrightarrow{\text{a.s.}} \phi_{\bar{\theta}}^* = \min_{x \in \mathcal{X}} \phi(x, \bar{\theta}), \quad n \rightarrow \infty.$$

If, moreover, H_{ϕ} is satisfied, then $\xi_n(x^{(i)}) \xrightarrow{\text{a.s.}} 0$ for all $i \neq i^*$.

Example. Suppose that $\eta(x, \theta) = [\theta_1 / (\theta_1 - \theta_2)] [\exp(-\theta_2 x) - \exp(-\theta_1 x)]$ in the model (2) with i.i.d. errors $\mathcal{N}(0, 1)$. The objective is to maximize $\eta(x, \bar{\theta})$ for $x \in \mathcal{X}$

consisting of 1001 points regularly spaced in $[0, 10]$. We take $\phi(x, \theta) = -\eta(x, \theta)$ and $\bar{\theta} = (0.7, 0.2)^\top$, so that $\eta(x, \bar{\theta})$ reaches its maximum value in \mathcal{X} (approximately 0.606, indicated by a dashed line in Fig. 1) at $x^* = 2.51$. The design points are generated by (6) for $n \geq 2$, with $\hat{\theta}^n$ the LS estimator and $x_1 = 1.25, x_2 = 6.6$. Three sequences are considered for $\{\lambda_n\}$: $\lambda_n^{(a)} = \log^2 n, \lambda_n^{(b)} = n/(1 + \log^2 n)$ and $\lambda_n^{(c)} = n^{1.1}, n \leq 1000$ (notice that $\lambda^{(b)} < \lambda^{(a)}$ on the horizon considered). Th. 5 is satisfied for $\lambda_n^{(a)}$ and $\lambda^{(b)}$, but $\lambda_n^{(c)}$ increases too fast and does not insure convergence of ξ_n to the delta measure at x^* , see Fig. 1 for a typical realization. Of course, the behaviour is even worse for the “best intention design” (also called “forced certainty equivalence” in the control literature) $x_{k+1} = \arg \min_{x \in \mathcal{X}} \phi(x, \hat{\theta}^k)$.




Fig. 1: Evolution of $\eta(x_n, \bar{\theta})$ as a function of n for three different sequences $\{\lambda_n\}$.

Similar results are obtained when the cost $\phi(x, \bar{\theta})$ to be minimized is not directly related to $\eta(x, \bar{\theta})$. Consider, for instance, a regulation problem where the objective is to set a function $\varphi(x, \bar{\theta})$ on a given target T , so that one may take $\phi(x, \theta)$ as a measure of the distance between $\varphi(x, \theta)$ and T , e.g., $\phi(x, \theta) = [\varphi(x, \theta) - T]^2$. There, “best intention design” (the “continuous reassessment method” in dose finding), or Robbins-Monro type procedures (see, e.g., Lai and Robbins 1978) can be used when $\varphi(x, \theta) = \eta(x, \theta)$. The adaptive design (6) may be convenient in more general circumstances where the function $\varphi(x, \theta)$ to be regulated differs from $\eta(x, \theta)$.

References

Antognini, A. B. and A. Giovagnoli (2005). On the large sample optimality of sequential designs for comparing two or more treatments. *Sequential Analysis* 24, 205–217.

- Chaudhuri, P. and P. Mykland (1993). Nonlinear experiments: optimal design and inference based likelihood. *Journal of the American Statistical Association* 88(422), 538–546.
- Chaudhuri, P. and P. Mykland (1995). On efficiently designing of nonlinear experiments. *Statistica Sinica* 5, 421–440.
- Cook, R. D. and V. Fedorov (1995). Constrained optimization of experimental design (invited discussion paper). *Statistics* 26, 129–178.
- Cook, R. D. and W. Wong (1994). On the equivalence between constrained and compound optimal designs. *Journal of the American Statistical Association* 89(426), 687–692.
- Dragalin, V. and V. Fedorov (2006). Adaptive designs for dose-finding based on efficacy-toxicity response. *Journal of Statistical Planning and Inference* 136, 1800–1823.
- Dragalin, V., V. Fedorov, and Y. Wu (2008). Adaptive designs for selecting drug combinations based on efficacy-toxicity response. *Journal of Statistical Planning and Inference* 138, 352–373.
- Fedorov, V. and P. Hackl (1997). *Model-Oriented Design of Experiments*. Berlin: Springer.
- Ford, I. and S. Silvey (1980). A sequentially constructed design for estimating a nonlinear parametric function. *Biometrika* 67(2), 381–388.
- Hu, I. (1998). On sequential designs in nonlinear problems. *Biometrika* 85(2), 496–503.
- Lai, T. (1994). Asymptotic properties of nonlinear least squares estimates in stochastic regression models. *Annals of Statistics* 22(4), 1917–1930.
- Lai, T. and H. Robbins (1978). Adaptive design in regression and control. *Proc. Nat. Acad. Sci. USA* 75(2), 586–587.
- Mikulecká, J. (1983). On a hybrid experimental design. *Kybernetika* 19(1), 1–14.
- Müller, W. and B. Pötscher (1992). Batch sequential design for a nonlinear estimation problem. In V. Fedorov, W. Müller, and I. Vuchkov (Eds.), *Model-Oriented Data Analysis II, Proceedings 2nd IIASA Workshop, St Kyrik (Bulgaria), May 1990*, pp. 77–87. Heidelberg: Physica Verlag.
- Pronzato, L. (2000). Adaptive optimisation and D -optimum experimental design. *Annals of Statistics* 28(6), 1743–1761.
- Pronzato, L. (2009a). Asymptotic properties of nonlinear estimates in stochastic models with finite design space. *Statistics & Probability Letters* 79, 2307–2313. DOI: 10.1016/j.spl.2009.07.025.
- Pronzato, L. (2009b). One-step ahead adaptive D -optimal design on a finite design space is asymptotically optimal. *Metrika*. (to appear, DOI: 10.1007/s00184-008-0227-y).
- Pronzato, L. (2010). Penalized optimal designs for dose-finding. *Journal of Statistical Planning and Inference* 140, 283–296. DOI: 10.1016/j.jspi.2009.07.012.
- Wu, C. (1985). Asymptotic inference from sequential design in a nonlinear situation. *Biometrika* 72(3), 553–558.
- Wynn, H. P. (1970). The sequential generation of D -optimum experimental designs. *Annals of Math. Stat.* 41, 1655–1664.

Filling and D-optimal Designs for the Correlated Generalized Exponential Model

Juan M. Rodríguez-Díaz, Teresa Santos-Martín, Milan Stehlík and Helmut Waldl

Abstract The aim of this paper is to provide guidelines for efficient statistical estimation of the parameters of the modified Arrhenius model for chemical kinetics. We study D -optimal and filling designs for this model, assuming correlated observations and exponential covariance with or without nugget effect. We consider both equidistant and exact designs for small samples, and study the behaviour of different types of filling designs when a greater number of observations is preferred.

1 Introduction

The aim of this paper is to provide guidelines for the statistically efficient estimation of the parameters of the modified Arrhenius model. This is used, for instance, for modelling the flux of methane in the troposphere or in chemical kinetics for reactions at membranes. As the troposphere, extensively studied because of greenhouse gas emission, provides a very exotic environment for kinetics of chemical reactions normally studied on earth, one can argue that the correlation parameter and modified Arrhenius equation can play a crucial role in this modelling. The Modified Arrhenius (MA) model is

$$Y = at^{-m}e^{-\beta/T} + \varepsilon = \eta(T, m, \beta) + \varepsilon, \quad (1)$$

where $a, \beta \geq 0, m$ are constants. This is equivalent to the Generalized Exponential (GE) model through the change of variable $X = 1/T$. This later model has been studied, for the case of uncorrelated errors, by Dette and Sperlich (1994) from a Bayesian point of view and by Rodríguez-Díaz and Santos-Martín (2009) for different efficiency functions, optimality criteria and restrictions on the design space.

The model has applications in experimental sciences, especially in chemical kinetics. The influence of temperature on the rate of the process is usually given in

Juan M. Rodríguez-Díaz and Teresa Santos-Martín

Department of Statistics,

University of Salamanca, Pl. de los Caídos, Salamanca, Spain, e-mail: juanmrod@usal.es, maysam@usal.es

Milan Stehlík and Helmut Waldl

Department of Applied Statistics,

Johannes Kepler University, Freistädterstr. 315, 4040 Linz, Austria, e-mail: Milan.Stehlik@jku.at, Helmut.Waldl@jku.at

terms of the Arrhenius equation for which optimal designs are given in Rodríguez-Aragón and López-Fidalgo (2005). However, in some cases the Arrhenius model does not seem to be adequate to describe the experimental results and the MA (or GE) model appears as the right choice (Laidler 1984). Other applications of model (1) in chemistry include the transition state theory (TST) of chemical reactions (International Union of Pure and Applied Chemistry (IUPAC) 2008). The Arrhenius-like expressions appearing in TST take various forms, e.g. $k = K_B T / h \exp(-\delta G / (RT))$, where δG is the Gibbs free energy of activation, K_B is Boltzmann's constant and h is Planck's constant.

The MA model is a so-called partially nonlinear regression model (Hill 1980). This particularly means that the D -optimal designs do not depend on the value of a , which is the linear trend parameter. A design will be a collection of points of the independent variable $\{x_1, \dots, x_n\}$, where n is the size of the design (exact design). In this paper we concentrate on how to distribute temperatures optimally in order to obtain statistically efficient estimators of trend parameters m and β and correlation parameter r . After fitting the best temperature range for measurements (very often it is not the whole interval) one may be interested in particular temperatures along this interval. In this case space-filling designs would be useful, as described in Sections 4 and 5. We consider two covariance structures:

- Cov1. Exponential covariance function $C(T_1, T_2, r) = \sigma^2 \exp(-|T_1 - T_2|r)$, $r > 0$ where r is the covariance parameter and $d = |T_1 - T_2|$ is the distance between the particular temperatures.
- Cov2. Exponential with inverted arguments $C(T_1, T_2, r) = \sigma^2 \exp(-|\frac{1}{T_1} - \frac{1}{T_2}|r)$, $r > 0$. The reason for using Cov2 is to avoid ill-conditioning of the covariance matrix for some situations. Cov2 is also a natural translation of Cov1 from the GE model, using the transformation $X = 1/T$.

For more discussion on the identifiability of the covariance parameters r and σ^2 see Müller and Stehlík (2009), where the role of the nugget effect is also discussed. It is well known that the parameters r and σ^2 are not individually identifiable, which is observed also in our simulation studies. The therein used maximum likelihood estimators suffer from severe bias.

Here we centre on D -optimality, which corresponds to the maximization of the criterion function $\Phi(M) = \det M$, the determinant of the Fisher information matrix. Let θ be the vector of trend parameters and r the covariance parameter. The Fisher information matrices are $M_\theta(n)$ and $M_r(n) = \frac{1}{2} \text{tr} \left\{ C^{-1}(r) \frac{\partial C(r)}{\partial r} C^{-1}(r) \frac{\partial C(r)}{\partial r} \right\}$. So for both parameters of interest we have $M_{\theta,r}(n) = \begin{pmatrix} M_\theta(n) & 0 \\ 0 & M_r(n) \end{pmatrix}$. This method comes from the widely developed uncorrelated setup and further development is needed before it can be applied routinely in practice. Theoretical justifications for using the Fisher information for D -optimal designs in the correlated setup designs in the correlated setup can be found in Abt and Welch (1998) and Pázman (2007). Further references on the Fisher information as a design criterion in the correlated setup are in Stehlík (2007).

The paper is organized as follows. Section 2 assumes m known and illustrates the analytical peculiarities of exact designs. We must use numerical techniques to obtain the D -optimal exact designs. In section 3 we consider m to be unknown, which is the most interesting case. Usually practitioners prefer taking observations at a larger number of points (more than three) and thus, from Section 4 onwards, *filling* designs more or less covering a chosen interval with several points are studied and their behaviour compared through examples. As can be seen, the upper bounds of temperature intervals are present in the support of the optimal designs, which encompasses the natural fact that reaction kinetics for higher temperatures are speedier.

2 Assuming m Known

When m is assumed to be known the only trend parameters are a and β . The non-correlated case has been already studied in Rodríguez-Díaz and Santos-Martín (2009), where it was proved that, for approximate designs, a 2-point design is optimal. The covariance structure Cov1 corresponds to the Ornstein-Uhlenbeck process, for which $M_r(n)$ is known, (see Müller and Stehlík 2010; Zagoraiou and Baldi-Antognini 2009): $M_r(n) = \sum_{i=1}^{n-1} d_i^2 (e^{2rd_i} + 1) / (e^{2rd_i} - 1)^2$, where $d_i = T_{i+1} - T_i$ are distances between neighbouring temperatures. For $r \rightarrow 0+$ we have $M_r(n) \rightarrow \infty$ which encompasses the fact that neighbouring points are important for efficient estimation of the correlation parameter (see e.g. Kiselák and Stehlík 2008 or Zagoraiou and Baldi-Antognini 2009).

Now let us consider two-point designs with the covariance Cov2, i.e. let $n = 2$, and $\{x, x + d\}$ be the design. The following Theorem, obtained by direct algebra, gives the guidelines for computing the two-point D-optimal design.

Theorem 1. *The exact two-point D-optimal design for the GE model with covariance Cov2 is $\xi_2 = \{x, x + d\}$, where $x = (-\beta d + m + \sqrt{\beta^2 d^2 + m^2}) / (2\beta)$ and d is the zero root of*

$$\beta^2 \left(4^{2m+1} d e^{2dr} r + 16^m d r + 2^{4m+1} + e^{4dr} \left(16^m d r - 2^{4m+1} + 16^m \sqrt{\beta^2 d^2 + m^2} \right) - 16^m \sqrt{\beta^2 d^2 + m^2} \right) d^2 + m \left(m + \sqrt{\beta^2 d^2 + m^2} \right) \left\{ 4^{2m+1} d e^{2dr} r + 16^m d r + 2^{4m+1} + e^{4dr} \left(16^m d r - 2^{4m+1} \right) \right\}.$$

For ξ_2 we have $d = \frac{2x(m-x\beta)}{2x\beta-m}$; $m, \beta, r > 0$; $\frac{m}{2\beta} < x < \frac{m}{\beta}$. For all m and r the unique solution of x depending on β is bounded by the asymptotes

$$ass_l = \begin{cases} \frac{1+4m-\sqrt{1+4m}}{4\beta} & \text{for } \beta_c \leq \beta \\ \frac{m-D\beta+\sqrt{m^2+D^2\beta^2}}{2\beta} & \text{for } \beta \leq \beta_c \end{cases} \quad ass_u = \frac{m}{\beta},$$

where $\beta_c = \frac{1+4m-\sqrt{1+4m}}{2D(\sqrt{1+4m}-1)}$ and $D = \frac{1.8006}{r}$. The asymptote ass_l is already a good approximation for the exact solution. The lower bound asymptote ass_l is for small x -values very similar to the exact solution for the model with uncorrelated data analyzed in Rodríguez-Díaz and Santos-Martín (2009), where $x = (1 + 2m - \sqrt{1 + 2m}) / (2\beta)$. Furthermore we have $\lim_{x \rightarrow \infty} d = D$. On the other hand, the design given by

$(d, x) = (2\text{ass}_l(m - \text{ass}_l\beta)/(2\text{ass}_l\beta - m), \text{ass}_l)$ guarantees a minimum efficiency of 84% (Rodríguez-Díaz et al. 2009).

When $m = 0$ the first optimal design point needs to be the lower bound of the design space Rodríguez-Díaz et al. 2009. The interesting feature of this setup is that there exists a positive d^* that maximizes $\det[M(2)(\theta, r)]$, and this behaviour is different from that observed for the OU process studied in Kiseřák and Stehlík (2008), Theorem 2. Therein it is observed that, when both parameters $\{\theta, r\}$ have to be estimated, there exist D -optimal equidistant designs with finite d when the number of design points n is greater than 3. But for $n = 2$ or 3 the designs collapse, which can be avoided by a so-called nugget effect (Stehlík et al. 2008).

The case of 2-point optimal designs with covariance Cov_2 when only β and a are parameters of interest has been studied in Rodríguez-Díaz et al. (2009). The results for 3-point designs with correlation led to complicated expressions for both equidistant and general cases.

3 Case of Unknown m

Let us now study the correlated case with unknown m which should also be estimated. Thus we have a three-parameter model. The case of non equidistant designs is again troublesome, thus let us take $n = 3$ and the equidistant design $\{x, x + d, x + 2d\}$ where the parameters to be estimated are a, β, m . The determinant of the information matrix will be $D_c = D_s 2d^2 e^{4dr} (1 + e^{2dr}) / (e^{2dr} - 1)^4$, where D_s does not depend on r . After some algebra we get

$$x = \frac{-3\beta e^{4dr} d + 3\beta d - 5e^{2dr} rd - e^{4dr} rd - 2rd + 2e^{4dr} + 3e^{4dr} m - 3m - 2}{3\beta(-1 + e^{4dr})},$$

whence the optimal d can be found numerically.

Example 1. For $a = \beta = m = 1$ and $r = 0.5$, the solution is $x = 0.205$ and $d = 0.855$

For a general 3-point design we get a complicated expression. A more detailed study of design when m is unknown and further analysis is in Rodríguez-Díaz et al. (2009).

4 Parabolic Designs

Parabolic Designs are a special case of filling designs. The name *Filling Designs* refers to different types of designs that cover the design space with a specific number of points including both extremes. In previous works López-Fidalgo and Wong (2002) study several alternatives with the name *Sequence Designs* for the Michaelis-Menten model, with the distances between consecutive points increasing from the beginning of the interval. Later on, Rodríguez-Aragón and López-Fidalgo (2005) used such designs for the Arrhenius model with the distances increasing from the centre c of the interval. Very recently Rodríguez-Díaz and Santos-Martín (2009) followed this last approach for studying the Modified Arrhenius model, introduc-

ing the name *Filling Designs*. For uniform, arithmetic, geometric and linear inverse designs the samples are either uniformly distributed (directly, or inversely through the model), or more dense around the center of the interval. However, sometimes it might be preferable to concentrate samples near a point different from the centre. Following this idea, a new type of filling design is presented, using a parabolic transformation to spread a uniformly distributed set of points, in a way depending on more parameters than the extremes of the interval and the number of points. The procedure is described in Figure 1.




Fig. 1: Construction of Parabolic Designs: first select a design on the x -axis and move it to the parabola (left), then project the points in the parabola again on the x -axis through rays from the projection point Q . The spread can be modified by the choice of Q , the parabola's curvature k and a suitable movement v of the initial design. The final step is the adaptation of these last points to the desired interval.

Let us take for instance a uniformly distributed initial design in $[-1, 1]$, maybe moved by parameter v , $p_i = -1 - v + 2(i - 1)/(n - 1)$, $i = 1, \dots, n$; then using the parabola $y = kx^2$ and the projection point $Q = (0, q)$ ($q > \max\{p_1^2, p_n^2\}$), the 'projected' design will be $P_i = p_i q / (q - k p_i^2)$ $i = 1, \dots, n$. Finally let us adapt this last design to fit in the interval $[A, B]$, $\{A + l(P_i - P_1)/(P_n - P_1)\}_{i=1, \dots, n}$. The parameters k and q regulate the dispersion of the points, and $-1 \leq v \leq 1$ controls the area where the concentration of points is greatest ($-1 \rightarrow$ near A , $1 \rightarrow$ near B , $0 \rightarrow$ around $(A + B)/2$).

Example 2. Joining the measurements of Vaghjiani and Ravishankara (1991) and Gierczak et al. (1997) we get a sample of 62 temperatures in the interval $[195, 420]$. The arithmetic design chooses the best interval to be $[257.9, 420]$, with determinant 7.94892×10^{10} . The parabolic design prefers $[237, 420]$, with a determinant value of 1.71427×10^{11} ; that is more than double. But the most important fact is the value of parameter $v = 1$ showing that the best design tries to concentrate the sample points mainly at the upper extreme of the interval. The design is shown in Figure 2.




Fig. 2: Parabolic optimal design for Gierczak example in $[195, 420]$

5 Illustrative Example

METHANE: CH_4 is an important greenhouse gas, the concentration of which in the troposphere is steadily increasing. To estimate the flux of methane into the atmosphere and its atmospheric lifetime, its rate of removal needs to be accurately determined. The main loss process for atmospheric methane is the reaction with the hydroxyl radical, $OH + CH_4 \rightarrow CH_3 + H_2O$. This reaction has been extensively studied, and it can be expressed in the MA form as $k(T) = aT^{-m}exp(-\beta/T)$, with $a = 2.80 \times 10^{-14}$, $m = -0.667$, and $\beta = 1575$ (Jet Propulsion Laboratory 2006). This three-parameter fit may be preferred for lower stratosphere and upper troposphere calculations. Considering independent and normally distributed errors with mean 0 and constant variance ($\sigma^2 = 1$) and m known, the parameters to be estimated are a and β . The locally D -optimal designs for the model are two-point designs (Rodríguez-Díaz and Santos-Martín 2009), and for these values of the parameters the designs depend on the upper bound of the design interval.

Table 1: Best t_1 for different designs when fixing the upper extreme, for different r

	Uniform Design			Arithmetic Design			Geometric Design		
	$r = .001$	$r = 0.05$	$r > 0.5$	$r = .001$	$r = 0.05$	$r > 0.5$	$r = .001$	$r = 0.05$	$r > 0.5$
195-300	195	228.88	245.15	195	229.90	241.61	197.67	235.87	238.81
233-343	233	254.24	273.80	233	254.60	269.70	233	262.68	266.09
278-378	278	274.98	296.24	278	274.81	291.62	278	284.56	287.44
223-420	223	300.07	322.23	227.61	299.09	316.91	237.20	310.71	312.16
295-660	295	440.33	454.70	295	431.85	445.47	295	448.89	438.22

Let us now assume Cov1. For a general 3-point design we have fixed one design point to be the upper bound and then we have computed the exact designs. For r greater than 0.5 these designs almost coincide with the 3-point designs of the uncorrelated case. Considering four-point designs, Table 1 shows the best first design point t_1 , when fixing the upper extreme of the interval in the correlated case for different r , for the Uniform, Arithmetic and Geometric designs respectively, when only θ is the parameter of interest. It can be observed that for small r (high covariance) the best efficiency is obtained taking the filling designs in the whole initial interval. However, for independent observations it is more efficient to take the observations in a part of the initial interval. We get similar results when r needs to be estimated as well, especially when the nominal value of r is small.



Fig. 3: Different designs in the interval [208.8,420]

62 measurements have been taken at 18 different temperatures. Now the design given by the 18 different points from Example 2 will be compared with an 18 point

equidistant design on the interval [208.8, 420] (see Figure 3). Mean and variance of the empirical distributions of the ML-estimates of each parameter for both designs can be seen in Table 2, showing that the designs are quite similar. As an estimate for the D -optimality criterion, the determinant of the inverted empirical covariance matrix of the parameter estimates was computed for each design: for the Gierczak design we got $\det M_{Gierczak} = 2.3984 \cdot 10^7$, while for the equidistant design we got $\det M_{equidist} = 6.3271 \cdot 10^6$. We noticed a remarkable estimation bias - especially for \hat{r}_{ML} - which should be analyzed in more detail. An expanded version of these examples can be found in Rodríguez-Díaz et al. (2009).

Table 2: Statistics of the ML parameter estimates

	true value	Gierczak design			equidistant design		
		mean	bias	variance	mean	bias	variance
m	2.82	2.829	.009056	$8.592 \cdot 10^{-5}$	2.829	.008723	$8.094 \cdot 10^{-5}$
β	987	1009	21.77	530.9	1008	21.17	493.1
σ^2	.2365	.1904	-.04615	.01021	.1936	-.04285	.01610
r	.03643	.06560	.02917	.003477	.07058	.03415	.008176

6 Conclusions and Discussion

Probably the main lesson we can learn is that the D -optimal design is analytically peculiar and these designs can be practically obtained only by numerical computation. However, especially two point locally D -optimal designs are of interest, since they may help us to find a reasonable range for filling designs. The latter are probably the only applicable designs when seeking for a higher number of design points. It is an interesting issue that very often the best designs do not use the whole design interval, but only a part of it. This idea should be taken into account by practitioners when they design their experiments. The second important observation is the large bias of the ML estimator of r . From the theoretical point of view this is not surprising since σ^2 and r are not simultaneously identifiable. Therefore an important issue will be to develop bias reduction methods.

Acknowledgements We acknowledge helpful discussions with Licesio J. Rodríguez-Aragón from the University of Castilla La Mancha (Spain), and useful comments from Prof. Šárka Klementová from the Dept. of Mathematics, Physics and Chemistry, Univ. of South Bohemia. This research was supported by ‘Acciones Integradas 2008-09’ (Project Nr. ES 18/2008), Aktion 53p19, the Spanish Junta de Castilla y León (Project ‘SA071A09’) and the Spanish Ministry of Education and Science (Projects ‘MTM 2007-672111-C03-01’ and ‘Ingenio Mathematica (i-MATH)’ No. CSD2006-00032, Consolider - Ingenio 2010).

References

Abt, M. and W. Welch (1998). Fisher information and maximum-likelihood estimation of covariance parameters in gaussian stochastic processes. *Canadian Journal of Statistics* 26, 127–137.

- Dette, H. and S. Sperlich (1994). A note on Bayesian D-optimal designs for a generalization of the exponential growth model. *South African Statistical Journal* 28, 103–117.
- Gierczak, T., R. Talukdar, S. Herndon, G. Vaghjiani, and A. Ravishankara (1997). Rate coefficients for the reactions of hydroxyl radical with methane and deuterated methanes. *Journal of Physical Chemistry A* 101, 3125–3134.
- Hill, P. (1980). D-optimal designs for partially nonlinear regression models. *Technometrics* 22, 275–276.
- International Union of Pure and Applied Chemistry (IUPAC) (2008). Transition state theory. Technical report, <http://goldbook.iupac.org/T06470.html>.
- Jet Propulsion Laboratory (2006). Chemical kinetics and photochemical data for use in atmospheric studies, NASA panel for data evaluation. Technical report, California Institute of Technology, Pasadena, California.
- Kisěĺák, J. and M. Stehlík (2008). Equidistant D-optimal designs for parameters of Ornstein-Uhlenbeck processes. *Statistics and Probability Letters* 78, 1388–1396.
- Laidler, K. J. (1984). The development of the Arrhenius equation. *Journal of Chemical Education* 61, 494–498.
- López-Fidalgo, J. and W. Wong (2002). Design issues for the Michaelis-Menten model. *Journal of Theoretical Biology* 215, 1–11.
- Müller, W. and M. Stehlík (2009). Issues in the optimal design of computer simulation experiments. *Applied Stochastic Models in Business and Industry* 25, 163–177.
- Müller, W. and M. Stehlík (2010). Compound optimal spatial designs. *Environmetrics DOI: 10.1002/env.1009*.
- Pázman, A. (2007). Criteria for optimal design for small-sample experiments with correlated observations. *Kybernetika* 43, 453–462.
- Rodríguez-Aragón, L. and J. López-Fidalgo (2005). Optimal designs for the arrhenius equation. *Chemometrics & Intelligent Laboratory Systems* 77, 131–138.
- Rodríguez-Díaz, J. and M. Santos-Martín (2009). Study of the best designs for modifications of the arrhenius equation. *Chemometrics & Intelligent Laboratory Systems* 95, 199–208.
- Rodríguez-Díaz, J., T. Santos-Martín, M. Stehlík, and H. Waldl (2009). Filling and d-optimal designs for the correlated generalized exponential models. Technical report, IFAS - JKU University.
- Stehlík, M. (2007). D-optimal designs and equidistant designs for stationary processes. In J. López-Fidalgo and B. Rodríguez-Díaz, J.M. and Torsney (Eds.), *MODA 8 – Advances in Model-Oriented Design and Analysis*, pp. 205–212. Heidelberg: Physica-Verlag.
- Stehlík, M., J. Rodríguez-Díaz, W. Müller, and J. López-Fidalgo (2008). Optimal allocation of bioassays in the case of parametrized covariance functions: an application in lung's retention of radioactive particles. *Test* 17, 56–68.
- Vaghjiani, G. and A. Ravishankara (1991). New measurement of the rate coefficient for the reaction of OH with methane. *Nature* 350, 406–409.
- Zagoraiou, M. and A. Baldi-Antognini (2009). Optimal designs for parameter estimation of the Ornstein-Uhlenbeck process. *Applied Stochastic Models in Business and Industry* 25, 583–600.

Designs for Dose Finding Studies on Safety and Efficacy

Katrin Roth

Abstract In early phase clinical trials usually only the safety and tolerability of a new drug are investigated at first. We suggest a model with which safety, measured on a categorical scale, and efficacy, a binary response, can be evaluated simultaneously. Subsequently we derive locally optimal designs for this model. Additionally, we apply this model in a sequential approach and compare its features with approaches considering only one endpoint.

1 Introduction

Phase I dose escalation studies are part of the clinical drug development process. At that stage of the development, little knowledge about how the drug and the human body interact is available. Traditionally, the primary goal of these studies is to find the maximum tolerated dose (MTD). The MTD typically is defined as the dose that induces an intolerable toxic event (dose limiting toxicity, DLT) with a probability less than $\frac{1}{3}$. Due to safety issues, dose escalation studies are performed sequentially. An approach widely used is the 3+3 design. Following this method, subjects are assigned in cohorts of three to one out of a sequence of specified doses. The first cohort is assigned to the lowest dose. The following cohorts are assigned to the next higher, the same or the next lower dose depending on the observed number of toxicities in the previous cohort or cohorts. A maximum of two cohorts are treated on the same dose step. The algorithm stops when we would either escalate to a dose that has already been declared as too toxic (at least one third of observed DLTs) or we would take further observations on a dose step where already 6 patients have been observed. The highest dose with less than one third observed DLTs is then declared the MTD. Details can be found in Ivanova (2006). The 3+3 design is safe in the sense that only a few patients experience DLTs or are treated with doses exceeding the MTD. However, the probability of finding the actual MTD can be

Katrin Roth

Bayer Schering Pharma AG, Müllerstr. 178, 13353 Berlin, and Otto-von-Guericke-Universität Magdeburg, e-mail: katrin.roth@bayerhealthcare.com

quite low and in most cases, the MTD is underestimated. Additionally, the number of subjects needed gets large if the true MTD is much larger than the starting dose. These properties of the 3+3 design have been investigated in various simulation studies, among others in Gerke and Siedentop (2007).

Even though the primary goal of dose escalation studies is usually finding the MTD based on a binary toxicity outcome, the information on the toxicity is commonly recorded on a categorical scale (cf. National Cancer Institute, 2006). By the dichotomization valuable information is ignored. Additionally, some first information on the efficacy of the drug is often recorded, but only used and analyzed in an exploratory way. Approaches where toxicity and safety, both measured on a binary scale, are considered simultaneously are suggested among others by Dragalin and Fedorov (2006). They introduce a bivariate model and derive designs for this model.

The purpose of this work is to improve the designs for dose escalation studies in such a way that both the categorical toxicity information as well as first efficacy information is accounted for. In Section 2 we introduce a bivariate model that allows for incorporating this information. Subsequently, we derive locally optimal designs for this model. In Section 3 we apply a sequential approach to the model. We compare the features of this approach with traditional designs by conducting a simulation study. We conclude with a discussion of the results.

2 Optimal Design in a Bivariate Model

2.1 Definition of the Model

Consider a bivariate response variable $\mathbf{Y} = (T, E)^T$, with T being the toxicity endpoint and E the efficacy endpoint. Without loss of generality let the efficacy endpoint be measured on a binary scale with possible outcomes 0 (no efficacy) and 1 (efficacy), while the toxicity endpoint is observed in $K+1$ categories $j = 0, 1, \dots, K$, where the higher category indicates stronger toxicity. These categories can e.g. be defined by the Common Terminology Criteria for Adverse Events (cf. National Cancer Institute, 2006). As adequate univariate modelling, the logistic model and the proportional odds model could be used for binary and ordered categorical outcomes, respectively. The bivariate modelling should be analogous, and thus the marginal distributions of the considered endpoints should follow a logistic and a proportional odds model.

Consider a single control variable x , namely the dose.

Let $P(T(x) = y_T)$ and $P(E(x) = y_E)$ denote the probability of the outcome being $y_T \in \{0, \dots, K\}$ and $y_E \in \{0, 1\}$ given the treatment dose x .

Define $F(x) := \exp(x) / \{1 + \exp(x)\}$ and let the marginal distribution of the efficacy endpoint be given by

$$P(E(x) = 1) = \frac{\exp\left(\frac{x-\mu}{\sigma}\right)}{1 + \exp\left(\frac{x-\mu}{\sigma}\right)} = F\left(\frac{x-\mu}{\sigma}\right).$$

For notational convenience let $x_\mu := \frac{x-\mu}{\sigma}$ and thus $P(E(x) = 1) = F(x_\mu)$.

The marginal distribution of the toxicity endpoint is given by

$$P(T(x) \geq j) = \frac{\exp\left(\frac{x-\alpha_j}{\beta}\right)}{1 + \exp\left(\frac{x-\alpha_j}{\beta}\right)} = F\left(\frac{x-\alpha_j}{\beta}\right).$$

Here let $x_{\alpha_j} := (x - \alpha_j)/\beta$, $\alpha = 1, \dots, K$ and thus $P(T(x) \geq j) = F(x_{\alpha_j})$. This gives consistency with adequate univariate modelling.

The joint distribution that yields the above marginal distributions is not necessarily unique. We use a specific joint distribution for the bivariate modelling which is rather simple to construct. It is given by the following functions.

Define $G(x, y) := F(x)F(y) [1 + \tau \{1 - F(x)\} \{1 - F(y)\}]$. Then

$$P(T(x) \geq j \wedge E(x) = 1) = G(x_{\alpha_j}, x_\mu).$$

This is a bivariate distribution function from the class of Farlie-Gumbel-Morgenstern distributions (cf. Kotz et al., 2000, §44.13), which arises quite naturally from the given univariate marginal distributions. It describes the relationship between dose and efficacy, and dose and each of the toxicity categories, respectively. The relationship is such that the probability for efficacy and toxicity of a certain grade, respectively, is monotonically increasing with the dose.

The additional parameter τ is introduced to take into account a possible dependence between both endpoints. It is restricted to the interval $[-1, 1]$ to ensure a valid probability distribution. For $\tau = 0$, both endpoints are independent, for $\tau > 0$ we have a positive correlation, and for $\tau < 0$ the correlation is negative.

2.2 Optimal Designs for This Model

Following well-known optimal design theory, we first present the information matrix for this model. Thereafter we give designs for specific parameter settings.

2.2.1 Information Matrices

The information matrix of a single observation at design point x given the parameter vector $\theta = (\mu, \alpha_1, \dots, \alpha_K, \sigma, \beta, \tau)^T$ is denoted by $\mathbf{M}(x, \theta)$, and can be derived from

$$\mathbf{M}(x, \theta) = E \left(\frac{\partial l}{\partial \theta} \frac{\partial l}{\partial \theta^T} \right),$$

where l denotes the log-likelihood function of a single observation $\mathbf{y} = (j, i)$.

For notational convenience, let us define

$$H(x, y) := F(x)[1 + \tau\{1 - F(x)\}\{1 - 2F(y)\}].$$

Let $\alpha_0 = -\infty$ and $\alpha_{K+1} = \infty$. Note that $H(x_{\alpha_0}, x_\mu) = 1$ and $H(x_{\alpha_{K+1}}, x_\mu) = 0$.

Then the information matrix for this model is given by

$$\mathbf{M}(x, \boldsymbol{\theta}) = \begin{pmatrix} \mathbf{V} \cdot \mathbf{D} \cdot \mathbf{H} \\ \mathbf{t}^T \end{pmatrix} \mathbf{P} (\mathbf{H}^T \cdot \mathbf{D} \cdot \mathbf{V}^T, \mathbf{t}),$$

where:

$$\mathbf{D} = \text{diag} [F(x_\mu)\{1 - F(x_\mu)\}/\sigma, F(x_{\alpha_1})\{1 - F(x_{\alpha_1})\}/\beta, \dots, F(x_{\alpha_K})\{1 - F(x_{\alpha_K})\}/\beta] \\ \sim (K+1) \times (K+1);$$

$$\mathbf{P} = \begin{pmatrix} \text{diag} \left(\frac{1}{p_{j0}} \right)_{j=0}^K & 0 \\ 0 & \text{diag} \left(\frac{1}{p_{j1}} \right)_{j=0}^K \end{pmatrix},$$

with $p_{ji} = P(T = j \text{ and } E = i)$, and

$\mathbf{H} = (\mathbf{H}_1 \mid \mathbf{H}_2)$, where the matrix \mathbf{H} is composed of two similar looking matrices \mathbf{H}_1 and \mathbf{H}_2 , both of dimension $(K+1) \times (K+1)$. Let

$$\mathbf{H}_0 = \text{diag} H(x_\mu, x_{\alpha_j}), \mathbf{h} = \left(H(x_{\alpha_j}, x_\mu) - H(x_{\alpha_{j+1}}, x_\mu) \right)_{j=0}^K \text{ and}$$

$$\mathbf{A} = \begin{pmatrix} -1 & 1 & & 0 \\ 0 & -1 & \ddots & \\ & & \ddots & 1 & 0 \\ 0 & & & -1 & 1 \end{pmatrix}.$$

Then $\mathbf{H}_2 = \begin{pmatrix} \mathbf{h}^T \\ \mathbf{H}_0 \mathbf{A} \end{pmatrix}$. \mathbf{H}_1 has the same structure and is given by $\mathbf{H}_1 = \begin{pmatrix} \mathbf{0} \\ \mathbf{A} \end{pmatrix} - \mathbf{H}_2$.

Additionally,

$$\mathbf{V} = \begin{pmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \dots & 1 \\ x_\mu & 0 & \dots & 0 \\ 0 & x_{\alpha_1} & \dots & x_{\alpha_K} \end{pmatrix} \sim (K+3) \times (K+1),$$

and $\mathbf{t} = (\mathbf{t}_1^T, -\mathbf{t}_1^T)$ with




Fig. 1: D-optimal design for the bivariate model with $\alpha_1 = 0, \alpha_3 = 2\alpha_2, \beta = 1, \sigma = 1$ and $\tau = 0.8$; left: optimal design points, right: optimal weights.

$$t_1 = \left(F(x_\mu)\{1 - F(x_\mu)\} [F(x_{\alpha_{j+1}})\{1 - F(x_{\alpha_{j+1}})\} - F(x_{\alpha_j})\{1 - F(x_{\alpha_j})\}] \right)_{j=0}^K.$$

M is of rank $K + 1$ and therefore not of full rank. This implies that e.g. a D-optimal design has to comprise at least two distinct designs points, since a one-point-design would lead to a singular information matrix.

Details of the derivation of the information matrix and its rank can be found in Roth, 2009b.

2.2.2 Locally D-optimal Designs

In this section we present locally D-optimal designs for specific parameter settings. These designs are derived numerically using the information matrix given above, and thus are only approximations to the optimal designs, but with high efficiency.

Consider the following setting: there is one binary endpoint and one categorical endpoint with 4 categories (i.e. $K = 3$). We consider a standardized model where, without loss of generality, $\mu = 0$ and $\sigma = 1$ (cf. Ford et al., 1992) and take $\beta = 1$ and $\tau = 0.8$. The value for α_1 is fixed to 0, α_2 is varied from α_1 to 10 and α_3 is given by $2\alpha_2$ to get equidistant categories. The D-optimal designs for these parameter constellations are given in Figure 1.

The number of design points varies from 2 to 4. It increases as the differences between the parameters α_j increase. It stands out that two of the four design points coincide with α_2 and α_3 for large values of these parameters, while the other design points spread out around μ and α_1 . The weights corresponding to the design points equal to α_2 and α_3 converge to equal values of approximately 0.15, while the other weights differ and are approximately 0.34 and 0.37.

For other parameter settings, like different number of categories, different ratios of σ and β and values of α_1 different from zero, a similar structure is seen. The number of design points always increases with the difference between the values of the α_j s, but never gets larger than the number of location and dispersion parameters

in the model. It is noteworthy that the number of design points is less than the total number of parameters, due to the multivariate responses. For the case of 2×2 categories, the maximum number of design points is 4, whereas for the case of 2×4 categories, there are 6 design points at a maximum.

The results presented above show that for this model with a reasonable number of categories, we can derive locally D-optimal designs with few design points, that are thus applicable in practice.

3 Sequential Approach and Simulation Study

We now use the model described above in the setting of the Sequential Locally Optimal Design (SLOD, cf. Roth, 2009a). SLOD generally works as follows. We start with a usual 3+3 design, and continue until we can estimate the parameters in a chosen model. Then we continue with the design that maximizes the cumulative information of the experiment with respect to a specified optimality criterion by adding a fixed number of additional subjects. The design space can be restricted to satisfy safety constraints. We re-estimate the model parameters and adjust the design space and the design after each cohort, and stop when the pre-specified maximum sample size is reached.

Here we apply SLOD based on a logistic, a proportional odds and the bivariate model derived above with 2×2 categories. We perform simulations to compare the properties of these approaches with the traditional 3+3 design.

We display some of the results of a simulation study. We only consider one dose response scenario and only display the results based on the D-criterion and a cohort size of one. We only used the pre-specified sequence of doses, thus being able to apply a simple grid search to find the best doses for the subsequent cohort. Additionally we introduced a safety constraint such that the highest admissible dose is one dose step above the currently estimated MTD.

We assumed a dose-response scenario using the doses 0.6, 1.2, 2, 3, 4, 5.3, 7, 9, 12.4, 16.5, 22 and 29.4 mg. The true dose toxicity relationship is either given by

- a logistic model with parameter $\alpha = 30$ and $\beta = 7.67$ or
- a proportional odds model with parameters $\alpha_1 = 14$, $\alpha_2 = 22$, $\alpha_3 = 30$ and $\beta = 7.67$ or
- the bivariate model described above with parameter $\alpha = 30$, $\beta = 7.67$, $\mu = 20$, $\sigma = 5$ and $\tau = 0$,

where the parameters α_j and β describe the dose toxicity relationship, whereas parameters μ and σ are associated with the dose efficacy relationship.

Defining DLTs in the proportional odds model as toxicities observed in the highest of the four categories, which corresponds to the DLT definition within the framework of the logistic model, the MTD is 22mg for all of the three models. Additionally, in the bivariate model, the minimum effective dose (minED), being the lowest dose where the estimated probability of efficacy is larger than $\frac{1}{3}$, is 18.5 mg.

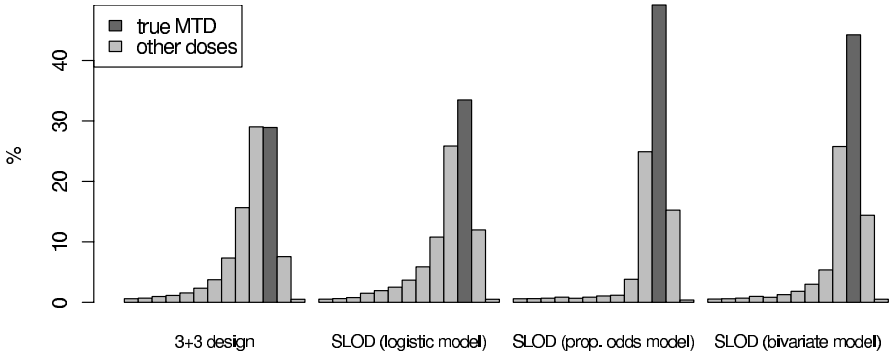


Fig. 2: Percentage of each dose being estimated as the *MTD* for the different methods.

For each of the four approaches, 10 000 simulation runs were conducted. In Figure 2, for each method, the percentage of correctly estimated *MTDs* and the distribution of the estimated *MTDs* is displayed.

Other features of the respective methods are given in Table 1. The average number of subjects (\bar{N}), the average number of *DLTs*, the average number of patients treated above the *MTD* and the MSE of the estimated *MTD* is displayed for all methods. For SLOD, the ratio of the upper and lower 95% confidence interval for the *MTD* is given.

Table 1: Average number of subjects (\bar{N}), of observed *DLTs* (\bar{N}_{DLT}), of subjects treated with doses above the *MTD* ($\bar{N}_{>MTD}$), mean squared error for the *MTD* (MSE *MTD*) and minimum effective dose (MSE minED), and median width of the 95% confidence intervals for the *MTD* (CI *MTD*) and the minimum effective dose (CI minED) for the different methods.

	SLOD			
	3+3 design	logistic model	prop. odds model	bivariate model
\bar{N}	38.43	35.35	37.94	37.97
\bar{N}_{DLT}	3.44	3.81	5.61	4.81
$\bar{N}_{>MTD}$	1.61	3.86	6.18	5.42
MSE <i>MTD</i>	73.05	69.16	39.39	47.25
CI <i>MTD</i>		1.95	1.96	2.03

For the approach based on the bivariate model, additional information on the performance in estimating the minimum effective dose can be obtained from the experiments. In these simulations, the MSE of the estimated minED is 8.83 and the ratio of upper and lower limit of the 95% confidence interval for the minED is 1.87.

The results of this simulation study show the SLOD generally performs better than the traditional 3+3 design with respect to finding the correct *MTD*. At the same time, the risk for the patients is slightly increased with SLOD. Considering the bivariate model, we observe that the precision of the estimated *MTD* is in the

same range as for the logistic model, but with the advantage of additionally being able to estimate the minED with considerable precision.

4 Discussion

We presented a bivariate model that is appropriate for simultaneously analyzing a categorical toxicity outcome and a binary efficacy outcome. The locally D-optimal designs in this model comprise relatively few design points and thus are applicable in practice. We presented a sequential approach where we can apply the optimal designs for this bivariate model, and thus introduced the possibility of incorporating additional information on both toxicity and efficacy in early phase clinical trials.

The simulations show that the suggested approach performs better than the traditional 3+3 design. Accounting for the efficacy endpoint in the design does not downgrade the information that can be obtained for the toxicity endpoint, yet valuable information on the efficacy endpoint can be gained.

Acknowledgements The author is very grateful to Rainer Schwabe (University of Magdeburg) and Thomas Schmelter (Bayer Schering Pharma AG) for their support.

References

- Dragalin, V. and Fedorov, V. (2006). Adaptive designs for dose-finding based on efficacy-toxicity response. *Journal of Statistical Planning and Inference* 136, 1800–1823.
- Ford, I., Torsney, B., and Wu, C. F. J. (1992). The use of a canonical form in the construction of locally optimal designs for non-linear problems. *Journal of the Royal Statistical Society, Series B* 54, 569–583.
- Gerke, O. and Siedentop, H. (2007). Optimal phase I dose escalation trial designs in oncology - A simulation study. *Statistics in Medicine* 27, 5329–5344.
- Ivanova, A. (2006). Dose-finding in oncology – nonparametric methods. In Ting, N., Editor, *Dose Finding in Drug Development*, 49–58. New York: Springer-Verlag.
- Kotz, S., Balakrishnan, N., and Johnson, N. (2000). *Continuous Multivariate Distributions, Vol. 1*. New York: Wiley.
- National Cancer Institute (2006). *Common Terminology Criteria for Adverse Events*.
- Roth, K. (2009a). Adaptive designs for dose escalation studies. In Ermakov, S., Melas, V., and Pepelyshev, A., editors, *Proceedings of the 6th Saint Petersburg Workshop on Simulation*, 133–142.
- Roth, K. (2009b). *Optimal Design for Dose Finding Studies on Safety and Efficacy*. PhD thesis, Otto-von-Guericke-Universität Magdeburg, Germany.

A Radial Scanning Statistic for Selecting Space-filling Designs in Computer Experiments

Olivier Roustant, Jessica Franco, Laurent Carraro and Astrid Jourdan

Abstract In the study of computer codes, filling space as uniformly as possible is important to describe the complexity of the investigated phenomenon. However, this property is not conserved by reducing the dimension. Some numeric experiment designs are conceived in this sense as Latin hypercubes or orthogonal arrays, but they consider only the projections onto the axes or the coordinate planes. We introduce a statistic which allows studying the good distribution of points according to all 1-dimensional projections. By angularly scanning the domain, we obtain a useful graphical representation. The advantages of this new tool are demonstrated on usual space-filling designs. Graphical, decisional and dimensionality issues are discussed.

1 Introduction

For the last 15 years or so, the design of experiments theory initiated by Fisher (1926) has experienced a revival for the analysis of costly industrial computer codes. This development has led to at least two major changes. First, these codes represent phenomena of an increasing complexity, which implies that the corresponding models are often nonlinear and/or nonparametric. Second, the experiment itself is different. Numerical experiments are simulations and, except for stochastic codes, produce the same response for identical conditions. Thus, replications are useless.

In this new paradigm, the experiment planning methods are different. For example, when the code is to be analyzed for the first time, one often tries to satisfy the following two requirements. Firstly, distribute the points in the space as uniformly as possible to catch non-linearities; this excludes repetitions also. Secondly, this space coverage should remain well-distributed even when the effective dimension is lowered. The first requirement was the starting point of research work in space-filling designs (SFD) (see e.g. Fang, Li, and Sudjianto 2006; Santner, Williams, and

O. Roustant

Ecole Nat. Sup. des Mines, Saint-Etienne, France, e-mail: roustant@emse.fr

J. Franco

Total, Pau, France, e-mail: jessica.franco@total.com

L. Carraro

Telecom Saint-Etienne, France, e-mail: laurent.carraro@telecom-st-etienne.fr

A. Jourdan

Ecole Int. des Sc. du Trait. de l'Inform., Pau, France, e-mail: astrid.jourdan@eisti.fr

Notz 2003). The second requirement stems from the observation that codes often depend only on a few influential variables, which may be either the given variables themselves or linear combinations of these variables. Note that dimension reducing techniques like KDR (Fukumizu, Bach, and Jordan 2009) or SIR (Li 1991) effectively identify the subspace generated by the influential variables. Hence, it is desirable that the space-filling property should be also satisfied in the projection onto subspaces. A good challenger is the uniform random design (UD). However, the sampled points are often gathered or unequally spaced. Other candidates are Latin hypercube designs (LHD), orthogonal arrays (OA), and low-discrepancy sequences (LDS) (Koehler and Owen 1996; Niederreiter 1987). Nevertheless LHDs and OAs consider the projections onto margins, which is not sufficient if, for example, the code is a function of one linear combination of the variables. Finally, some LDSs behave poorly in projection.

The aim of this article is to introduce a statistic based on *all* 1-dimensional orthogonal projections to check uniformity of an experimental design. Some work can be found in the literature of random numbers generators testing. For instance, the spectral test measures the maximal distance of points contained in oblique parallel planes (Knuth 1997; Ripley 1987). However, the test applies to designs that exhibit a lattice structure, which is not always the case in computer experiments since SFDs are often obtained by sampling or scrambling. There are also some uniformity tests to detect particular features such as clustering (see e.g. L'Ecuyer and Simard 2007). Nevertheless, the orthogonal projection of a uniform distribution onto an oblique direction is not uniform, which requires an adaptation.

The article is structured as follows. In section 2, we introduce the radial scanning statistic and its the associated visualization tool, and address graphical and decisional issues. In section 3, we show examples of applications for selecting space-filling designs. In section 4, dimensionality issues and further researches are discussed. Software is available at <http://www.emse.fr/roustant>.

2 The Radial Scanning Statistic

As a motivating example, let us consider a simulator depending on 8 variables x_1, \dots, x_8 on the cubic domain $[-1, 1]^8$. As a first stage, a reasonable way to study its behaviour is to evaluate it at few points defined by a space-filling design. For instance, let us choose an 80-point¹ Sobol LDS Niederreiter (1987). As shown in Figure 1, this is not always a good choice. Indeed, looking at the orthogonal projections onto the axis corresponding to the angle $\frac{3\pi}{4}$, we observe that if the code is a function of $x_2 - x_7$, the information brought by the 80 design points comes down to only 16 different values. This problem was automatically detected by what we call the “radial scanning statistic” (RSS). This tool scans angularly every 2-dimensional (2D) cubic domain and detects the angles for which the projected points are not

¹ We follow the rule argued by Loepky, Sacks, and Welch (2008) to select the sample size as approximately 10 times the problem dimension.




Fig. 1: Defect detection of 8D-Sobol low discrepancy sequence of length 80: projected points onto dimensions (2,7) [left], RSS curve [middle] and projected points onto the worst axis [right].

“well” distributed. Here, the 8D design is first projected orthogonally onto all the 28 coordinate planes, which provides 2D designs for which we can compute the RSS. The worst case is detected for the dimensions (2,7). In that coordinate plane, note that the two detected angles $\frac{\pi}{4}, \frac{3\pi}{4}$ correspond to situations where there are many replications in the projection sample.

We now turn to the definition of the radial scanning statistic. It can be theoretically defined for any k -dimensional cubic domain. For the sake of simplicity, we present the 2D case. Let $\Omega = [-1, 1]^2$ and let $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}$ be an experimental design. For a given angle θ , denote by L_θ the straight line in the direction θ and by $\Pi_\theta(\mathbf{x}^{(1)}), \dots, \Pi_\theta(\mathbf{x}^{(n)})$ the orthogonal projection of design points onto L_θ . The RSS tests the null hypothesis H_0 : “ $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}$ is a sample of the uniform distribution in Ω ”. Ideally we would like that, for any θ , the projected points $\Pi_\theta(\mathbf{x}^{(i)})$ are close to their theoretical distribution μ_θ under H_0 . This is a standard *goodness-of-fit* (GOF) problem. As we see next, μ_θ is perfectly known, so the problem reduces to testing for the uniform distribution after applying the Probability Integral Transformation F_θ (see e.g. D’Agostino and Stephens 1986, Chapter 4), where F_θ is the cdf of μ_θ . Finally the RSS is defined as follows:

Definition 1. Let S be a GOF statistic for the uniform distribution. For an angle $\theta \in [0, 2\pi]$, the **radial scanning statistic** (RSS) is defined by:

$$R^S(\theta) = S \left\{ F_\theta \circ \Pi_\theta(\mathbf{x}^{(1)}), \dots, F_\theta \circ \Pi_\theta(\mathbf{x}^{(n)}) \right\} \tag{1}$$

For simplicity’s sake, RSS will also refer to the *family* of statistics $R^S(\theta)$ indexed by θ ; the corresponding polar curve is called an *RSS curve*.

The cdf F_θ of μ_θ is known analytically in two or higher dimensions. First remark that μ_θ is not the uniform distribution, except when θ corresponds to a coordinate axis. For instance when $\theta = \frac{\pi}{4}$, $\Pi_\theta(\mathbf{x}^{(i)}) = \frac{\sqrt{2}}{2}(x_1^{(i)} + x_2^{(i)})$ so that μ_θ is a triangular distribution. In the general case, the projection onto L_θ is a linear combination of independent random variables of uniform distribution, which leads to a traditional

problem of probabilities first solved by Lagrange in the 18th century. The result is the following (see proof and discussion in Shiu 1987):

Proposition 1 *If X is a random vector uniformly distributed over the hypercube $\Omega = [-1, 1]^d$ and Z is the projection of X onto the straight line generated by a unitary vector \mathbf{a} such that $a_j \neq 0, \forall j \in \{1, \dots, d\}$, then the cdf of Z is given by:*

$$F_Z(z) = \frac{1}{\prod_{j=1}^d 2a_j} \sum_{\mathbf{s} \in \{-1, 1\}^d} \varepsilon(\mathbf{s}) \frac{(x + \langle \mathbf{s}, \mathbf{a} \rangle)_+^d}{d!}, \tag{2}$$

where $\varepsilon(\mathbf{s}) = \prod_{j=1}^d s_j$, $\langle \cdot, \cdot \rangle$ is the usual scalar product and $y_+ = \max(y, 0)$. As a result, for a given axis, Z admits a piecewise linear density whose nodes correspond to the projections of the domain corners.

2.1 Selecting a Goodness-of-fit Test for the Uniform Distribution

A key point is the choice of S . First remark that a single statistical test will not detect all departures from uniformity. In our case, it seems reasonable to focus on replications or *clustering*, since they result in losses of information, as explained in the previous section. Then our choice should be guided by the power of the test. For instance, the usual *edf* statistics such as Kolmogorov-Smirnov (KS) or Cramér-Von Mises (CVM) fail to detect clustering (L'Ecuyer and Simard 2007). On the other hand, the spacings transformation used to increase power in testing random numbers generators results in over-detection: only uniform random designs will pass the test. A compromise is reached with the statistics based on *spacings* (D'Agostino and Stephens 1986). If U_1, \dots, U_n is a sample from a uniform distribution on $[0, 1]$, and $U_{(1)}, \dots, U_{(n)}$ denotes the sorted values in increasing order, the spacings D_i are defined as:

$$D_i = U_{(i)} - U_{(i-1)}, \quad i = 1, \dots, n+1 \tag{3}$$

where, by convention, $U_{(0)} = 0$ and $U_{(n+1)} = 1$. An example is the Greenwood statistic $G_n = \sum_{i=1}^{n+1} D_i^2$, equivalent² to $V_n = \sum_{i=1}^{n+1} (D_i - \frac{1}{n+1})^2$. As an illustration, we plot the RSS curve of a 7×7 factorial design obtained with KS, CVM and Greenwood statistics (Figure 2). Remark that KS is a L^∞ -type statistic, while CVM and Greenwood are L^2 -type. In each case, the worst angle corresponds to the coordinate axis onto which the 49 values project to 7 different values, but the only case of rejection is for Greenwood. In addition, the other undesirable angles $\frac{\pi}{4}, \frac{3\pi}{4}$ are also detected.

The reason why such a difference is visible for detection is that the statistics based on cdf work roughly on differences between (empirical and theoretical) *cumulative* probabilities $U_{(i)} - \frac{i}{n+1}$, while those based on spacings work on the corresponding differences $D_i - \frac{1}{n+1} = (U_{(i)} - \frac{i}{n+1}) - (U_{(i-1)} - \frac{i-1}{n+1})$. In the first case, local deviations from uniformity are cumulated, and may be smoothed too much to be detected. Another advantage of spacings is that the corresponding statistics look

² Indeed it is easily shown that $V_n = G_n - \frac{1}{n+1}$.




Fig. 2: Defect detection of a 2D Factorial design of size 49 with three GOF statistics: CVM [left], KS [middle] and Greenwood [right]. The solid circle corresponds to the threshold at level 5%.

robust to small changes of the domain boundaries. As a conclusion, we recommend using statistics based on spacings.

2.2 Graphical Properties

When looking at the examples in Figures 1 and 2, it appears that the RSS curve is very irregular. Nevertheless, it is continuous in most cases as shown in the proposition below – an elementary proof is based on (1).

Proposition 2 *Denote by S the GOF statistic used for testing uniformity, considered as a function of the ordered sample values. If S is continuous, then R^S is continuous everywhere. If S is differentiable, then R^S admits left and right derivatives everywhere and is differentiable at $\theta \in [0, 2\pi]$ if and only if the projections onto L_θ are all different.*

As a consequence, continuity is ensured for all GOF statistics used in this paper. The differentiability condition is suitable at least for Greenwood and CVM statistics. In practice, we need to use a small discretization step. From our experience, a reasonable value may be 0.5 degrees.

2.3 Decisional Issues

To build a unique statistical test with the radial scanning statistic raises several issues.

First, many designs encountered in computer experiments are deterministic, and statistical tests should be interpreted carefully. Actually, a correct interpretation is possible by considering the *probability of rejection* p_r . For a design generated at random, p_r is the usual *p-value*, defined as $p_r(\theta) = P(R^S(\theta) > R^S_{\text{obs}}(\theta) | H_0)$, where $R^S_{\text{obs}}(\theta)$ is the value taken by $R^S(\theta)$ at the random design. For a given θ , $p_r(\theta)$ is a $U(0, 1)$ random variable, and H_0 is rejected at level α if $p_r(\theta) < \alpha$. Now if the design is deterministic, p_r is defined as $p_r(\theta) = P\{R^S(\theta) > r^S_{\text{obs}}(\theta)\}$, where $r^S_{\text{obs}}(\theta)$ is now deterministic. Then rejecting at level α simply means that the proportion of UD that behave worse is less than α .

More problematic is the *multiple testing* problem over angles and dimensions. Indeed, for each coordinate plane onto which the design is first projected (see §§2 and 3), we have a family of statistical tests indexed by an angle $\theta \in [0, \pi]$ with null hypothesis $H_0(\theta) = H_0$: “ $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}$ is a sample of the uniform distribution in Ω ”. Multiple testing has been studied intensely since the work of Benjamini and Hochberg (1995). In our case, a direct application seems to be difficult, due to the dependencies between $R^S(\theta)$. A partial solution could be proposed by considering some global statistics, as $\max_{\theta} R^S(\theta)$, which does not depend on θ .

3 Usage and Applications

For the sake of simplicity, the radial scanning statistic was presented in two dimensions. Nevertheless, Proposition 1 is very general and Definition 1 is easily extended to higher dimensions: in 3D, the RSS will give a surface in spherical coordinates and, in general, a hyper-surface indexed by 1D directions. In this section, we detail several applications of the 2D and 3D RSS to defect detection of d -dimensional designs. The general case is discussed in the conclusion.

First remark that uniformity is preserved by orthogonal projections onto coordinate subspaces. Thus, for a d -dimensional design $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}$, the 2D RSS can be computed from the 2D design $(x_{j_1}^{(i)}, x_{j_2}^{(i)})$ for all pairs (j_1, j_2) of dimensions. For three dimensions, we have triplets. The RSS can be used in 2 steps:

1. Automatic defect detection³;
2. Defect visualization in a coordinate plane (or 3D subspace).

This scheme was followed for the 8D Sobol sequence of Figure 1 (see §2). As a second example, consider the 3D random OA of size 49 of Figure 3 (right). It was obtained from a deterministic OA (midpoint OA), by sampling uniformly one point for every cell defined by the 49 points. This OA has strength 2, which implies that its projections onto 1D and 2D coordinate subspaces are well distributed, as can be seen in Figure 3 (left). As a result, no particular feature is noted by the 2D RSS. However, a serious defect is detected by the 3D RSS: there are 4 clusters among the projected points in one direction. This is because the 49 points of the underlying midpoint OA are contained in 4 parallel oblique planes. By adding noise to the data, the perfect alignments disappear but the clusters remain.

Automatic defect detection can be used to compare SFDs. For instance, let us choose the best 8D design of size 80 among popular designs. In Table 1, we have indicated the worst value of the Greenwood statistic over all pairs of dimensions and angles for several LHDs and LDSs. The result for UDs is also mentioned, and will be used as a benchmark. The worst design is clearly the Halton sequence, due to the presence of alignments as well as a large empty region. Among LDSs, the best one is Sobol. This design is not convenient (see Figure 1), but the result can be much improved by “scrambling” (Koehler and Owen 1996). With Owen scrambling, the

³ The RSS typically detects alignments, clustering or empty regions. In computer experiments these features can be viewed as *defects*, since they result in losses of information.




Fig. 3: A 3D random OA of size 49 (2D projections, left) and its 3D RSS plot (right). The worst direction and the corresponding orthogonal projections of design points are represented by solid lines.

Sobol sequences behave as well as uniform designs with respect to 1-dimensional projections. This is also the case for the Audze-Eglais LHD.

Table 1: Worst value of the Greenwood statistic for 8-dimensional SFDs of size 80

Design type ^a	Statistic value ^b
Uniform	0.039 (0.003)
Maximin Latin hypercube	0.048
Audze-Eglais Latin hypercube	0.037
Halton sequence	0.244
Sobol sequence	0.101
Sobol sequence, with Owen scrambling	0.041 (0.006)
Sobol sequence, with Faure-Tezuka scrambling	0.088 (0.010)
Sobol sequence, with Owen + Faure-Tezuka scrambling	0.041 (0.006)

^aLHDs are taken from <http://www.spacefillingdesigns.nl>. Halton and Sobol sequences are computed with the R package `randtoolbox` (<http://www.r-project.org>). ^bFor stochastic designs, the first number is the mean of the results over 100 simulations, and the second (in brackets) their standard deviation.

4 Conclusion and Further Research

The radial scanning statistic is a tool devoted to checking uniformity of d -dimensional designs according to all 1-dimensional orthogonal projections, including non-factorial axes. Despite the decisional issues encountered, the 2D and 3D RSS succeed in automatically detecting the main defects of several 8D popular LHDs and low discrepancy sequences, and in comparing them. Thus, it may be used to select the SFDs that satisfy the initial requirement of filling space *after* dimension reduction.

Further research will address dimensionality issues. As shown with the OA of Figure 3, some defects cannot be seen by projections onto small dimensional subspaces. Therefore, the development of d -dimensional RSS with $d \geq 4$ is under consideration. For computational reasons, an exhaustive search is impossible and an al-

ternative is to use an optimization technique. Another direction is to adapt the radial scanning statistic to orthogonal projections onto oblique 2D or higher subspaces.

Acknowledgements We wish to thank A. Antoniadis, the members of the DICE Consortium (<http://www.dice-consortium.fr>), the participants of ENBIS-DEINDE 2007, as well as two referees for their useful comments. We also thank Chris Yukna for his help in editing.

References

- Benjamini, Y. and Y. Hochberg (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society, Series B* 57, 289–300.
- D’Agostino, R. and M. Stephens (1986). *Goodness-of-fit Techniques*. New York; Marcel Dekker.
- Fang, K.-T., R. Li, and A. Sudjianto (2006). *Design and Modeling for Computer Experiments*. London: Chapman & Hall.
- Fisher, R. A. (1926). The arrangement of field experiments. *Journal of the Ministry of Agriculture* 33, 503–513.
- Fukumizu, K., F. Bach, and M. Jordan (2009). Kernel dimension reduction in regression. *Annals of Statistics* 37, 1871–1905.
- Knuth, D. (1997). *The Art of Computer Programming, Volume 2: Seminumerical Algorithms, 3rd edition*. Reading MA: Addison-Wesley.
- Koehler, J. and A. Owen (1996). Computer experiments. Amsterdam: Elsevier. *Handbook of Statistics* 13, 261–308.
- L’Ecuyer, P. and R. Simard (2007). Testu01: A C library for empirical testing of random number generators. *ACM Transactions on Mathematical Software* 33(4), Article 22.
- Li, K. (1991). Sliced inverse regression for dimension reduction (with discussion). *Journal of the American Statistical Association* 86, 316–342.
- Loeppky, J., J. Sacks, and W. Welch (2008). Choosing the sample size of a computer experiment: A practical guide. Technical report, NISS.
- Niederreiter, H. (1987). Low-discrepancy and low-dispersion sequences. *Journal of Number Theory* 30, 51–70.
- Ripley, B. (1987). *Stochastic Simulation*. New York: Wiley.
- Santner, T., B. Williams, and W. Notz (2003). *The Design and Analysis of Computer Experiments*. New York: Springer-Verlag.
- Shiu, E. S. W. (1987). Convolution of uniform distributions and ruin probability. *Scandinavian Actuarial Journal* 70, 191–197.

Optimal Designs for Two-Colour Microarray Experiments for Estimating Interactions

Katharina Schiffl and Ralf-Dieter Hilgers

Abstract In recent years microarray experiments have become one of the most prominent tools for analyzing gene expressions. However, microarrays are expensive; thus carefully planning of these experiments is fundamental. Since in practical applications researchers are often interested in interactions in multi-factor settings, we derive optimal designs for their estimation in this paper. We will show the optimality of candidate designs using equivalence theorems. The resulting designs help to ensure precise results with minimal resources.

1 Introduction

Nowadays microarray experiments form a widely used tool in gene expression analysis due to the fact that they can screen thousands of genes simultaneously. One of their main goals is to identify differently expressed genes that can be made accountable for a certain disease. Throughout this work we will focus on two channel microarrays. They can hybridize two samples on the same array; mRNA transcripts from two biological samples, called targets, are extracted and labelled with green (Cy3) or red (Cy5) dyes, respectively, and are placed on the microarray. After hybridization, a laser measures the dye fluorescence of each colour for all genes on the array. These intensities correspond to the gene expressions of the considered genes. Higher intensities indicate higher gene expressions. A detailed description of microarray experiments can be found in Wit and McClure (2004).

Since microarrays are expensive, it is fundamental to use appropriate designs to

Katharina Schiffl

RWTH Aachen University, Department of Medical Statistics, 52074 Aachen, Germany, e-mail: kschiffl@ukaachen.de

Ralf-Dieter Hilgers

RWTH Aachen University, Department of Medical Statistics, 52074 Aachen, Germany, e-mail: rhilgers@ukaachen.de

get precise results with minimal resources. Optimal designs assign the samples to the microarrays in such a way as to ensure unbiased estimates with minimal variances of the effects of interest. Thereby, microarray experiments correspond to incomplete block designs with block size two. Design issues for microarray experiments have been investigated intensively in the last years (see, for example, Kerr and Churchill 2001) although most authors focus on one factor of interest. Some authors investigate main effects and first-order interactions. See, for example Glonek and Solomon (2004) or Banerjee and Mukerjee (2007), who consider factorial designs for microarray experiments under the baseline parametrization. Furthermore, Kerr (2006) and Grossmann and Schwabe (2007) derive efficient designs for the estimation of main effects and two-way interactions when all factors have two levels. Another reference is Stanzel and Hilgers (2007), who give approximate designs for the estimation of two-factor interactions. However, in clinical trials scientists are often interested in many factors and their interactions. For example, in experiments with different drugs the interactions of these drugs are often of primary interest. Therefore, the interesting question of optimal designs for estimating interactions in multi factor settings arises and will be considered in this paper. We extend the investigations of Stanzel and Hilgers (2007), who focused on two factors of interest. Section 2 introduces the underlying statistical model describing microarray experiments, gives a definition of the optimality criteria used and presents a short overview of the methods we apply. Optimal designs for the estimation of interactions are derived in Section 3. The paper concludes with a brief discussion.

2 Preliminaries

Many authors have focused on the statistical analysis and modelling of microarray experiments. Kerr, Martin, and Churchill (2000) analyzed two-colour microarray data by analysis of variance (ANOVA) and recommended a model describing the logarithms of the measured intensities dependent on treatment, array, dye and gene effects including interactions, namely

$$\log(y) = T\tau + A\alpha + D\delta + G\gamma + \varepsilon, \quad (1)$$

where $y = (y_1, \dots, y_{2ag})$ is the vector of all observed dye intensities. These dye intensities depend on the treatment effect $\tau = (\tau_1, \dots, \tau_k)$, the array effect $\alpha = (\alpha_1, \dots, \alpha_a)$, the dye used $\delta = (\delta_{\text{green}}, \delta_{\text{red}})$, and the gene investigated $\gamma = (\gamma_1, \dots, \gamma_g)$. The $2ag \times (k + a + 2 + g)$ design matrix is written $[T \mid A \mid D \mid G]$. The error terms $\varepsilon = (\varepsilon_1, \dots, \varepsilon_{2ag})$ are assumed to be independently identically distributed with mean zero and variance σ^2 .

Kerr's work has been extended by many authors. For instance, Landgrebe, Bretz, and Brunner (2006) analyzed the logarithmic ratios of dye intensities of each microarray separately for each gene. Instead of the two observations

$$\begin{aligned} \log(y_{ij\text{green}}) &= \tau_i + \alpha_j + \delta_{\text{green}} + \varepsilon_{ij\text{green}} \quad \text{and} \\ \log(y_{lj\text{red}}) &= \tau_l + \alpha_j + \delta_{\text{red}} + \varepsilon_{lj\text{red}}, \end{aligned}$$

Landgrebe et al. considered the log ratio

$$\log\left(\frac{y_{ij\text{green}}}{y_{lj\text{red}}}\right) = \tau_i - \tau_l + \delta_{\text{green}} - \delta_{\text{red}} + \varepsilon_{ij\text{green}} - \varepsilon_{lj\text{red}}.$$

Therefore, they investigated

$$z = X\tau + W\delta + \eta \tag{2}$$

where $z = (z_1, \dots, z_a)$ is the vector of all log ratios of the dye intensities of each array dependent on the treatment effect $\tau = (\tau_1, \dots, \tau_k)$ and the dye $\delta = (\delta_{\text{green}}, \delta_{\text{red}})$. $[X \mid W]$ is the design matrix, where each row of X consists of exactly one 1 and one -1 , with all other entries equal to zero. W is equal to $\begin{pmatrix} 1 & 1 & \dots & 1 \\ -1 & -1 & \dots & -1 \end{pmatrix}^T$ and η is the independently identically distributed error with mean zero and equal variance σ^2 . Throughout this paper we will consider model (2) with the novel aspect that n factors of interest are involved. Thus, we denote by $\tau = (\tau_{11\dots 1}, \dots, \tau_{11\dots k_n}; \dots; \tau_{k_1 k_2 \dots k_{n-1} 1}, \dots, \tau_{k_1 k_2 \dots k_n})$ the vector of all combinations of factor levels of interest, where k_i is the number of factor levels of factor i , $i \in \{1, \dots, n\}$. For $n = 1$ the model reduces to model (2).

Optimal designs for a given contrast set C minimize a particular score function of the variance

$$\text{Var}\left\{C^T \begin{pmatrix} \hat{\tau} \\ \hat{\delta} \end{pmatrix}\right\} = \sigma^2 C^T ([X \mid W]^T [X \mid W])^{-1} C. \tag{3}$$

Throughout this paper we will consider the family of ϕ_p -optimality criteria for $p \in (-\infty, 1]$. Thus the score function

$$\left(\frac{1}{r} \sum_{j=1}^r \lambda_j^{-p}\right)^{-\frac{1}{p}}$$

is minimized for the positive eigenvalues $\lambda_1, \dots, \lambda_r$ of the variance-covariance matrix of the estimated parameters. For $p = 0$ the expression $\left(\prod_{j=1}^r \lambda_j\right)^{1/r}$ is minimized to yield D-optimality; for $p = -1$ the A-optimality criterion results. Since we are not interested in the dye effect, the last two rows of all contrast matrices C contain only zero entries. Therefore, from now on, we will ignore the dye effect in model 2 although all results go through with the dye effect present. One method in approximate theory to show ϕ_p -optimality of a given candidate design is the following equivalence theorem Pukelsheim 1993:

Theorem 1. A design is ϕ_p -optimal, $p \in (-\infty, 1]$, for estimation of the contrast set $C^T \tau$ if and only if there exists a generalized inverse $G = (X^T P X)^-$ of $X^T P X$ that satisfies the normality inequalities

$$x^T G C (C^T G C)^+ (C^T G C)^{1-p} (C^T G C)^+ C^T G^T x \leq \text{Tr} \left((C^T G C)^+ (C^T G C)^{1-p} \right) \tag{4}$$

for all possible design points $x \in \mathcal{X}$. \mathcal{X} is the set of $k_1 \cdot \dots \cdot k_n$ dimensional column vectors with exactly one entry equal to 1 and exactly one entry equal to -1 , all other entries are zero, $\mathcal{X} = \{x \in \{-1, 0, 1\}^{k_1 \cdot \dots \cdot k_n} : \exists! i \text{ with } x_i = 1 \wedge \exists! j \text{ with } x_j = -1\}$. $\exists!$ stands for the expression “there is one and only one”. P is the diagonal matrix containing the optimal weights for all design points listed in the design matrix X . $(C^T G C)^+$ denotes the Moore-Penrose Inverse of $(C^T G C)$.

Stanzel and Hilgers (2007) derived optimal designs for the estimation of two-way interactions, i.e. they considered contrast matrices denoted by $P_{k_1}^T \otimes P_{k_2}^T$, where k_1 and k_2 are the levels of the two factors of interest. P_k is recursive defined by $P_2 := [1, -1]$

and $P_k := \left[\begin{array}{c|c} 1_{k-1} & -I_{k-1} \\ \hline 0_{\binom{k-2}{2}} & P_{k-1} \end{array} \right]$ for all $k \in \mathbb{N}$, whereas 1_k and 0_k are k -dimensional column vectors with all entries equal to 1 and 0, respectively. I_k denotes the $k \times k$ identity matrix and \otimes denotes the Kronecker product.

Hinkelmann and Kempthorne (2005) denoted three-way interactions as $P_{k_1}^T \otimes P_{k_2}^T \otimes P_{k_3}^T$, where k_1, k_2 and k_3 are the levels of the three factors of interest. Interactions in multi-factor settings are denoted by $P_{k_1}^T \otimes P_{k_2}^T \otimes \dots \otimes P_{k_n}^T$. In the following section, we consider the important question of optimal designs for interactions. We will also extend the results of Stanzel and Hilgers (2007) in the two-factor setting and consider contrast sets $P_{k_1}^T \otimes C$ for arbitrary contrast matrices C .

3 Optimal Designs

We derive ϕ_p -optimal designs, $p \in (-\infty, 1]$, for the estimation of interactions in multi-factor settings for model 2. Since we are not interested in the dyes, in this section we will ignore the dye effect in model 2. Obviously, all results can be shown similarly with the dye effect; only, in some cases, dye swaps should be added to provide optimality.

3.1 Interactions in Multi-factor Settings for the Estimation of All Pairwise Comparisons

The following theorem gives optimal designs for the estimation of interactions $C_n = P_{k_1}^T \otimes P_{k_2}^T \otimes \dots \otimes P_{k_n}^T$ in multi-factor settings, if we are interested in all pairwise comparisons. In practice, these results are most interesting for $n = 3$ factors.

Theorem 2. Consider the multi-factor model (2) with $n \in \mathbb{N}$ factors of interest: k_i , $i \in \{1, \dots, n\}$, denotes the number of levels of factor i , w.l.o.g. $2 \leq k_1 \leq \dots \leq k_n$. The ϕ_p -optimal design, $p \leq 1$, for estimation of the interactions $C_n^T = P_{k_1} \otimes P_{k_2} \otimes \dots \otimes P_{k_n}$ is the design with design matrix $X_n = P_{k_1} \otimes I_{k_2} \otimes \dots \otimes I_{k_n}$ and equal weights $P^{(n)} = \frac{1}{k_2 \dots k_n \binom{k_1}{2}} I_{k_2 \dots k_n} \binom{k_1}{2}$.

The proof is relegated to the Appendix. At this point we know the optimal design if our primary interest is only the multi-factor interactions. First, the factor with the smallest number of levels is chosen. Then, all pairwise comparisons of these factor levels are allocated on microarrays for each factor level combination of the remaining factors. In many practical applications the number of factor levels does not exceed $k_i = 3$, $i \in \{1, \dots, n\}$. See, for example, Churchill (2002). The derived optimal designs also perform well for estimation of the main effects, although these are not our primary interest. In particular, the constructed designs are also optimal for one main effect (see Stanzel and Hilgers 2007).

3.2 Interactions in Two-factor Settings

In addition to the estimation of the interactions for all pairwise comparisons, biologists are often interested in further contrast sets, for instance comparisons with a control treatment, Helmert contrasts or “all to next” contrasts. Considering two factors with k_1 and k_2 factor levels, optimal designs for the contrasts $P_{k_1}^T \otimes C$ are investigated for arbitrary contrast matrices C .

Theorem 3. Suppose that the design with the $a \times k_2$ design matrix X and the $a \times a$ weight matrix P is ϕ_p -optimal, $p \leq 1$, for estimation of the contrasts $C^T \tau$ in the one-factor model (2) without dye. Then, the design with design matrix $\tilde{X} = I_{k_1} \otimes X$ and weight matrix $\tilde{P} = \frac{1}{k_1} (I_{k_1} \otimes P)$ is ϕ_p -optimal in the two-factor model (2) without dye for the estimation of the contrast set $\tilde{C} \tau$ with $\tilde{C} = P_{k_1}^T \otimes C$, if $a_{ij} \leq 0$, $i \neq j$ and $a_{ii} \leq \frac{k_1 - 1}{2k_1} \text{const}$ with $\text{const} = \text{Tr}\{(C^T GC)^+ (C^T GC)^{1-p}\}$ and $A = GC(C^T GC)^+ (C^T GC)^{1-p} (C^T GC) + C^T G^T$.

The proof of this theorem is similar to the proof of Theorem 2. Considering $C = P_{k_2}^T$ the conditions in Theorem 3 reduce to $k_2 \leq k_1$. Therefore, for $C = P_{k_2}^T$ Theorem 3 yields the same results as stated in Stanzel and Hilgers (2007). However, Theorem 3 can also be applied for other contrast sets. For instance, the conditions in Theorem 3 also reduce to $k_2 \leq k_1$, if the treatment control comparisons are of interest, i.e. $\tau_0 - \tau_i$, $i \in \{1, \dots, k_2\}$ are of interest, where τ_0 denotes the effect of the control treatment.

4 Discussion

Interactions are of high interest in many applications. It is therefore important to investigate optimal designs for their estimation. In this paper we have constructed optimal designs for estimating interactions in multi-factor settings. In practice our results have the potential to lead to precise results with minimal resources, so reducing the number of microarrays required in applications.

Acknowledgements We thank the referees for many constructive suggestions. This research was financially supported by the BMBF within the joint research project SKAVOE (Foerderkennzeichen: 03HIPAB4).

Appendix: Proof of Theorem 2

We prove Theorem 2 by induction for fixed p using Theorem 1. Stanzel and Hilgers (2007) provide the basis of the induction for $n = 2$. Define $G_n := (X_n^T P^{(n)} X_n)^+$, $E_n := G_n C_n (G_n^T G_n C_n)^+ (C_n^T G_n C_n)^{1-p} (C_n^T G_n C_n)^+ C_n^T G_n^T$ & $\text{const}_n := \text{Tr}\{(C_n^T G_n C_n)^+ (C_n^T G_n C_n)^{1-p}\}$. Let J_k be the $k \times k$ matrix with all entries equal to one. Using $E_2 = 2^{p-1} k_2^{1-2p} k_1^{-p} (k_1 - 1)^{1-p} \{(I_{k_1} - \frac{1}{k_1} J_{k_1}) \otimes (I_{k_2} - \frac{1}{k_2} J_{k_2})\}$ and $\text{const}_2 = 2^p k_1^{-p} (k_1 - 1)^{1-p} (k_2 - 1) k_2^{-2p}$ for $k_1 \leq k_2$ (see Stanzel and Hilgers 2007), the following expressions can easily be shown:

$$E_n = 2^{p-1} k_1^{-p} (k_1 - 1)^{1-p} \prod_{i=2}^n k_i^{1-2p} \{(I_{k_1} - \frac{1}{k_1} J_{k_1}) \otimes \dots \otimes (I_{k_n} - \frac{1}{k_n} J_{k_n})\},$$

$$\text{const}_n = 2^p k_1^{-p} (k_1 - 1)^{1-p} \prod_{i=2}^n k_i^{-2p} (k_i - 1) = \left(\frac{k_1 - 1}{2k_1}\right)^{-p} \prod_{i=1}^n k_i^{-2p} (k_i - 1).$$

Since X_{n-1} with weights $P^{(n-1)}$ is ϕ_p -optimal for estimation of C_{n-1} , we know that the normality inequalities $x^T E_{n-1} x \leq \text{const}_{n-1}$ hold for all $x \in \mathcal{X}_{n-1} = [x \in \{-1, 0, 1\}^{k_1 \dots k_{n-1}} | \exists! i : x_i = -1 \wedge \exists! j : x_j = 1]$. Therefore, we have to show $y^T E_n y \leq \text{const}_n$ or equivalently $y^T (E_{n-1} \otimes (I_{k_n} - \frac{1}{k_n} J_{k_n})) y \leq (k_n - 1) \text{const}_{n-1}$ for all $y \in \mathcal{X}_n = [y \in \{-1, 0, 1\}^{k_1 \dots k_n} | \exists! i : y_i = -1 \wedge \exists! j : y_j = 1]$. Suppose $1 \leq h, l \leq k_1 k_2 \dots k_{n-1}$ and $1 \leq i, j \leq k_n$ and partition

$$y = (y_1; \dots; y_{k_1 k_2 \dots k_{n-1}}) = (y_{11}, \dots, y_{1k_n}; \dots; y_{k_1 k_2 \dots k_{n-1} 1}, \dots, y_{k_1 k_2 \dots k_{n-1} k_n}).$$

The term $x^T E_{n-1} x \leq \text{const}_{n-1}$ for $x \in \mathcal{X}_{n-1}$ with $x_h = 1$ and $x_l = -1$ is equivalent to $e_{hh}^{(n-1)} + e_{ll}^{(n-1)} - e_{hl}^{(n-1)} - e_{lh}^{(n-1)} \leq \text{const}_{n-1}$. We will distinguish three cases. Firstly, suppose $h = l$ and $i \neq j$, w.l.o.g. $i < j$ and $y_{hi} = 1$ as well as $y_{lj} = -1$. Then

$$y^T k_n (E_{n-1} \otimes (I_{k_n} - \frac{1}{k_n} J_{k_n})) y = 2k_n e_{hh}^{(n-1)}.$$

Due to the structure of E_n we know

$$e_{hh}^{(n)} = \left(\frac{k_1 - 1}{2k_1} \right)^{1-p} \prod_{i=1}^n k_i^{-2p} (k_i - 1).$$

Therefore we have to show

$$\begin{aligned} 2k_n \left(\frac{k_1 - 1}{2k_1} \right)^{1-p} \prod_{i=1}^{n-1} k_i^{-2p} (k_i - 1) &\leq \left(\frac{k_1 - 1}{2k_1} \right)^{-p} \prod_{i=1}^{n-1} k_i^{-2p} (k_i - 1) (k_n - 1) \\ \Leftrightarrow k_n (k_1 - 1) &\leq k_1 (k_n - 1). \end{aligned}$$

This inequality is fulfilled for $k_1 \leq k_n$. For $h \neq l$ and $i = j$ the inequality can be shown in a similar way. Assuming $h \neq l$ and $i \neq j$, we get

$$\begin{aligned} y^T k_n (E_{n-1} \otimes (I_{k_n} - \frac{1}{k_n} J_{k_n})) y &= (k_n - 1) (e_{hh}^{(n-1)} + e_{ll}^{(n-1)}) + 2e_{lh}^{(n-1)} \\ &\leq (k_n - 1) \text{const}_{n-1} + 2k_n e_{hl}^{(n-1)}. \end{aligned}$$

Due to the structure of E_{n-1} we know, for $h \neq l$,

$$e_{lh}^{(n-1)} \in \left\{ \frac{(\prod_{i=2}^n k_i^{1-2p}) \cdot (k_1 - 1)^{q_1} (k_2 - 1)^{q_2} \dots (k_{n-1} - 1)^{q_{n-1}} (-1)^{q_0}}{2^{1-p} k_1^p (k_1 - 1)^{p-1} \cdot k_1 k_2 \dots k_{n-1}} \mid \text{for } i \in \{1, \dots, n-1\} \text{ is } q_i \in \{0, 1\} \text{ and } q_0 = (n-1) - \sum_{i=1}^{n-1} q_i \right\}.$$

Obviously, $e_{hl}^{(n-1)} \leq 0$ for q_0 odd. Hence, let q_0 be even and define

$$h_n(k_1, \dots, k_n) := 2^{p-1} k_1^{-p} (k_1 - 1)^{1-p} \left(\prod_{i=2}^n k_i^{1-2p} \right).$$

Therefore, $\text{const}_{n-1} = 2h_{n-1}(k_1, \dots, k_{n-1}) \left(\prod_{i=2}^{n-1} \frac{k_i - 1}{k_i} \right)$ and we have to prove the following inequalities

$$\begin{aligned} &\frac{2(k_n - 1)h(k_1, \dots, k_{n-1})(k_1 - 1)(k_2 - 1) \dots (k_{n-1} - 1)}{k_1 k_2 \dots k_{n-1}} \\ &+ \frac{2h(k_1, \dots, k_{n-1})(k_1 - 1)^{q_1} \dots (k_{n-1} - 1)^{q_{n-1}} (-1)^{q_0}}{k_1 k_2 \dots k_{n-1}} \\ &\leq 2(k_n - 1)h(k_1, \dots, k_{n-1}) \left(\prod_{i=2}^{n-1} \frac{k_i - 1}{k_i} \right) \frac{k_1}{k_1} \\ \Leftrightarrow &(k_1 - 1)^{q_1} (k_2 - 1)^{q_2} \dots (k_{n-1} - 1)^{q_{n-1}} (-1)^{q_0} \leq (k_2 - 1)(k_3 - 1) \dots (k_n - 1). \end{aligned}$$

This inequality is fulfilled for $k_n \geq k_1 \geq 2$ and for all $q_i \in \{0, 1\}$ for $i \in \{1, \dots, n-1\}$ and $q_0 = n - \sum_{i=1}^n q_i$, since all $\frac{\binom{k_i-1}{q_i}}{k_i-1} \leq 1$. This completes the proof.

References

- Banerjee, T. and R. Mukerjee (2007). Optimal factorial designs for cDNA microarray experiments. *Annals of Applied Statistics* 2, 366–385.
- Churchill, G. (2002). Fundamentals of experimental design for cDNA microarrays. *Nature Genetics* 32, 490–495.
- Glonek, G. and P. J. Solomon (2004). Factorial and time course designs for cDNA microarray experiments. *Biostatistics* 5, 89–111.
- Grossmann, H. and R. Schwabe (2007). The relationship between optimal designs for microarray and paired comparison experiments. (Preprint).
- Hinkelmann, K. and O. Kempthorne (2005). *Design and Analysis of Experiments, Vol. 2*. New York: Wiley.
- Kerr, K. F. (2006). Efficient 2^k factorial designs for blocks of size 2 with microarray applications. *Journal of Quality Technology* 38, 309–317.
- Kerr, M. K. and G. Churchill (2001). Experimental design for gene expression microarrays. *Biostatistics* 2, 183–201.
- Kerr, M. K., M. Martin, and G. A. Churchill (2000). Analysis of variance for gene expression microarray data. *Journal of Computational Biology* 7, 819–837.
- Landgrebe, J., F. Bretz, and E. Brunner (2006). Efficient design and analysis of two colour factorial microarray experiments. *Computational Statistics and Data Analysis* 50, 499–517.
- Pukelsheim, F. (1993). *Optimal Design of Experiments*. New York: Wiley.
- Stanzel, S. and R. Hilgers (2007). The within-B-swap (BS) design is A- and D-optimal for estimating the linear contrast for the treatment effect in 3-factorial cDNA microarray experiments. In J. López-Fidalgo, J. Rodríguez-Díaz, and B. Torsney (Eds.), *mODa 8 - Advances in Model-Oriented Design and Analysis*, pp. 197–204. Heidelberg: Physica.
- Wit, E. and J. McClure (2004). *Statistics for Microarrays*. Chichester: Wiley.

Discrimination Between Random and Fixed Effect Logistic Regression Models

Chiara Tommasi, Maria Teresa Santos-Martín and Juan Manuel Rodríguez-Díaz

Abstract D_3 - and KL-optimum designs are computed for discriminating between univariate logistic regression models with or without random effects. Both these competing optimum designs are constructed numerically. The main problem in finding them is the computation of some integrals at each step of the numerical procedure. In order to improve the convergence speed of this numerical procedure some integral approximations are suggested.

1 Introduction

In biosciences random effects play a fundamental role and the interest in finding optimum designs in this context is growing. For instance, Mentré, Mallet, and Baccar (1997), Patan and Bogacka (2007) and Graßhoff, Holling, and Schwabe (2009) discuss optimum designs for random effect regression models.

For fixed effect binary regression models, especially for logistic models, optimum designs have been extensively studied in the literature. See Abdelbasit and Plackett (1983), Minkin (1987) and Sitter and Fainaru (1997), among many others. Recently, Ouwens, Frans, and Martijn (2006) have studied optimum designs for logistic models with random intercepts. In all these papers the goal was to find (locally) optimum designs for estimating the unknown parameters of the models (or some functions of them). Sometimes, however, whether the model has random or fixed effects may be unknown. In this case the aim of the experiment should be dual, to choose between the rival models and to estimate the parameters of the

Chiara Tommasi

Department of Economics, Business and Statistics, University of Milan, Italy, e-mail: chiara.tommasi@unimi.it

M.T. Santos-Martín and J.M. Rodríguez-Díaz

Department of Statistics, University of Salamanca, Spain e-mail: maysam@usal.es, e-mail: juanmrod@usal.es

chosen model. In this paper optimum designs for discriminating between fixed- and random-effect univariate logistic regression models are computed. More specifically, two different optimality criteria are considered: the well-known D_s -criterion (see for instance Silvey 1980) and the KL -criterion, which is based on the Kullback-Leibler discrepancy between the rival models and was proposed by López-Fidalgo, Trandafir, and Tommasi (2007). Neither of the corresponding optimum designs can be found analytically, thus numerical procedures must be used. The main problem in constructing D_s - and KL -optimum designs is the computation of some integrals. If such integrals are computed numerically the numerical procedures become very slow, so that some approximations are needed. Theoretical details for deriving these approximations to integrals are given in the Appendix.

The outline of the paper is as follows: in Section 2 the models and the notation used throughout the work are introduced. In Sections 3 and 4, D_s - and KL -optimality criteria are specialized for discriminating between logistic regression models with or without random effects. Finally, in Section 5 an example is provided.

2 Logistic Regression Model

Let $\mathbf{b} = (b_0, b_1)'$ be the vector of the so called fixed effects, where “'” denotes transposition, and let $\beta = (\beta_0, \beta_1)'$ be the vector of random coefficients, which come from a Normal distribution with mean vector \mathbf{b} and dispersion matrix $\mathbf{V} = \text{diag}(v_0^2, v_1^2)$. The case of non-diagonal \mathbf{V} is somewhat more complex because it involves one parameter more, but it does not change substantially the theoretical results of this paper.

Let Y be a binary random response with success probability given by one of the following models:

1. The random effect logistic regression model

$$\begin{cases} P(Y = 1|\beta) = F(\beta_0 + \beta_1 x) = \frac{e^{\beta_0 + \beta_1 x}}{1 + e^{\beta_0 + \beta_1 x}} \\ \beta \sim N_2(\mathbf{b}, \mathbf{V}); \end{cases} \quad (1)$$

2. The common (fixed effect) logistic regression model

$$P(Y = 1; \mathbf{b}) = F(b_0 + b_1 x) = \frac{e^{b_0 + b_1 x}}{1 + e^{b_0 + b_1 x}}. \quad (2)$$

Let us denote model (1) by $f_1(y, x, \theta_1)$ and model (2) by $f_2(y, x, \theta_2)$ where $\theta_1 = (\mathbf{b}', \mathbf{v}')'$ and $\theta_2 = \mathbf{b}$ are vectors of unknown parameters. Here $\mathbf{v} = (v_0, v_1)'$ denotes the vector of standard deviations of β .

The main aim of this paper is to find optimum designs for discriminating between models (1) and (2). These two rival models are nested since the fixed effect model is a special case of the random effect model when $v_0 = v_1 = 0$. Two possible cri-

teria of optimality for handling this discrimination problem are the D_s -criterion for estimating \mathbf{v} ($s = 2$) and the KL -criterion. Both optimality criteria are motivated by the likelihood ratio test for the hypothesis $H_0 : v_0 = v_1 = 0$, because the corresponding optimum designs are expected to yield good power for testing this hypothesis (see for instance Dette and Titoff 2009 and López-Fidalgo, Trandafir, and Tommasi 2007).

3 D_s -Optimality Criterion

In order to construct a D_s -optimum design, the Fisher information matrix must be computed.

It is well known that for fixed effect models the Fisher information matrix coincides with the information matrix corresponding to the linearized binary regression model. This equivalence holds even if the coefficients are random variables (Gaussian or not). In order to have n independent observations, y_1, \dots, y_n , it is assumed that one observation per individual is taken, as in Graßhoff, Holling, and Schwabe (2009). With this assumption, the log-likelihood corresponding to $f_1(y, x, \theta_1)$ is

$$\log L(\theta_1) = \sum_{i=1}^n \log \int [F(\beta_0 + \beta_1 x_i)]^{y_i} [1 - F(\beta_0 + \beta_1 x_i)]^{1-y_i} \phi(\beta, \theta_1) d\beta,$$

where the integration is taken over \mathbb{R}^2 and $\phi(\beta, \theta_1)$ denotes the pdf of β . The Fisher information matrix is $I(\xi) = \{I_{kj}(\xi)\}$ where

$$I_{kj}(\xi) = E \left[-\frac{\partial^2}{\partial \theta_{1k} \partial \theta_{1j}} \log L(\theta_1) \right], \quad k, j = 1, \dots, 4.$$

The random-effect logistic regression model (1) can be rewritten as $Y_i = E(Y_i) + \varepsilon_i$ where $E(Y_i) = \int F(\beta_0 + \beta_1 x_i) \phi(\beta; \theta_1) d\beta$, $E(\varepsilon_i) = 0$ and $Var(\varepsilon_i) = \int F(\beta_0 + \beta_1 x_i) \phi(\beta; \theta_1) d\beta \cdot [1 - \int F(\beta_0 + \beta_1 x_i) \phi(\beta; \theta_1) d\beta]$. If this model is linearized at some nominal values of the parameters then the corresponding information matrix is $M(\xi) = \int_{\mathcal{X}} \mathbf{g}(x) \mathbf{g}(x)' d\xi(x)$, where $\mathbf{g}(x) = [g_1(x), g_2(x), g_3(x), g_4(x)]'$ and

$$g_j(x) = \frac{\int F(\beta_0 + \beta_1 x_i) \frac{\partial}{\partial \theta_{1j}} \phi(\beta; \theta_1) d\beta}{\sqrt{\int F(\beta_0 + \beta_1 x_i) \phi(\beta; \theta_1) d\beta \cdot [1 - \int F(\beta_0 + \beta_1 x_i) \phi(\beta; \theta_1) d\beta]}}.$$

After some algebra it can be proved that $M(\xi)$ coincides with the Fisher information matrix $I(\xi)$.

For the logistic regression model with Gaussian random effects,

$$g_1(x) = \frac{I_1(x)}{v_0 \sqrt{I_0(x)[1 - I_0(x)]}}, \quad g_2(x) = \frac{I_2(x)}{v_1 \sqrt{I_0(x)[1 - I_0(x)]}},$$

$$g_3(x) = \frac{I_3(x) - I_0(x)}{v_0 \sqrt{I_0(x)[1 - I_0(x)]}} \quad \text{and} \quad g_4(x) = \frac{I_4(x) - I_0(x)}{v_1 \sqrt{I_0(x)[1 - I_0(x)]}},$$

where

$$I_0(x) = \frac{1}{2\pi} \int e^{h_1(\tilde{\beta};x)} d\tilde{\beta}, \quad I_1(x) = \frac{1}{2\pi} \int \tilde{\beta}_0 e^{h_1(\tilde{\beta};x)} d\tilde{\beta}, \quad I_2(x) = \frac{1}{2\pi} \int \tilde{\beta}_1 e^{h_1(\tilde{\beta};x)} d\tilde{\beta},$$

$$I_3(x) = \frac{1}{2\pi} \int \tilde{\beta}_0^2 e^{h_1(\tilde{\beta};x)} d\tilde{\beta}, \quad I_4(x) = \frac{1}{2\pi} \int \tilde{\beta}_1^2 e^{h_1(\tilde{\beta};x)} d\tilde{\beta},$$

$$h_1(\tilde{\beta};x) = b_0 + v_0 \tilde{\beta}_0 + b_1 x + v_1 \tilde{\beta}_1 x - \frac{1}{2} \tilde{\beta}_0^2 - \frac{1}{2} \tilde{\beta}_1^2 - \log \left[1 + \exp \left(b_0 + v_0 \tilde{\beta}_0 + b_1 x + v_1 \tilde{\beta}_1 x \right) \right],$$

$$\tilde{\beta}_0 = (\beta_0 - b_0)/v_0 \quad \text{and} \quad \tilde{\beta}_1 = (\beta_1 - b_1)/v_1.$$

If $\mathbf{g}_b(x) = [g_1(x), g_2(x)]'$ and $\mathbf{g}_v(x) = [g_3(x), g_4(x)]'$, then the Fisher information matrix is

$$M(\xi) = \int_{\mathcal{X}} \mathbf{g}(x) \mathbf{g}(x)' d\xi(x) = \begin{bmatrix} M_b(\xi) & M_{b,v}(\xi) \\ M_{v,b}(\xi) & M_v(\xi) \end{bmatrix}, \quad (3)$$

where $M_b(\xi) = \int_{\mathcal{X}} \mathbf{g}_b(x) \mathbf{g}_b(x)' d\xi(x)$, $M_{b,v}(\xi) = \int_{\mathcal{X}} \mathbf{g}_b(x) \mathbf{g}_v(x)' d\xi(x)$, $M_{v,b}(\xi) = M_{b,v}'(\xi)$ and $M_v(\xi) = \int_{\mathcal{X}} \mathbf{g}_v(x) \mathbf{g}_v(x)' d\xi(x)$.

Let the inverse of the Fisher information matrix be

$$M^{-1}(\xi) = \begin{bmatrix} M^{b,b}(\xi) & M^{b,v}(\xi) \\ M^{v,b}(\xi) & M^{v,v}(\xi) \end{bmatrix},$$

when the asymptotic covariance matrix of the maximum likelihood estimator of \mathbf{v} is

$$M^v(\xi) = [M_v(\xi) - M_{v,b}(\xi) M_b^{-1}(\xi) M_{b,v}(\xi)]^{-1}.$$

A D_s -optimum design minimizes the determinant of $M^v(\xi)$ (or its logarithm), or equivalently maximizes the criterion function

$$\Psi_{D_s}(\xi) = \log \frac{|M(\xi)|}{|M_b(\xi)|},$$

since $|M_v(\xi) - M_{v,b}(\xi) M_b^{-1}(\xi) M_{b,v}(\xi)| = |M(\xi)|/|M_b(\xi)|$.

It is well known that $\xi_{D_s}^*$ is a D_s -optimum design if and only if it fulfills the inequality

$$\mathbf{g}(x)' M^{-1}(\xi_{D_s}^*) \mathbf{g}(x) - \mathbf{g}_b(x)' M_b^{-1}(\xi_{D_s}^*) \mathbf{g}_b(x) - s \leq 0, \quad x \in \mathcal{X}, \quad (4)$$

where $s = 2$ as \mathbf{v} is bidimensional. This inequality is also useful for constructing D_s -optimum designs through the first order algorithm, since the left-hand side of (4) is the directional derivative of $\Psi_{D_s}(\xi)$ at $\xi_{D_s}^*$ in the direction of $\xi_x - \xi_{D_s}^*$, where ξ_x is the design which concentrates all weight at the point x . This algorithm converges very slowly if the integrals $I_j(x)$, $j = 0, \dots, 4$, are evaluated numerically. In order to improve the convergence rate, the approximations to integrals (15) - (19) may be applied, since $h_1(\tilde{\beta};x)$ is a concave function of $\tilde{\beta}$ for any $x \in \mathcal{X}$.

4 KL-Optimality Criterion

Let us recall that the Kullback-Leibler discrepancy between the rival models $f_1(y, x, \theta_1)$ and $f_2(y, x, \theta_2)$, assuming $f_1(y, x, \theta_1)$ as the “true” completely known model, is

$$\begin{aligned} \mathcal{J}[f_1(y, x, \theta_1), f_2(y, x, \theta_2)] &= \int f_1(y, x, \theta_1) \log \frac{f_1(y, x, \theta_1)}{f_2(y, x, \theta_2)} dy \tag{5} \\ &= \left[\log \frac{I_0(x)}{1 - I_0(x)} - b_0 - b_1 x \right] I_0(x) + \log[1 - I_0(x)] + \log \left(1 + e^{b_0 + b_1 x} \right). \end{aligned}$$

The corresponding *KL*-criterion function is

$$I_{2,1}(\xi) = \min_{\theta_2 \in \Omega_2} \int_{\mathcal{X}} \mathcal{J}[f_1(y, x, \theta_1), f_2(y, x, \theta_2)] \xi(dx).$$

A design ξ_{21}^* which maximizes $I_{2,1}(\xi)$ is a *KL*-optimum design. Exchanging the role of $f_1(y, x, \theta_1)$ and $f_2(y, x, \theta_2)$ in (5) the Kullback-Leibler distance between $f_2(y, x, \theta_2)$ and $f_1(y, x, \theta_1)$ is $\mathcal{J}[f_2(y, x, \theta_2), f_1(y, x, \theta_1)]$ and the corresponding *KL*-optimality criterion $I_{1,2}(\xi)$ may be defined. In this paper, however, $I_{1,2}(\xi) = 0$ since the fixed effect model is nested within the random-effect model and thus only $I_{2,1}(\xi)$ can be computed. For this reason the random effect model has been considered as “true”. López-Fidalgo, Trandafir, and Tommasi (2007) proved that a design ξ_{21}^* is a *KL*-optimum design if and only if $\psi(x, \xi_{21}^*) \leq 0$ for any $x \in \mathcal{X}$, where

$$\psi(x, \xi) = \mathcal{J}[f_1(y, x, \theta_1), f_2(y, x, \hat{\theta}_2)] - \int_{\mathcal{X}} \mathcal{J}[f_1(y, x, \theta_1), f_2(y, x, \hat{\theta}_2)] d\xi(x) \tag{6}$$

is the directional derivative of $I_{2,1}(\xi)$ at ξ in the direction of $\xi_x - \xi$ and $\hat{\theta}_2$ is the assumed unique element of

$$\Omega_2(\xi) = \left\{ \hat{\theta}_2 : \hat{\theta}_2(\xi) = \arg \min_{\theta_2 \in \Omega_2} \int_{\mathcal{X}} \mathcal{J}[f_1(y, x, \theta_1), f_2(y, x, \theta_2)] d\xi(x) \right\}.$$

In order to compute numerically a *KL*-optimum design (through the first-order algorithm) the Kullback-Leibler discrepancy plays a fundamental role since it appears in the expression of the directional derivative (6). Again, in order to improve the convergence speed of the first-order algorithm, Laplacian approximation (15) has been used for the integral $I_0(x)$.

5 Some Results

For the nominal values of the parameters $b_0 = 1, b_1 = 2, v_0 = 0.2$ and $v_1 = 0.3$, the D_S - and *KL*-optimum designs are

$$\xi_{D_s}^* = \left\{ \begin{array}{cccc} 0 & 0.245 & 0.637 & 1 \\ 0.181 & 0.295 & 0.286 & 0.238 \end{array} \right\} \quad \text{and} \quad \xi_{21}^* = \left\{ \begin{array}{ccc} 0 & 0.433 & 1 \\ 0.207 & 0.471 & 0.322 \end{array} \right\},$$

respectively. The *KL*-optimum design has only three support points, while the D_s -optimum design is supported at four points. However, since the Fisher information matrix of $\xi_{D_s}^*$ has an eigenvalue close to zero, both $\xi_{D_s}^*$ and ξ_{21}^* enable us to efficiently estimate only the parameters of the common logistic regression model. The efficiency of the *KL*-optimum design with respect to $\xi_{D_s}^*$ cannot be computed because ξ_{21}^* is not a regular design, i.e. $M(\xi_{21}^*)$ is a singular matrix. The efficiency of the D_s -optimum design with respect to ξ_{21}^* is $\text{Eff}_{21}(\xi_{D_s}^*) = I_{2,1}(\xi_{D_s}^*)/I_{2,1}(\xi_{21}^*) = 0.659$. A more detailed study of efficiencies will be developed in our future work, taking into consideration also different nominal values for the parameters.

From the computational point of view, the construction of *KL*-optimal designs is less time-demanding than that of D_s -optimal designs. The reason may be the different number of integrals involved in the expressions for directional derivatives (4) and (6).

6 Appendix

Theorem 1. *Let $p(x, y) = a_0x^2 + a_1y^2 + a_2xy + a_3x + a_4y$ such that $(x, y) \in \mathbb{R}^2$, $a_0 < 0$, $a_1 < 0$, $a_2, a_3, a_4 \in \mathbb{R}$ and $4a_0a_1 - a_2^2 > 0$, then*

$$\int \int e^{p(x,y)} dx dy = K, \tag{7}$$

where

$$K = \frac{2\pi e^{-\frac{a_0a_2^2+a_1a_3^2-a_2a_3a_4}{4a_0a_1-a_2^2}}}{\sqrt{4a_0a_1 - a_2^2}}. \tag{8}$$

Proof.

Let (x, y) be a bivariate Gaussian random vector such that $E(x) = \mu_x$, $E(y) = \mu_y$, $\text{Var}(x) = \sigma_x^2$, $\text{Var}(y) = \sigma_y^2$ and $\text{Cov}(x, y) = \sigma_{xy}$ and let $f(x, y)$ denote the pdf of (x, y) . The analytical solution (7) follows from the equality $\int \int f(x, y) dx dy = 1$.

Corollary 1. *Let $h(x, y)$ be a concave function of $(x, y) \in \mathbb{R}^2$ and let $h'_x, h'_y, h''_{xx}, h''_{yy}$ and h''_{xy} denote the first- and the second-order partial derivatives of $h(x, y)$ evaluated at the point (x_0, y_0) . Then*

$$\int \int e^{h(x,y)} dx dy \approx \int \int e^{p(x,y)} dx dy = k_0 K, \tag{9}$$

where

$$k_0 = \exp \left[h(x_0, y_0) - x_0 h'_x - y_0 h'_y + \frac{1}{2} x_0^2 h''_{xx} + \frac{1}{2} y_0^2 h''_{yy} + x_0 y_0 h''_{xy} \right] \tag{10}$$

and K is given by equation (8) with $a_0 = h''_{xx}/2$, $a_1 = h''_{yy}/2$, $a_2 = h''_{xy}$, $a_3 = h'_x - x_0 h''_{xx} - y_0 h''_{xy}$ and $a_4 = h'_y - y_0 h''_{yy} - x_0 h''_{xy}$.

Proof

This corollary follows from Theorem 1 by replacing the function $h(x, y)$ with its second-order Taylor series expansion at the point (x_0, y_0) in the left-hand side of equation (9).

Remark 1. Using the notation of Theorem 1, from the identities $\iint x f(x, y) dx dy = \mu_x$, $\iint y f(x, y) dx dy = \mu_y$, $\iint x^2 f(x, y) dx dy = \sigma_x^2 + \mu_x^2$ and $\iint y^2 f(x, y) dx dy = \sigma_y^2 + \mu_y^2$ it follows that

$$\iint x e^{p(x,y)} dx dy = -K \frac{2a_1 a_3 - a_2 a_4}{4a_0 a_1 - a_2^2}, \tag{11}$$

$$\iint y e^{p(x,y)} dx dy = -K \frac{2a_0 a_4 - a_2 a_3}{4a_0 a_1 - a_2^2}, \tag{12}$$

$$\iint x^2 e^{p(x,y)} dx dy = K \frac{(2a_1 a_3 - a_2 a_4)^2 - 2a_1(4a_0 a_1 - a_2^2)}{(4a_0 a_1 - a_2^2)^2} \tag{13}$$

and

$$\iint y^2 e^{p(x,y)} dx dy = K \frac{(2a_0 a_4 - a_2 a_3)^2 - 2a_0(4a_0 a_1 - a_2^2)}{(4a_0 a_1 - a_2^2)^2}, \tag{14}$$

with K given in (8). Under the same hypothesis as in Corollary 1, approximations for $\iint x e^{h(x,y)} dx dy$, $\iint y e^{h(x,y)} dx dy$, $\iint x^2 e^{h(x,y)} dx dy$ and $\iint y^2 e^{h(x,y)} dx dy$ are given on multiplying the right-hand side of each equation (11)-(14) by k_0 given in (10).

Remark 2. If the function $h(x, y)$ is approximated around its maximum point, i.e. (x_{\max}, y_{\max}) , then (9) becomes

$$\iint e^{h(x,y)} dx dy \approx 2\pi \tilde{I}_0(x), \tag{15}$$

where $\tilde{I}_0(x) = e^{h(x_{\max}, y_{\max})} / \sqrt{H}$ and $H = h''_{xx} h''_{yy} - h''_{xy}{}^2$ is the determinant of the Hessian matrix of $h(x, y)$ evaluated at (x_{\max}, y_{\max}) . The integral approximation (15) is called a Laplacian approximation and is more accurate for small values of v_0 and v_1 (e.g. values less than 1). See for instance Demidenko (2004), pp. 340 and 400.

In addition, it can be proved that

$$\iint x e^{h(x,y)} dx dy \approx x_{\max} 2\pi \tilde{I}_0(x), \tag{16}$$

$$\iint y e^{h(x,y)} dx dy \approx y_{\max} 2\pi \tilde{I}_0(x) \tag{17}$$

$$\iint x^2 e^{h(x,y)} dx dy \approx \left(x_{\max}^2 - \frac{h''_{yy}}{H} \right) 2\pi \tilde{I}_0(x), \tag{18}$$

and

$$\iint y^2 e^{h(x,y)} dx dy \approx \left(y_{\max}^2 - \frac{h''_{xx}}{H} \right) 2\pi \tilde{I}_0(x). \quad (19)$$

Acknowledgements This research was supported by the Spanish Junta de Castilla y León (Project ‘SA071A09’) and the Spanish Ministry of Education and Science (Projects ‘MTM 2007-672111-C03-01’ and ‘Ingenio Mathematica (i-MATH)’ No. CSD2006-00032, Consolider - Ingenio 2010).

References

- Abdelbasit, K. and R. Plackett (1983). Experimental design for binary data. *Journal of the American Statistical Association* 78, 90–98.
- Demidenko, E. (2004). *Mixed Models: Theory and Applications*. New York: Wiley.
- Dette, H. and S. Titoff (2009). Optimal discrimination designs. *Annals of Statistics* 37, 2056–2082.
- Graßhoff, U., H. Holling, and R. Schwabe (2009). On optimal design for a heteroscedastic model arising from random coefficients. In *Proceedings of the 6th St. Petersburg Workshop on Simulation*, St. Petersburg, pp. 387–392. VVM com. Ltd.
- López-Fidalgo, J., C. Trandafir, and C. Tommasi (2007). An optimal experimental design criterion for discriminating between non-normal models. *Journal of the Royal Statistical Society, Series B* 69, 231–242.
- Mentré, F., A. Mallet, and D. Baccar (1997). Optimal design in random-effects regression models. *Biometrika* 84, 429–442.
- Minkin, S. (1987). Optimal designs for binary data. *Journal of the American Statistical Association* 82, 1098–1103.
- Ouwens, M., T. Frans, and B. Martijn (2006). A maximin criterion for the logistic random intercept model with covariates. *Journal of Statistical Planning and Inference* 136, 962–981.
- Patan, M. and B. Bogacka (2007). Efficient sampling windows for parameter estimation in mixed effects models. In *mODa 8 - Advances in Model-Oriented Design and Analysis*, New York, pp. 147–155. Physica-Verlag.
- Silvey, S. (1980). *Optimal Design*. New York: Chapman & Hall.
- Sitter, R. and I. Fainaru (1997). Optimal designs for the logit and probit models for binary data. *Canadian Journal of Statistics* 25, 175–190.

Estimation and Optimal Designing under Latent Variable Models for Paired Comparisons Studies via a Multiplicative Algorithm

Bernard Torsney

Abstract We consider:

1. The problem of estimating the parameters of latent variable models such as the Bradley Terry or Thurstone Model by the method of maximum likelihood, given data from a paired comparisons experiment. The parameters of these models can be taken to be weights which are positive and sum to one;
2. The problem of determining approximate locally optimal designs for good estimation of these parameters; i.e of determining optimal design weights which are also positive and sum to one.

1 Paired Comparisons

1.1 Introduction

We have two alternative examples of a general problem, namely determining weights optimally. Much theory for this problem, e.g. optimality conditions and numerical techniques have been developed in the optimal design arena. So this can be transported to the estimation problem. We can extend techniques to this case. In section 1 we introduce the notion of paired comparisons studies and latent variable models. In section 2 the parameter estimation problem is outlined with optimality results and a general class of multiplicative algorithms outlined in sections 3 and 4 respectively. A specific algorithm is applied to the Bradley Terry log-likelihood in section 5 and locally optimal designing is considered in section 6.

We consider paired comparison experiments in which J treatments or products are compared in pairs. In a simple form a subject is presented with two treatments

Bernard Torsney

Department of Statistics, University of Glasgow, Glasgow G12 8QW,UK e-mail: ben@stats.gla.ac.uk

and asked to indicate which he/she prefers or considers better. In reality the subject will be an expert tester; for example, a food taster in examples arising in food technology. The link with optimal design theory (apart from the fact that a specialised design, paired comparisons, is under consideration) is that, the parameters of latent variable models for the resultant data are like weights. Hence the theory characterising and the methods developed for finding optimal design weights can be applied to characterising and finding the maximum likelihood estimators of these latent variable 'weights'.

1.2 The Data

In a simple experiment a set of such testers is available and each is presented with one pair from a set of J treatments, say T_1, T_2, \dots, T_J . The number of comparisons, n_{ij} of T_i to T_j , we assume has been predetermined. Sufficient summary data comprises the set $\{O_{ij} : i = 1, \dots, J; j = 1, \dots, J; i < j \text{ or } i > j\}$, where O_{ij} is the observed frequency with which T_i is preferred to T_j . Of course $O_{ij} + O_{ji} = n_{ij}$

Bradley and El-Helbawy (1976) introduce an example involving 8 coffee types. 26 pairwise comparisons were made on each pair, i.e. $n_{ij} = 26$. So $O_{ij} + O_{ji} = 26$ and $N = \sum_i \sum_j O_{ij} = 728$.

The coffees are the eight combinations arising from a 2^3 factorial structure, the factors being Brew Strength, Roast Colour, Coffee Brand. We are not exploiting this structure and leave them arbitrarily labelled.

1.3 Models

1.3.1 A General Model

In the absence of other information the most general model here is to propose

$$O_{ij} \sim Bi(n_{ij}, \theta_{ij})$$

where,

$$\theta_{ij} = P(T_i \text{ is preferred to } T_j).$$

Apart from the constraint $O_{ij} + O_{ji} = n_{ij}$, independence between frequencies is an expected assumption. So, apart from the constraint $\theta_{ij} + \theta_{ji} = 1$, these define unrelated binomial parameters. The maximum likelihood estimator of θ_{ij} is O_{ij}/n_{ij} (the proportion of times T_i is preferred to T_j in these n_{ij} comparisons), and formal inferences can be based on the asymptotic properties of these.

1.3.2 Latent Variable Models

These are more restricted models in that they impose interrelations between the θ_{ij} . Assuming that $F(\cdot)$ is a symmetric distribution function, then

$$\theta_{ij} = F(\lambda_i - \lambda_j) = F\{\log_e(p_i/p_j)\}$$

where $p_i = \exp(\lambda_i)$. The symmetry of $F(\cdot)$ ensures that

$$\theta_{ij} + \theta_{ji} = F(\lambda_i - \lambda_j) + F(\lambda_j - \lambda_i) = 1.$$

The p_i or λ_i can be viewed as indices or quality characteristics, one for each treatment. The implication of the model is that the difference in quality between two treatments has distribution function $F(\cdot)$.

Two primary examples of this model are the Bradley Terry and Thurstone models. Respectively these take $F(\cdot)$ to be the Logistic and the Normal distributions. In the Logistic case θ_{ij} has the simplistic form: $\theta_{ij} = p_i/(p_i + p_j)$; see Thurstone (1927), Bradley and Terry (1952), also Kuk (1995).

2 Parameter Estimation

The likelihood of the data is

$$L = \prod_{r < s} \prod [F\{\log_e(p_r/p_s)\}]^{O_{rs}} [F\{\log_e(p_s/p_r)\}]^{O_{sr}}.$$

We focus on the parameters p_i and denote the likelihood by $L(p)$.

However we cannot estimate these as free parameters. This arises from the fact that we only have observations on comparisons between treatments, and is reflected in the property that θ_{ij} is invariant to proportional changes in p_i and p_j . In consequence the p_i are only unique up to a constant multiple; (likewise the λ_i up to a constant shift). In keeping with this they are positive as the relationship $p_i = \exp(\lambda_i)$ implies $p_i > 0$. Mathematically speaking θ_{ij} and hence $L(p)$ is a homogeneous function of degree zero in the p_i i.e. $L(cp) = L(p)$, where c is a scalar constant. So $L(p)$ is constant on rays running out from the origin. It will therefore be maximised along one specific ray. We can identify this ray by finding a particular optimising p^* . This we can do by imposing a constraint on p . Possible constraints are $\sum_i p_i = 1$ or $\prod_i p_i = 1$, or $g(p) = 1$ where $g(p)$ is a surface which cuts each ray exactly once. In the case $J = 2$ a suitable $g(p)$ is defined by $p_2 = h(p_1)$, where $h(\cdot)$ is a decreasing function which cuts the two main axes, as in the case of $h(p_1) = 1 - p_1$, or has these as asymptotes, as in the case of $h(p_1) = 1/p_1$. In general a suitable choice of $g(p)$ is one which is positive and homogeneous of some degree h . Note that other alternatives are $\sum_i p_i = C$ or $\prod_i p_i = C$, where C is any positive constant; e.g. $C = J$ or $C = 100$.

The choice of $\prod_i p_i = 1$, being equivalent to $\sum_i \ln(p_i) = 0$, confers on $\lambda_i = \ln(p_i)$ the notion of a main effect. However we will opt for the choice of $\sum_i p_i = 1$, which conveys the notion of p_i as a weight. We wish to maximise the likelihood or log-likelihood subject to this constraint and to non-negativity too. This is an example of the following general problem:

Problem (\mathfrak{P})

Maximise $\phi(p)$ subject to $p_i \geq 0$, $\sum_i p_i = 1$.

We wish to maximise $\phi(p)$ with respect to a probability distribution.

For the estimation problem we will take $\phi(p) = \ln\{L(p)\}$.

There are many examples of this problem arising in various areas of statistics, especially in the area of optimal regression design. We can exploit optimality results and algorithms developed in this area. The feasible region is an open but bounded set. Thus there should always be a solution to this problem allowing for the possibility of an unbounded maximum, multiple solutions and solutions at vertices (i.e. $p_t = 1, p_i = 0, i \neq t$).

3 Optimality Conditions

We assume that $\phi(\cdot)$ is differentiable. Let

$$F_j = d_j - p^T d = d_j - \sum_i p_i d_i, \text{ where } d_j = \partial\phi/\partial p_j.$$

We call F_j the j th vertex directional derivative of $\phi(\cdot)$ at p .

Note that $\sum_j p_j F_j = 0$, so that, in general, some F_j are negative and some are positive.

Given $\phi(\cdot)$ is differentiable at p^* , then a necessary condition for $\phi(p^*)$ to be a local maximum of $\phi(\cdot)$ in the feasible region of Problem (\mathfrak{P}) is

$$F_j^* = 0 \text{ for } p_j^* > 0,$$

$$F_j^* \geq 0 \text{ for } p_j^* = 0.$$

If $\phi(\cdot)$ is concave on its feasible region, then these first order stationarity conditions are both necessary and sufficient. This is the general equivalence theorem in optimal design. See Whittle (1973), Kiefer (1974). In fact the second condition is redundant for this estimation problem, while, given homogeneity of degree zero of $L(p)$, the first reduces to standard first order conditions: $d_j^* = 0$.

4 Algorithms

4.1 Multiplicative Algorithm

Problem (\mathfrak{P}) has a distinct set of constraints, namely the variables p_1, p_2, \dots, p_J must be nonnegative and sum to 1. Let $f(d, \delta)$ be a function satisfying (for $\delta > 0$):

- $f(d, \delta) > 0$,
- $\frac{\partial f(d, \delta)}{\partial d} > 0$ (for $\delta > 0$),
- $f(d, 0) = \text{constant}$

(e.g. $f(d, \delta) = \Phi(\delta d)$ or $f(d, \delta) = d^\delta$ (if $d > 0$.)

An iteration which neatly submits to these and has some suitable properties is the multiplicative algorithm:

$$p_j^{(r+1)} = \frac{p_j^{(r)} f(d_j^{(r)})}{\sum_i p_i^{(r)} f(d_i^{(r)})}$$

where $d_j^{(r)} = \left. \frac{\partial \phi}{\partial p_j} \right|_p = p^{(r)}$, while $f(d)$ is positive and strictly increasing in d and may depend on one or more free parameters.

4.2 Properties of the Algorithm

Under the conditions imposed on $f(\cdot, \cdot)$, the above iterations possess the following properties which are considered in more detail in Torsney (1988), Torsney and Alahmadi (1992) and Mandal and Torsney (2000):

1. $p^{(r)}$ is always feasible.
2. $F_\phi\{p^{(r)}, p^{(r+1)}\} \geq 0$, with equality when the d_j 's corresponding to nonzero p_j 's have a common value $d (= \sum_i p_i d_i)$, in which case $p^{(r)} = p^{(r+1)}$.
So an iterate $p^{(r)}$ is a fixed point of the iteration if derivatives $d_j^{(r)}$ corresponding to nonzero $p_j^{(r)}$ are equal; i.e. if corresponding vertex directional derivatives $F_j^{(r)}$ are zero.
3. If $\delta = 0$ there is no change in $p^{(r)}$, given $f(d, \delta) = \text{constant}$
4. So the algorithm should be monotonic for small positive δ .

5 Fitting Bradley-Terry Models

Our criterion is

$$\phi(p) = \ln\{L(p)\}$$

Since $L(p)$ is a homogeneous function of degree zero $\sum_i p_i d_i = 0$. In fact $d_j = F_j$. So there are always positive and negative d_j unless all are zero. We require a function $f(d, \delta)$ which is defined for positive and negative d , where we take d to represent a partial derivative. Noting that all p_j^* must be positive a suitable choice should be governed by the fact that at the optimum $d_j^* = 0, j = 1, 2, \dots, J$.

We opt for $f(d, \delta) = \Phi(\delta d)$, so that iterations prove to be

$$p_j^{(r+1)} = \frac{p_j^{(r)} \Phi(\delta d_j^{(r)})}{\sum_i p_i^{(r)} \Phi(\delta d_i^{(r)})}$$

Coffee Example.

In this case $J = 8$ coffee types were compared yielding a total of $N = 728$ observations; i.e. $\sum \sum O_{ij} = 728$. A suitable δ is $\delta = 1/N$. In effect we are standardising the sample size to 1, through replacing observed by relative frequencies in the log-likelihood, and then taking $\delta = 1$.

Torsney (2004) reported the following results. Starting from $p_j^{(0)} = 1/J$, the numbers of iterations needed to achieve $\max |d_j| = \max |F_j| \leq 10^{-n}$, for $n = 0, 1, \dots, 7$ respectively are 17, 21, 25, 32, 38, 45, 51, 59. The optimal p^* is (0.190257, 0.122731, 0.155456, 0.106993, 0.091339, 0.149406, 0.080953, 0.102865).

Iterations were monotonic.

6 Local Optimal Designing

We have not introduced any design variables. However we can pose the question: how many comparisons n_{ij} there should be between T_i and T_j ? This of course is an exact design problem. The easier approximate design problem poses the question: what proportion λ_{ij} of such comparisons there should be?

This depends on our model. We focus on the Bradley Terry Model. The parameters are now p_1, p_2, \dots, p_J . We wish good estimation of these. The information matrix is

$$M(\lambda) = \sum \sum_{i < j} \lambda_{ij} w_{ij} v_{ij} v_{ij}^T$$

where $v_{ij} = (e_i - e_j), e_i$ being the i th unit vector $w_{ij} = 1/(p_i + p_j)^2$.

We note the following properties:

1. $M(\lambda)$ has the form of the information matrix of a weighted linear model with weights w_{ij} . This happens with a wide range of generalised linear models.
2. $M(\lambda)$ depends on the p_i 's (but only through the w_{ij} 's).

We need provisional values for them. A conventional choice is $p_j = 1/J$.

However we have maximum likelihood estimates. This does not seem to have been considered in the literature before.

3. $M(\lambda)$ is singular. This is another manifestation of the fact that we only have observations on comparisons between treatments. We can only estimate differences between treatments. This has implications for choice of design criteria. We must restrict consideration to good estimation of such differences (or other contrasts). This issue too appears to have been ignored in the literature.

Two feasible classes are:

$$D_L - \text{criteria} : \Psi(M) = -\log \det(LM^+L^T),$$

$$A_L - \text{criteria} : \Psi(M) = -\text{trace}(LM^+L^T).$$

Here M^+ denotes the Moore-Penrose inverse of M and L defines a set of $(k - 1)$ linearly independent differences between the p_i parameters.

The D_L -criterion would be invariant to any such choice of L .

In general a locally optimal design problem is, for given p , to choose λ optimally subject to $\lambda_{ij} \geq 0, \sum \lambda_{i < j} = 1$, i.e. solve Problem (P) for $\phi(\lambda) = \Psi\{M(\lambda)\}$ for some $\Psi\{\cdot\}$.

We need derivatives with respect to λ_{ij} , which we denote by d_{ij} , for optimality checking and numerical purposes. We have:

$$\begin{aligned} \text{for the } D_L\text{-criterion,} \quad & d_{ij} = w_{ij}v_{ij}^T M^+ L^T (LM^+L^T)^{-1} LM^+ v_{ij} \\ \text{for the } A_L\text{-criterion,} \quad & d_{ij} = w_{ij}v_{ij}^T M^+ L^T LM^+ v_{ij}. \end{aligned}$$

Of note is that these are positive, as is the case with all standard design criteria. For the multiplicative algorithm a feasible choice is $f(d, \delta) = d^\delta$, the original form of this function when the algorithm was first conceived for determining optimal designs. The choices of δ we opt for here correspond to choices which have been shown to be monotonic for the standard D -criterion and A -criterion, namely $\delta = 1, 1/2$ respectively.

Coffee example

We choose to determine locally optimal designs at the current maximum likelihood estimates; i.e. at $p^* = (0.190257, 0.122731, 0.155456, 0.106993, 0.091339, 0.149406, 0.080953, 0.102865)$. We use the following choices of $f(d, \delta)$: for the D_L -criterion: $f(d, \delta) = d$; for the A_L -criterion: $f(d, \delta) = d^{1/2}$.

Iterations begin at $\lambda_{ij}^{(0)} = 1/(J(J - 1))$.

We take L to be the matrix defining the 7 differences $p_1 - p_j, j = 2, 3, \dots, 8$.

We summarise the implications if a further experiment is to be run and parameter values are in the region of the maximum likelihood estimates: for D_L -optimality no comparisons would be made between coffee types 1 and 3 and between coffee types 1 and 6; under both designs maximum weight is put on the comparisons between

coffee types 1 and 7, which have the largest and smallest estimated Bradley Terry parameters; the A_L -optimal weights of the 7 comparisons with the first coffee type exceed 0.07 while the remainder are less than 0.03, which is in keeping with the focus of the choice of L on differences with this coffee type.

For comparison we note that uniform weights of $1/28 = 0.0357143$.

7 Discussion

There are several extensions of this work in respect of both parameter estimation and local optimal designing (arguably new): for rankings; for “no preference” options; for factorially structured treatments.

References

- Bradley, R. A. and Terry, M. E. (1952). The rank analysis of incomplete block designs I, The method of paired comparisons. *Biometrika* 39, 324-345.
- Bradley, R. A. and El-Helbawy, A. T. (1976). Treatment contrasts in paired comparisons: basic procedures with applications to factorials. *Biometrika* 63, 255-262.
- Kiefer, J. (1974) General equivalence theory for optimum designs (approximate theory) *Annals of Statistics* 2, 849-879.
- Kuk, A.C. Y. (1995). Modelling paired comparison data with large numbers of draws and large variability of draw percentages among players. *Statistician* 44, 523-528.
- Mandal, S. and Torsney, B. (2000) Algorithms for the construction of optimising distributions. *Communications in Statistics (Theory and Methods)* 29, 1219-1231.
- Thurstone, L. L. (1927b). The method of paired comparisons for social values. *Journal of Abnormal and Social Psychology* 21, 384-400.
- Torsney, B. (1988) Computing Optimizing Distributions with Applications in Design, Estimation and Image Processing. In *Optimal Design and Analysis of Experiments* (Y. Dodge, V.V. Fedorov and H.P. Wynn eds). Amsterdam: North Holland, 361-370.
- Torsney, B. (2004) Fitting Bradley Terry models using a multiplicative algorithm. In *Proceedings in Computational Statistics (COMPSTAT2004)* (J. Antoch, ed.). Heidelberg: Physica-Verlag, 214-226.
- Torsney, B. and Alahmadi, A. M. (1992) Further developments of algorithms for constructing optimizing distributions. In *Model Oriented Data Analysis (Proceedings of 2nd IIASA- Workshop, St. Kyrik, Bulgaria, 1990)* (V. Fedorov, W.G. Muller, I.N. Vuchkov eds). Heidelberg: Physica-Verlag, 121-129.
- Whittle, P. (1973). Some general points in the theory of optimal experimental design *Journal of the Royal Statistical Society B* 35, 123-130.

Pointwise Consistency of the Kriging Predictor with Known Mean and Covariance Functions

Emmanuel Vazquez and Julien Bect

Abstract This paper deals with several issues related to the pointwise consistency of the kriging predictor when the mean and the covariance functions are known. These questions are of general importance in the context of computer experiments. The analysis is based on the properties of approximations in reproducing kernel Hilbert spaces. We fix an erroneous claim of Yakowitz and Szidarovszky (J. Multivariate Analysis, 1985) that, under some assumptions, the kriging predictor is pointwise consistent for all continuous sample paths.

1 Introduction

The domain of *computer experiments* is concerned with making inferences about the output of an expensive-to-run numerical simulation of some physical system, which depends on a vector of factors with values in $\mathbb{X} \subseteq \mathbb{R}^d$. The output of the simulator is formally an unknown function $f : \mathbb{X} \rightarrow \mathbb{R}$. For example, to comply with ever-increasing standards regarding pollutant emissions, numerical simulations are used to determine the level of emissions of a combustion engine as a function of its design parameters (Villemonteix 2008). The emission of pollutants by an engine involves coupled physical phenomena whose numerical simulation by a finite-element method, for a fixed set of design parameters of the engine, can take several hours on high-end servers. It then becomes very helpful to collect the answers already provided by the expensive simulator, and to construct from them a simpler computer model, that will provide approximate but cheaper answers about a quantity of interest. This approximate model is often called a *surrogate*, or a *metamodel*, or an *emulator* of the actual simulator f . The quality of the answers given by the approximate model depends on the quality of the approximation, which depends, in turn and in

Emmanuel Vazquez · Julien Bect
SUPELEC, 3 rue Joliot-Curie, 91192 Gif-sur-Yvette, France
e-mail: emmanuel.vazquez@supelec.fr and julien.bect@supelec.fr

part, on the choice of the evaluation points of f , also called *experiments*. The choice of the evaluation points is usually called the *design of experiments*. Assuming that f is continuous, it is an important question to know whether the approximate model behaves consistently, in the sense that if the evaluation points x_n are chosen sequentially in such a way that a given point $x \in \mathbb{X}$ is an accumulation point of $\{x_n, n \geq 1\}$, then the approximation at x converges to $f(x)$.

Since the seminal paper of Sacks et al. (1989), *kriging* has been one of the most popular methods for building approximations in the context of computer experiments (see, e.g., Santner et al. 2003). In the framework of kriging, the unknown function f is seen as a sample path of a stochastic process ξ , which turns the problem of approximation of f into a prediction problem for the process ξ . In this paper, we shall assume that the mean and the covariance functions are known. Motivated by the analysis by Vazquez and Bect (2009) of the *expected improvement* algorithm, a popular kriging-based optimization algorithm, we discuss several issues related to the pointwise consistency of the kriging predictor, that is, the convergence of the kriging predictor to the true value of ξ at a fixed point $x \in \mathbb{X}$. These issues are barely documented in the literature, and we believe them to be of general importance for the asymptotic analysis of sequential design procedures based on kriging.

The paper is organized as follows. Section 2 introduces notation and various formulations of pointwise consistency, using the reproducing kernel Hilbert space (RKHS) attached to ξ . Section 3 investigates whether L^2 -pointwise consistency at x can hold when x is not in the adherence of the set $\{x_n, n \geq 1\}$. Conversely, assuming that x is in the adherence, Section 4 studies the set of sample paths $f = \xi(\omega, \cdot)$ for which pointwise consistency holds. In particular, we fix an erroneous claim of Yakowitz and Szidarovszky (1985)—namely, that the kriging predictor is pointwise consistent for all continuous sample paths under some assumptions.

2 Several Formulations of Pointwise Consistency

Let ξ be a second-order process defined on a probability space $(\Omega, \mathcal{A}, \mathbb{P})$, with parameter $x \in \mathbb{X} \subseteq \mathbb{R}^d$. Without loss of generality, it will be assumed that the mean of ξ is zero and that $\mathbb{X} = \mathbb{R}^d$. The covariance function of ξ will be denoted by $k(x, y) := \mathbb{E}[\xi(x)\xi(y)]$, and the following assumption will be used throughout the paper:

Assumption 1. The covariance function k is continuous.

The kriging predictor of $\xi(x)$, based on the observations $\xi(x_i)$, $i = 1, \dots, n$, is the orthogonal projection

$$\hat{\xi}(x; x_n) := \sum_{i=1}^n \lambda^i(x; x_n) \xi(x_i) \quad (1)$$

of $\xi(x)$ onto $\text{span}\{\xi(x_i), i = 1, \dots, n\}$. The variance of the prediction error, also called the *kriging variance* in the literature of geostatistics (see, e.g., Chilès and

Delfiner 1999), or the *power function* in the literature of radial basis functions (see, e.g., Wu and Schaback 1993), is

$$\begin{aligned} \sigma^2(x; \underline{x}_n) &:= \text{var} \left[\xi(x) - \hat{\xi}(x; \underline{x}_n) \right] \\ &= k(x, x) - \sum_i \lambda^i(x; \underline{x}_n) k(x, x_i). \end{aligned}$$

For any $x \in \mathbb{R}^d$, and any sample path $f = \xi(\omega, \cdot)$, $\omega \in \Omega$, the values $\xi(\omega, x) = f(x)$ and $\hat{\xi}(\omega, x; \underline{x}_n)$ can be seen as the result of the application of an evaluation functional to f . More precisely, let δ_x be the Dirac measure at $x \in \mathbb{R}^d$, and let $\lambda_{n,x}$ denote the measure with finite support defined by $\lambda_{n,x} := \sum_{i=1}^n \lambda^i(x; \underline{x}_n) \delta_{x_i}$. Then, for all $\omega \in \Omega$, $\xi(\omega, x) = \langle \delta_x, f \rangle$ and $\hat{\xi}(\omega, x; \underline{x}_n) = \langle \lambda_{n,x}, f \rangle$. Pointwise consistency at $x \in \mathbb{R}^d$, defined in Section 1 as the convergence of $\hat{\xi}(\omega, x; \underline{x}_n)$ to $\xi(x)$, can thus be seen as the convergence of $\lambda_{n,x}$ to δ_x in some sense.

Let \mathcal{H} be the RKHS of functions generated by k , and \mathcal{H}^* its dual space. Denote by $(\cdot, \cdot)_{\mathcal{H}}$ (resp. $(\cdot, \cdot)_{\mathcal{H}^*}$) the inner product of \mathcal{H} (resp. \mathcal{H}^*), and by $\|\cdot\|_{\mathcal{H}}$ (resp. $\|\cdot\|_{\mathcal{H}^*}$) the corresponding norm. It is well-known (see, e.g., Wu and Schaback 1993) that

$$\|\delta_x - \lambda_{n,x}\|_{\mathcal{H}^*}^2 = \|k(x, \cdot) - \sum_i \lambda^i(x; \underline{x}_n) k(x_i, \cdot)\|_{\mathcal{H}}^2 = \sigma^2(x; \underline{x}_n).$$

Therefore, the convergence $\lambda_{n,x} \rightarrow \delta_x$ holds strongly in \mathcal{H}^* if and only if the kriging predictor is $L^2(\Omega, \mathcal{A}, P)$ -consistent at x ; that is, if $\sigma^2(x; \underline{x}_n)$ converges to zero. Since k is continuous, it is easily seen that $\sigma^2(x; \underline{x}_n) \rightarrow 0$ as soon as x is adherent to $\{x_n, n \geq 1\}$. Indeed,

$$\sigma^2(x, \underline{x}_n) \leq E[(\xi(x) - \xi(x_{\varphi_n}))^2] = k(x, x) + k(x_{\varphi_n}, x_{\varphi_n}) - 2k(x, x_{\varphi_n}),$$

with $(\varphi_n)_{n \in \mathbb{N}}$ a non-decreasing sequence such that $\forall n \geq 1, \varphi_n \leq n$ and $x_{\varphi_n} \rightarrow x$. As explained by Vazquez and Bect (2009), it is sometimes important to work with covariance functions such that the converse holds. That leads to our first open issue, which will be discussed in Section 3:

Problem 1. Find necessary and sufficient conditions on a continuous covariance k such that $\sigma^2(x; \underline{x}_n) \rightarrow 0$ implies that x is adherent to $\{x_n, n \geq 1\}$.

Moreover, since strong convergence in \mathcal{H}^* implies weak convergence in \mathcal{H}^* , we have

$$\lim_{n \rightarrow \infty} \sigma^2(x; \underline{x}_n) = 0 \implies \forall f \in \mathcal{H}, \quad \lim_{n \rightarrow \infty} \langle \lambda_{n,x}, f \rangle = \langle \delta_x, f \rangle = f(x). \quad (2)$$

Therefore, if x is adherent to $\{x_n, n \geq 1\}$, pointwise consistency holds for all sample paths $f \in \mathcal{H}$. However, this result is not satisfying from a Bayesian point of view since $P\{\xi \in \mathcal{H}\} = 0$ if ξ is Gaussian (see, e.g., Lukic and Beder 2001, Driscoll's theorem). In other words, modeling f as a Gaussian process means that f cannot be expected to belong to \mathcal{H} . This leads to our second problem:

Problem 2. For a given covariance function k , describe the set of functions \mathcal{G} such that, for all sequences $(x_n)_{n \geq 1}$ in \mathbb{R}^d and all $x \in \mathbb{R}^d$,

$$\lim_{n \rightarrow \infty} \sigma^2(x; \underline{x}_n) = 0 \implies \forall f \in \mathcal{G}, \quad \lim_{n \rightarrow \infty} \langle \lambda_{n,x}, f \rangle = f(x). \quad (3)$$

An important question related to this problem, to be discussed in Section 4, is to know whether the set \mathcal{G} contains the set $C(\mathbb{R}^d)$ of all continuous functions. Before proceeding, we can already establish a result which ensures that considering the kriging predictor is relevant from a Bayesian point of view.

Theorem 1. *If ξ is Gaussian, then $\{\xi \notin \mathcal{G}\}$ is P-negligible.*

Proof. If ξ is Gaussian, it is well-known that $\hat{\xi}(x; \underline{x}_n) = E[\xi(x) \mid \mathcal{F}_n]$ a.s., where \mathcal{F}_n denotes the σ -algebra generated by $\xi(x_1), \dots, \xi(x_n)$. Note that $(E[\xi(x) \mid \mathcal{F}_n])$ is an L^2 -bounded martingale sequence and therefore converges, a.s. and in L^2 -norm, to a random variable ξ_∞ (see, e.g., Williams 1991). \square

3 Pointwise Consistency in L^2 -Norm and the No-Empty-Ball Property

The following definition has been introduced by Vazquez and Bect (2009):

Definition 1. A random process ξ has the No-Empty-Ball (NEB) property if, for all sequences $(x_n)_{n \geq 1}$ in \mathbb{R}^d and all $x \in \mathbb{R}^d$, the following assertions are equivalent:

- i) x is an adherent point of the set $\{x_n, n \geq 1\}$,
- ii) $\sigma^2(x, \underline{x}_n) \rightarrow 0$ when $n \rightarrow +\infty$.

The NEB property implies that there can be no empty ball centered at x if the prediction error at x converges to zero—hence the name. Since k is continuous, the implication 1.i \implies 1.ii is true. Therefore, Problem 1 amounts to finding necessary and sufficient conditions on k for ξ to have the NEB property.

Our contribution to the solution of Problem 1 will be twofold. First, we shall prove that the following assumption, introduced by Yakowitz and Szidarovszky (1985), is a sufficient condition for the NEB property:

Assumption 2. The process ξ is second-order stationary and has spectral density S , with the property that S^{-1} has at most polynomial growth.

In other words, Assumption 2 means that there exist $C > 0$ and $r \in \mathbb{N}^*$ such that $S(u)(1 + |u|^r) \geq C$, almost everywhere on \mathbb{R}^d . Note that this is an assumption on k , which prevents it from being too regular. In particular, the so-called *Gaussian covariance*,

$$k(x, y) = s^2 e^{-\alpha \|x-y\|^2}, \quad s > 0, \alpha > 0, \quad (4)$$

does not satisfy Assumption 2. In fact, and this is the second part of our contribution, we shall show that ξ with covariance function (4) does not possess the NEB property. Assumption 2 still allows consideration of a large class of covariance functions, which includes the class of (non-Gaussian) exponential covariances

$$k(x, y) = s^2 e^{-\alpha \|x-y\|^\beta}, \quad s > 0, \alpha > 0, 0 < \beta < 2, \quad (5)$$

and the class of Matérn covariances (popularized by Stein 1999).

To summarize, the main result of this section is:

Proposition 1.

- i) If Assumption 2 holds, then ξ has the NEB property.
- ii) If ξ has the Gaussian covariance given by (4), then ξ does not possess the NEB property.

The proof of Proposition 1 is given in Section 5. To the best of our knowledge, finding necessary and sufficient conditions for the NEB property—in other words, solving Problem 1—is still an open problem.

4 Pointwise Consistency for Continuous Sample Paths

An important question related to Problem 2 is to know whether the set \mathcal{G} contains the set $C(\mathbb{R}^d)$ of all continuous functions. Yakowitz and Szidarovszky (1985, Lemma 2.1) claim, but fail to establish, the following:

Claim 1. Let Assumption 2 hold. Assume that $\{x_n, n \geq 1\}$ is bounded, and denote by \mathbb{X}_0 its (compact) closure in \mathbb{R}^d . Then, if $x \in \mathbb{X}_0$,

$$\forall f \in C(\mathbb{R}^d), \quad \lim_{n \rightarrow \infty} \langle \lambda_{n,x}, f \rangle = f(x).$$

Their incorrect proof has two parts, the first of which is correct; it says in essence that, if $x \in \mathbb{X}_0$ (i.e., if x is adherent to $\{x_n, n \geq 1\}$), then

$$\forall f \in \mathcal{S}(\mathbb{R}^d), \quad \lim_{n \rightarrow \infty} \langle \lambda_{n,x}, f \rangle = f(x), \quad (6)$$

where $\mathcal{S}(\mathbb{R}^d)$ is the vector space of rapidly decreasing functions¹. In fact, this result

¹ Recall that $\mathcal{S}(\mathbb{R}^d)$ corresponds to those $f \in C^\infty(\mathbb{R}^d)$ for which

$$\sup_{|v| \leq N} \sup_{x \in \mathbb{R}^d} (1 + |x|^2)^N |(D^v f)(x)| < \infty$$

for $N = 0, 1, 2, \dots$, where D^v denotes differentiation of order v .

stems from the weak convergence result (2), once it has been remarked that² $\mathcal{S}(\mathbb{R}^d) \subset \mathcal{H}$ under Assumption 2.

The second part of the proof of Claim 1 is flawed because the extension of the convergence result from $\mathcal{S}(\mathbb{R}^d)$ to $C(\mathbb{R}^d)$, on the ground that $\mathcal{S}(\mathbb{R}^d)$ is dense in $C(\mathbb{R}^d)$ for the topology of the uniform convergence on compact sets, does not work as claimed by the authors. To get an insight into this, let $f \in C(\mathbb{R}^d)$, and let $(\phi_k) \in \mathcal{S}(\mathbb{R}^d)^\mathbb{N}$ be a sequence that converges to f uniformly on \mathbb{X}_0 . Then we can write

$$\begin{aligned} |\langle \lambda_{n,x}, f \rangle - f(x)| &\leq |\langle \lambda_{n,x}, f - \phi_k \rangle| + |\langle \lambda_{n,x} - \delta_x, \phi_k \rangle| + |\phi_k(x) - f(x)| \\ &\leq (1 + \|\lambda_{n,x}\|_{\text{TV}}) \sup_{\mathbb{X}_0} |f - \phi_k| + |\langle \lambda_{n,x} - \delta_x, \phi_k \rangle|, \end{aligned}$$

where $\|\lambda_{n,x}\|_{\text{TV}} := \sum_{i=1}^n |\lambda^i(x; \underline{x}_n)|$ is the total variation norm of $\lambda_{n,x}$, also called the *Lebesgue constant* (at x) in the literature of approximation theory. If we assume that the Lebesgue constant is bounded by $K > 0$, then we get, using (6),

$$\limsup_{n \rightarrow \infty} |\langle \lambda_{n,x}, f \rangle - f(x)| \leq (1 + K) \sup_{\mathbb{X}_0} |f - \phi_k| \xrightarrow[k \rightarrow \infty]{} 0.$$

Conversely, if the Lebesgue constant is not bounded, the Banach-Steinhaus theorem asserts that there exists a dense subset G of $(C(\mathbb{R}^d), \|\cdot\|_\infty)$ such that, for all $f \in G$, $\sup_{n \geq 1} |\langle \lambda_{n,x}, f \rangle| = +\infty$ (see, e.g., Rudin 1987, Section 5.8).

Unfortunately, little is known about Lebesgue constants in the literature of kriging and kernel regression. To the best of our knowledge, whether the Lebesgue constant is bounded remains an open problem—although there is empirical evidence in De Marchi and Schaback (2008) that the Lebesgue constant could be bounded in some cases under Assumption 2.

Thus, the best result that we can state for now is a fixed version of Claim 1. Note that the foregoing discussion is still valid if Assumption 2 is replaced by the weaker assumption that \mathcal{H} is dense in $(C(\mathbb{R}^d), \|\cdot\|_\infty)$. Kernels with this property have been called *universal kernels* by Steinwart (2001).

Theorem 2. *Let k be a universal kernel on \mathbb{X} . Assume that $\{x_n, n \geq 1\}$ is bounded, and denote by \mathbb{X}_0 its (compact) closure in \mathbb{R}^d . Then, for all $x \in \mathbb{X}_0$, the following assertions are equivalent:*

- i) $\forall f \in C(\mathbb{R}^d), \lim_{n \rightarrow \infty} \langle \lambda_{n,x}, f \rangle = f(x)$,
- ii) *the Lebesgue constant at x is bounded.*

The class of all universal kernels is wider than that of all kernels satisfying Assumption 2, and is not restricted to translation-invariant kernels—or equivalently, kernels associated to stationary processes; see Steinwart (2001) for examples.

² Indeed, under Assumption 2, we have $\forall f \in \mathcal{S}(\mathbb{R}^d)$,

$$\|f\|_{\mathcal{H}}^2 = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} |\tilde{f}(u)|^2 S(u)^{-1} du \leq \frac{1}{C(2\pi)^d} \int_{\mathbb{R}^d} |\tilde{f}(u)|^2 (1 + |u|^r) du < +\infty,$$

where \tilde{f} is the Fourier transform of f (see, e.g., Wu and Schaback 1993).

Note also that the Gaussian covariance (4) is a universal kernel Steinwart (2001, Example 1). Numerical experiments in De Marchi and Schaback (2008) suggest that the Lebesgue constant could be unbounded for this model in some cases, which would imply by Theorem 2 that the kriging predictor is not pointwise consistent for all continuous sample paths.

5 Proof of Proposition 1

Assume that $x \in \mathbb{R}^d$ is not adherent to $\{x_n, n \geq 1\}$. Then, there exists a $C^\infty(\mathbb{R}^d)$ compactly supported function f such that $f(x) \neq 0$ and $f(x_i) = 0, \forall i \in \{1, \dots, n\}$. For such a function, the quantity $\langle \lambda_{n,x}, f \rangle$ cannot converge to $f(x)$ since

$$\langle \lambda_{n,x}, f \rangle = \sum_{i=1}^n \lambda^i(x; \underline{x}_n) f(x_i) = 0 \neq f(x).$$

Under Assumption 2, $\mathcal{S}(\mathbb{R}^d) \subset \mathcal{H}$, as explained in Section 4. Thus, $f \in \mathcal{H}$; and it follows that $\lambda_{n,x}$ cannot converge (weakly, hence strongly) to δ_x in \mathcal{H}^* . This proves the first assertion of Proposition 1.

In order to prove the second assertion, pick any sequence $(x_n)_{n \geq 1}$ such that the closure \mathbb{X}_0 of $\{x_n, n \geq 1\}$ has a non-empty interior. We will show that $\sigma(x; \underline{x}_n) \rightarrow 0$ for all $x \in \mathbb{R}^d$. Then, choosing $x \notin \mathbb{X}_0$ proves the claim.

Recall that $\widehat{\xi}(x; \underline{x}_n)$ is the orthogonal projection of $\xi(x)$ onto $\text{span}\{\xi(x_i), i = 1, \dots, n\}$ in $L^2(\Omega, \mathcal{A}, \mathbb{P})$. Using the fact that the mapping $\xi(x) \mapsto k(x, \cdot)$ extends linearly to an isometry³ from $\overline{\text{span}}\{\xi(y), y \in \mathbb{R}^d\}$ to \mathcal{H} , we get that

$$\sigma(x; \underline{x}_n) = \|\xi(x) - \widehat{\xi}(x; \underline{x}_n)\| = d_{\mathcal{H}}(k(x, \cdot), H_n),$$

where $d_{\mathcal{H}}$ is the distance in \mathcal{H} , and H_n is the subspace of \mathcal{H} generated by $k(x_i, \cdot), i = 1, \dots, n$. Therefore

$$\lim_{n \rightarrow \infty} \sigma(x; \underline{x}_n) = \lim_{n \rightarrow \infty} d_{\mathcal{H}}(k(x, \cdot), H_n) = d_{\mathcal{H}}(k(x, \cdot), H_\infty),$$

where $H_\infty = \overline{\cup_{n \geq 1} H_n}$. Any function $f \in H_\infty^\perp$ satisfies $f(x_i) = (f, k(x_i, \cdot)) = 0$ and therefore vanishes on \mathbb{X}_0 , since \mathcal{H} is a space of continuous functions. Corollary 3.9 of Steinwart, Hush, and Scovel (2006) leads to the conclusion that $f = 0$ since \mathbb{X}_0 has a non-empty interior. We have proved that $H_\infty^\perp = \{0\}$, hence that $H_\infty = \mathcal{H}$ since H_∞ is a closed subspace. As a consequence, $\lim_{n \rightarrow \infty} \sigma(x; \underline{x}_n) = d_{\mathcal{H}}(k(x, \cdot), H_\infty) = 0$, which completes the proof. \square

³ often referred to as Loève's isometry (see, e.g., Lukic and Beder 2001)

References

- Chilès, J.-P. and P. Delfiner (1999). *Geostatistics: Modeling Spatial Uncertainty*. New York: Wiley.
- De Marchi, S. and R. Schaback (2008). Stability of kernel-based interpolation. *Advances in Computational Mathematics*. doi: 10.1007/s10444-008-9093-4.
- Lukic, M. N. and J. H. Beder (2001). Stochastic processes with sample paths in reproducing kernel Hilbert spaces. *Transactions of the American Mathematical Society* 353, 3945–3969.
- Rudin, W. (1987). *Real and Complex Analysis, 3rd ed.* New York: McGraw-Hill.
- Sacks, J., W. J. Welch, T. J. Mitchell, and H. P. Wynn (1989). Design and analysis of computer experiments. *Statistical Science* 4, 409–435.
- Santner, T. J., B. J. Williams, and W. I. Notz (2003). *The Design and Analysis of Computer Experiments*. New-York: Springer-Verlag.
- Stein, M. L. (1999). *Interpolation of Spatial Data: Some Theory for Kriging*. New York: Springer-Verlag.
- Steinwart, I. (2001). On the influence of the kernel on the consistency of support vector machines. *Journal of Machine learning Research* 2, 67–93.
- Steinwart, I., D. Hush, and C. Scovel (2006). An explicit description of the reproducing kernel Hilbert spaces of Gaussian RBF kernels. *IEEE Transactions on Information Theory* 52, 4635–4643.
- Vazquez, E. and J. Bect (2009). On the convergence of the expected improvement algorithm. Preprint available on arXiv, <http://arxiv.org/abs/0712.3744v2>.
- Villemonteix, J. (2008). *Optimisation de Fonctions Coûteuses*. Ph. D. thesis, Université Paris-Sud XI, Faculté des Sciences d’Orsay.
- Williams, D. (1991). *Probability with Martingales*. Cambridge: Cambridge University Press.
- Wu, Z. and R. Schaback (1993). Local error estimates for radial basis function interpolation of scattered data. *IMA Journal of Numerical Analysis* 13, 13–27.
- Yakowitz, S. J. and F. Szidarovszky (1985). A comparison of kriging with nonparametric regression methods. *Journal of Multivariate Analysis* 16, 21–53.

Information in a Two-stage Adaptive Optimal Design for Normal Random Variables having a One Parameter Exponential Mean Function

Ping Yao and Nancy Flournoy

Abstract This paper explores the characteristics of information derived from sequentially implementing estimated optimal designs. In such sequential experiments, called *adaptive optimal designs*, each stage uses an optimal design estimated from the data obtained in all prior stages. The measure that is used in adaptive optimal designs to construct treatment allocation procedures is, by definition, neither the observed nor the expected (Fisher) information. We explore these information measures in the context of a two-stage adaptive optimal design under a simple model. Specifically, random variables are assumed to be normal with a one parameter exponential mean function. With this model, some explicit results are obtained.

1 Introduction

Chernoff (1953) suggested that optimal designs for nonlinear functions be approximated by guessing the parameter values; this may be inefficient when the guess is far from the actual parameter value. In *adaptive optimal design*, sequential experiments use an optimal design estimated from all prior stages. This approach was suggested by Box and Hunter Box and Hunter (1963), White (1975), Silvey (1980), Dragalin, Fedorov, and Wu (2008), among others. Its appeal is that if an adaptive optimal design converges to the true optimal design, heuristically arguing, the overall design will become more efficient with additional stages.

In the adaptive optimal design literature, in place of constructing a likelihood from the joint density for responses and design points, responses have been treated

Ping Yao

Public Health and Health Education Program, 209K Wirtz Hall, Northern Illinois University, DeKalb, IL 60115 USA, e-mail: pyao@niu.edu

Nancy Flournoy

Department of Statistics, 146 Middlebush Hall, University of Missouri, Columbia, MO 65203 USA
e-mail: flournoyn@missouri.edu

as independent conditional, on the treatment - both for selecting the next design point and for evaluating the design's efficiency. Silvey (1980) and others point out that the information measure they employ is not, by definition, Fisher's information. While conditioning is generally accepted for analysis, the role of conditioning in adaptive design construction has not been clarified.

Denote the likelihood after s stages by \mathcal{L}_s ; $-\frac{d^2}{d\theta^2} \log \mathcal{L}_s |_{\theta=\hat{\theta}}$ is the *observed information* (see Efron and Hinkley 1978 and Lindsay and Li 1997); $Var(\frac{d}{d\theta} \log \mathcal{L}_s)$ is the *expected or Fisher's information*. The information measure \mathcal{M} given in Section 4 by (6) and used in the adaptive optimal design literature (e.g., Dragalin, Fedorov, and Wu (2008)) is not, by definition, either of these.

Section 2 introduces the nonlinear model and the two-stage design. Section 3 presents results on the stage 2 design point. In Section 4, properties of the information measures are investigated. Section 5 concludes with a brief discussion.

2 A Two-stage Design for Normal Random Variables having a One Parameter Exponential Mean Function

Let $Y = \eta(X) + \varepsilon$, $0 \leq X \leq b < \infty$. Assume $\eta(X) = \exp(-\theta X)$, $\theta > 0$. Assume $\varepsilon \sim \mathcal{N}(0, 1)$. Suppose n subjects are treated at $X = x_1$ (fixed) and independent responses $y_1 = (y_{11}, \dots, y_{1n})$ are observed. The likelihood at the end of stage 1 is

$$\mathcal{L}_1(\theta, x_1, y_1) = (2\pi)^{-n/2} \exp \left\{ -\frac{1}{2} \sum_{j=1}^n (y_{1j} - \eta(x_1))^2 \right\}.$$

Let $\bar{y}_i = n^{-1} \sum_{j=1}^n y_{ij}$, $j \geq 1, i = 1, 2$. Then if $\bar{y}_1 \leq 0$, the score function, $\frac{d}{d\theta} \log \mathcal{L}_1 = -\sum_1^n (y_{1j} - e^{-\theta x_1}) x_1 e^{-\theta x_1}$, is positive; if $0 < \bar{y}_1 < 1$, setting the score function equal to zero yields the maximum likelihood estimate (MLE). Then because the term $-\log \bar{y}_1/x_1$ becomes negative if $\bar{y}_1 > 1$,

$$\hat{\theta}_1 = \begin{cases} \infty & \text{if } \bar{y}_1 \leq 0, \\ -\log \bar{y}_1/x_1, & \text{if } 0 < \bar{y}_1 < 1, \\ 0 & \text{if } \bar{y}_1 \geq 1. \end{cases} \tag{1}$$

The information with respect to $f(y_{1j}|x_1, \theta)$ after stage 1, $j = 1, \dots, n$ is

$$\mathcal{M}(x_1, \theta) = -E_{y_{11}|x_1} \left(\frac{d^2}{d\theta^2} \log f(y_{11}|x_1, \theta) \right) = x_1^2 \exp \{-2\theta x_1\}.$$

Select the stage 2 design point as

$$x_2 = \arg \max_x -E_{y_{21}|x} \left(\frac{d^2}{d\theta^2} \log f(y_{21}|x, \theta) \right) \Big|_{\theta=\hat{\theta}_1} = \arg \max_x (x^2 \exp \{-2\hat{\theta}_1 x\}),$$

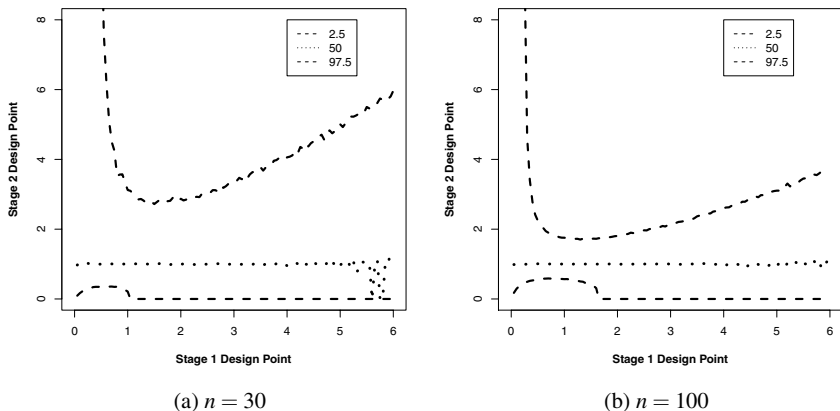


Fig. 1: Stage 2 Design Points: 2.5th, 50th and 97.5th Percentiles ($\theta = 1$; $b = 100$)

which is

$$x_2 = \begin{cases} 0 & \text{if } \bar{y}_1 \leq 0, \\ -x_1 / \log \bar{y}_1 & \text{if } 0 < \bar{y}_1 < 1, \\ b & \text{if } \bar{y}_1 \geq 1. \end{cases} \tag{2}$$

Assuming cohorts of equal size, observe $y_2 = (y_{21}, \dots, y_{2n})$. Then assuming responses given the treatment are independent of the past, i.e., $f(y_2|x_2, x_1, y_1, \theta) = f(y_2|x_2, \theta)$, the likelihood after stage 2 is

$$\mathcal{L}_2(x_1, x_2, y_1, y_2, \theta) = f(y_2|x_2, \theta)f(x_2|x_1, y_1, \theta)f(y_1|x_1, \theta).$$

Because x_2 is completely determined by x_1 and \bar{y}_1 , $f(x_2|x_1, y_1, \theta) = 1$; now the total likelihood and score function for two stages, respectively, can be written as

$$\mathcal{L}_2(x_1, x_2, y_1, y_2, \theta) = \prod_{i=1}^2 f(y_i|x_i, \theta), \quad \frac{\partial}{\partial \theta} \log \mathcal{L}_2 = \sum_{i=1}^2 \frac{\partial}{\partial \theta} \log f(y_i|x_i, \theta). \tag{3}$$

3 Properties of the Stage 2 Design Point

For the model with $b = 100$ and $\theta = 1$, Figure 3 displays the simulated 2.5th, 50th and 97.5th percentiles of x_2 as a function of x_1 for $n = 30$ and 100. Ten thousand replicates of x_2 were simulated for each plotted value of x_1 . The median of x_2 is close to $\operatorname{argmax}_x \{ \mathcal{M}(x, \theta) \} = 1$ regardless of x_1 , but the range from the 2.5th to 97.5th percentile of x_2 depends strongly on the initial design point.

$P(x_2 = 0) \geq 0.025$ when $x_1 > 2$ for $n = 100$ and $x_1 \geq 1.2$ for $n = 30$. The minimum of the 97.5th percentile of x_2 occurs for values of x_1 somewhat larger than one, more so for $n = 30$ than for $n = 100$; for smaller x_1 , the 97.5th percentile of x_2 rises steeply to b ; for larger values of x_1 , the 97.5th percentile rises much more slowly toward b .

Given θ , the variance of x_2 can be calculated numerically from

$$E(x_2|x_1) = bP(\bar{y}_1 > 1|x_1) - x_1 \int_{0^+}^{1^-} (\log \bar{y}_1)^{-1} f(\bar{y}_1|x_1) d\bar{y}_1;$$

$$E(x_2^2|x_1) = b^2P(\bar{y}_1 > 1|x_1) + x_1^2 \int_{0^+}^{1^-} (\log \bar{y}_1)^{-2} f(\bar{y}_1|x_1) d\bar{y}_1,$$

Theorem 1. $\hat{\theta} \xrightarrow{a.s.} \theta$ and $x_2 \xrightarrow{a.s.} \theta^{-1}$ as $n \rightarrow \infty$.

Proof. Let Z denote a standard normal random variable. Then

$$P(\bar{y}_1 \leq 0|x_1) = P\left(Z < -\sqrt{n}e^{-\theta x_1}\right) \xrightarrow{n \rightarrow \infty} 0,$$

$$P(\bar{y}_1 \geq 1|x_1) = P\left\{Z > \sqrt{n}\left(1 - e^{-\theta x_1}\right)\right\} \xrightarrow{n \rightarrow \infty} 0. \tag{4}$$

The results follow from the law of large numbers and continuity of the transformations from \bar{y}_1 to x_2 within the range $0 < \bar{y}_1 < 1$. \square

Using (4) with the Delta Method, we obtain

Theorem 2. For the two stage design under the model $y = e^{-x\theta} + \varepsilon$, $0 \leq x \leq b < \infty$, where $\varepsilon \sim \mathcal{N}(0, 1)$ and x_1 is given, as $n \rightarrow \infty$,

$$\sqrt{n}(x_2 - \theta^{-1}) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma^2), \tag{5}$$

where $\sigma^2 = x_1^{-2}\theta^{-4}e^{2\theta x_1}$.

4 Information Measures

The *observed information* (per-subject), as defined by Efron and Hinkley (1978) , is

$$-\frac{1}{2n} \frac{d^2}{d\theta^2} \log \mathcal{L}_2 = -\frac{1}{2n} \sum_{i=1}^2 \frac{d^2}{d\theta^2} \log f(y_i|x_i, \theta)$$

$$= \frac{1}{2n} \sum_{i=1}^2 \sum_{j=1}^n \left(2e^{-\theta x_i} - y_{ij}\right) x_i^2 e^{-\theta x_i} \Big|_{\theta=\hat{\theta}}.$$

The *adaptive optimal design information* averages $\mathcal{M}(x_1, \theta)$ and $\mathcal{M}(x_2, \theta)$:

$$\mathcal{M}(\{x_i\}_1^2, \theta) = -\frac{1}{2n} \sum_{i=1}^2 E_{y_i|x_i} \left(\frac{d^2}{d\theta^2} \log f(y_i|x_i, \theta) \right) = \frac{1}{2} \sum_{i=1}^2 x_i^2 e^{-2\theta x_i}. \quad (6)$$

Note the adaptive optimal information is a function of \bar{y}_1 through x_2 , whereas the observed information is a function of both \bar{y}_1 and \bar{y}_2 . It follows from the strong law of large numbers and continuity of the transformation from \bar{y}_1 for $\bar{y}_1 \in (0, 1)$ that

$$-\frac{d^2}{d\theta^2} \log f(y_i|x_i, \theta) = \mathcal{M}(x_i, \theta) - \frac{\sum_{j=1}^n \varepsilon_{ij}}{2n} x_i^2 e^{-\theta x_i} \xrightarrow{a.s.} \mathcal{M}(x_i, \theta) \quad \text{as } n \rightarrow \infty. \quad (7)$$

Because cross-product terms $E(\log f(y_i|x_i, \theta) \log f(y_j|x_j, \theta))$ are zero [see Hall and Heyde (1980), page 8], Fisher's information can be written as

$$\begin{aligned} \text{Var} \left(\frac{d}{d\theta} \log \mathcal{L}_2 \right) &= \sum_{i=1}^2 E \left[\sum_{j=1}^n (y_{ij} - e^{-\theta x_i}) x_i e^{-\theta x_i} \right]^2 \\ &= nE \left(E \left[(y_{2j} - e^{-\theta x_2})^2 x_2^2 e^{-2\theta x_2} \right] \middle| x_1, x_2, y_1 \right) + n x_1^2 e^{-2\theta x_1}. \end{aligned}$$

But $E \left((y_{2j} - e^{-\theta x_2})^2 \middle| x_1, x_2, y_1 \right) = \text{Var}(y_{2j} | y_1, x_2, x_1) = 1$, so

$$\begin{aligned} \frac{1}{2n} \text{Var} \left(\frac{d}{d\theta} \log \mathcal{L}_2 \right) &= \frac{1}{2} E \left(x_2^2 e^{-2\theta x_2} \right) + \frac{1}{2} x_1^2 e^{-2\theta x_1} \\ &= \frac{1}{2} E(\mathcal{M}(x_2, \theta)) + \frac{1}{2} E(\mathcal{M}(x_1, \theta)) = E(\mathcal{M}(\{x_i\}_1^2, \theta)). \end{aligned}$$

Theorem 3 provides a large sample approximation of Fisher's information.

Theorem 3.

$$\lim_{n \rightarrow \infty} E \left(x_2^2 e^{-2\theta x_2} \right) \xrightarrow{p} \theta^{-2} e^{-2}.$$

Proof. Expand $\mathcal{M}(x_2, \theta)$ into a Taylor series of order two about x :

$$\begin{aligned} x_2^2 e^{-2\theta x_2} &= x^2 e^{-2\theta x} + (x_2 - x) \left[2x e^{-2\theta x} (1 - \theta x) \right] \\ &\quad + (x_2 - x)^2 e^{-2\theta x} (1 - 4x\theta + 2\theta^2 x^2) + o((x_2 - x)^2). \end{aligned} \quad (8)$$

Evaluating x at $\text{argmax}_x \{ \mathcal{M}(x, \theta) \} = \theta^{-1}$, the second term on the right of (8) is zero yielding

$$E \left(x_2^2 e^{-2\theta x_2} \right) = \theta^{-2} e^{-2} - \text{Var}(x_2) e^{-2} + o((x_2 - \theta^{-1})^2).$$

The error term goes to zero by Theorem 1 and $\text{Var}(x_2) \rightarrow 0$ as $\text{Var}(\bar{y}_1) \rightarrow 0$. \square

$\text{Var} \left(\frac{d}{d\theta} \log \mathcal{L}_2 \right)$ can be approximated by $\frac{1}{2} \hat{\theta}^{-2} e^{-2} + \frac{1}{2} x_1^2 e^{-2\hat{\theta} x_1}$.

4.1 A Simulated Illustration

Again taking $b = 100$ and $\theta = 1$, Figure 4.1 shows $-\frac{d^2}{d\theta^2} \log f(y_1|x_1, \theta)$ as a function of the stage 1 design point. The median of $-\frac{d^2}{d\theta^2} \log f(y_1|x_1, \theta)$ is $\mathcal{M}(x, \theta)$ by (7), which also equals Fisher’s information since x_1 is given. The median values of $-\frac{d^2}{d\theta^2} \log f(y_1|x_1, \theta)$ attain their maximum of 0.135 at $\operatorname{argmax}_x \{\mathcal{M}(x, \theta)\} = 1$; the 97.5th percentiles are maximum at $x > 1$; the 2.5th percentiles are negative for many stage one design points.

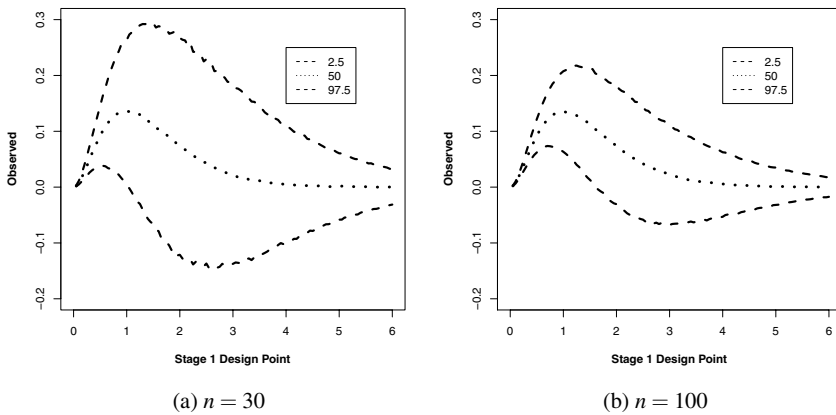


Fig. 2: Stage 1: 2.5th, 50th and 97.5th Percentiles of $-\frac{d^2}{d\theta^2} \log f(y_1|x_1, \theta)$; ($\theta = 1$; $b = 100$)

Now focus on the increment in the information measures that comes from stage 2 of the experiment. Both the observed information for stage 2 and $\mathcal{M}(x_2, \theta)$ are random, as they are functions \bar{y}_1 via x_2 . Information measures obtained during stage 2 were calculated from 10,000 simulated replicates for each plotted value of x_1 . The 2.5th, 50th and 97.5th percentiles are shown in Figure 4.1.

The $\max_{x_1} \{\mathcal{M}(x_2, \theta)\} = 0.135$, which is the asymptotic Fisher’s information. The 97.5th percentiles of $\mathcal{M}(x_2, \theta)$ attain 0.135 at all but the highest values of x_1 for $n = 100$ and 30. In contrast, the 97.5th percentile of $-\frac{d^2}{d\theta^2} \log f(y_2|x_2, \theta)$ is greater than 0.135 except for values of x_1 somewhat less than one. Furthermore, $-\frac{d^2}{d\theta^2} \log f(y_2|x_2, \theta)$ is negative with high probability.

The median of the adaptive optimal information, $\mathcal{M}(x_2, \theta)$, attains its maximum value when $x_1 = 1$ for $n = 100$ and 30. In addition, the median of $\mathcal{M}(x_2, \theta)$ comes closer to 0.135 at $x_1 = 1$ as the sample size increases. Indeed, the median of $\mathcal{M}(x_2, \theta)$ is close to 0.135 for a range of values of x_1 that includes $x_1 = 1$; this range is larger for $n = 100$ than for $n = 30$. For $n = 30$, the 2.5th percentile of

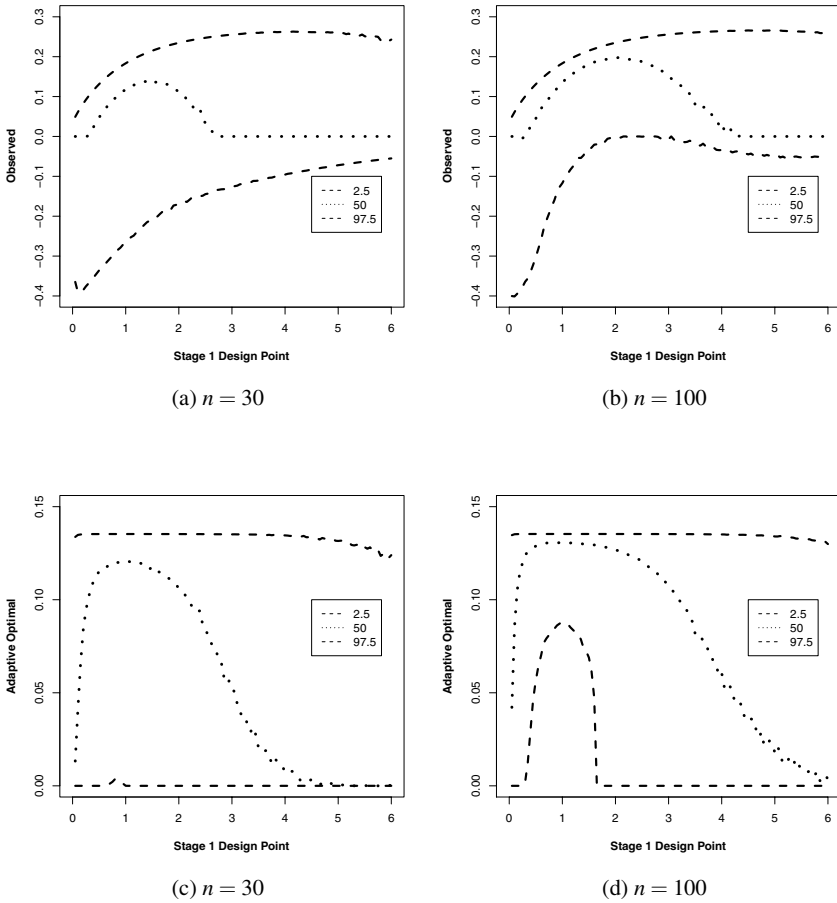


Fig. 3: Stage 2 Information: 2.5th, 50th and 97.5th Percentiles of $-\frac{d^2}{d\theta^2} \log f(y_2|x_2, \theta)$ in (a,b) and $\mathcal{M}(x_2, \theta)$ in (c,d); ($\theta = 1$; $b = 100$)

$\mathcal{M}(x_2, \theta)$ is zero, except for a very small blip for x_1 just less than one; however, for $n = 100$, the 2.5th percentile of $\mathcal{M}(x_2, \theta)$ is nearly quadratic for $x_1 \in (0.2, 1.8)$ with its maximum approximately 50% of 0.135.

The improvement of the adaptive optimal with sample size is impressive, particularly in a neighborhood of ± 0.8 of $\operatorname{argmax}_x \{\mathcal{M}(x_1, \theta)\} = 1$. Convergence of $-\frac{d^2}{d\theta^2} \log f(y_2|x_2, \theta)$ to the adaptive optimal appears to be slow.

5 Discussion

We have explored information measures for a two-stage adaptive optimal design in the context of a regression model with normal errors and exponential mean function. An exact expression for the second stage design point is obtained. The second stage design point is shown to be consistent as the cohort size tends to infinity, and asymptotically normal; also the variance of its asymptotic distribution is obtained.

Exact expressions for $-d^2 \log f(y_i|x_i, \theta)/d\theta^2$ and $\mathcal{M}(x_i, \theta)$, $i = 1, 2$, are given. Values of $-d^2 \log f(y_2|x_2, \theta)/d\theta^2$ are shown by (7) to fluctuate randomly, asymmetrically, around $\mathcal{M}(x_2, \theta)$, yet to converge to $\mathcal{M}(x_2, \theta)$ as $n \rightarrow \infty$. Efron and Hinkley (1978) and Lindsay and Li (1997) argue that the observed information is to be preferred over Fisher's information, but our simulations call their argument into question. Fisher's information equals $E(\mathcal{M}(x_2, \theta))$ which may be obtained numerically. A simple large sample approximation of Fisher's information is given. Our illustration suggests $\mathcal{M}(x_2, \theta)$ is converging to the asymptotic value of Fisher's information from below.

Acknowledgements Thanks to the referees whose comments led to some significant corrections in and improvements to this paper. Special thanks to Dr. Valerii V. Fedorov for suggesting this topic.

References

- Box, G. and W. Hunter (1963). Sequential design of experiments for nonlinear models. In *Proceedings of IBM Scientific Computing Symposium on Statistics*, Yorktown Heights, New York.
- Chernoff, H. (1953). Locally optimal designs for estimating parameters. *Annals of Mathematical Statistics* 24, 586–602.
- Dragalin, V., V. Fedorov, and Y. Wu (2008). Adaptive designs for selecting drug combinations based on efficacy-toxicity response. *Journal of Statistical Planning and Inference* 2, 352–373.
- Efron, B. and D. Hinkley (1978). Assessing the accuracy of the maximum likelihood estimate: observed versus Fisher information (with discussion). *Biometrika* 65, 457–483.
- Hall, P. and C. Heyde (1980). *Martingale Limit Theory and its Application*. New York: Academic Press.
- Lindsay, B. and B. Li (1997). On second order optimality of the observed Fisher information. *Annals of Statistics* 25, 2172–2199.
- Silvey, S. (1980). *Optimal Design: An Introduction to the Theory for parameter Estimation*. London: Chapman and Hall.
- White, L. (1975). *The Optimal Design of Experiments for Estimation of Nonlinear Models*. Ph. D. Dissertation, University of London.

Index

- A_L -optimality, 219
- A-optimality, 121, 199
- accurate interpolator, 125
- adaptive design, 47, 69, 91, 165, 166, 229
- adjusted p value, 160
- analysis of data, 15, 111
- ANOVA, 49, 158, 198
- approximate design, 61
- approximation to integral, 206
- Arrhenius model (modified), 173
- asymptotic properties, 166
- automatic defect detection, 194

- barycentric interpolation, 122
- Bayesian
 - approach, 99
 - design, 118
 - estimation, 149, 166
 - viewpoint, 173, 223
- beta distribution, 126
- beta regression, 60
- biased coin design, 22
- binary response, 41
- bivariate model, 182
- bivariate probit, 65
- black ball, 84
- Bonferroni, 160
- bootstrap, 53
- Bradley-Terry model, 217
- breakdown point, 138

- centre point, 105
- Chebyshev polynomials, 121
- chemical kinetics, 173
- clinical trial, 17, 41, 198
 - multicentre, 2
- combination drug therapy, 157

- compositional data, 57
- compound design, 20
- computational cost, 95
- computer experiments, 89, 121, 149, 189, 221
- concomitant variable, 17
- conditioning argument, 41
- confounded effect, 105
- consecutive first-order reactions, 61
- consistency, 69, 166, 222
- control theory, 96
- convergence, 229
- convex combination of designs, 11
- coordinate exchange algorithm, 108
- correlated errors, 42, 145, 174
- correlation function, 151
- cost, 29
- cost-based normalization, 76, 167
- covariance
 - function, 222
 - kernel, 145
 - structure, 174
- cubic model, 12

- D_1 -optimality, 11
- D_A (D_L) optimality, 18
- D_L (D_A) optimality, 219
- D_S -optimality, 9, 121, 206
- D-efficiency, 106, 114
- d-fullness, 138
- D-optimality, 108, 121, 165, 199
 - analytical peculiarity, 179
 - Bayesian, 118
 - correlated errors, 174
 - exact, 101
 - local, 59, 68, 76, 101, 114, 166, 174, 183, 205
- design construction, 210

- design measure, 114
 - non-standard interpretation, 145
- desirability function, 26
- deterministic model, 153, 193
- deterministic strategy, 91
- Dirichlet
 - distribution, 63
 - model, 59
- discrimination between models, 9, 206
- dose finding, 181
- dose-response model, 114, 157, 186
- dropout, 6
- drug development, 67, 157
- drug interaction, 198

- E-optimality, 121
- efficacy, 66
- EOHSA, 158
- equidistant design, 176
- equivalence theorem, 14, 115, 208
- ethics, 67, 82, 85, 114
- exchange algorithm, 69
- expected improvement algorithm, 89, 222
- expensive, 221

- factorial design, 30, 158
- Farlie-Gumbel-Morgenstern distribution, 183
- finite horizon optimization, 91
- finite-element method, 221
- finite-time strategy, 92
- first-order algorithm, 69, 76
- fixed effect, 205
- food tasting experiment, 214
- formally incorrect model, 75
- fractional-factorial design, 105

- gamma distribution, 2
- Gaussian process, 89, 149
- gene expression, 197
- generalized exponential model, 173
- generalized linear model, 139
- global ranking, 51
- goodness-of-fit, 191

- heteroscedastic errors, 131
- high-throughput screening, 158
- higher-order model, 13

- identifiability, 139, 140
- identifiability parameter, 139
- imbalance, 2, 18
- inferential viewpoint, 18
- information, 230

- information matrix, 66, 77, 99, 115, 151, 174, 184, 207
 - approximate, 101, 129, 147
- initial design point, 231
- intelligence
 - measurement of, 97
- inverse-linear grid, 77
- item-response model, 97
- Ito (Îto), 79

- KL-optimality, 206
- kriging, 90, 121, 222
- Kullback-Leibler distance, 16

- lack of fit, 35, 105
- large class of alternatives, 37
- latent variable, 215
- Latin hypercube, 190
- least trimmed squares (LTS), 137
- likelihood, 150
 - trimmed, 138
- linear logistic test model, 98
- LOF-optimality, 36
- logistic normal model, 62
- logistic regression, 205, 215
 - five-parameter model, 113
 - random effect, 207
- lognormal model, 135
- logratio transformation, 58
- loss, 18
 - of design efficiency, 19

- Matérn covariance, 90, 225
- maximin, 34
- maximum likelihood estimate, 99, 151, 179, 214
- maximum tolerated dose, 181
- measurement model, 76
- methane, 173
- microarray, 197
- misspecified model, 117
- mixed correlated responses, 65
- mixed-effect model, 133
- model checking, 34
- moderate sample size, 71
- multi-factor interactions, 200
- multi-response optimization, 25
- multiplicative algorithm, 217
- multivariate ranking, 49

- nested model, 206
- non-centrality parameter, 10
- non-linear model, 130, 137, 189
 - parameterization, 116

- partially, 174
- non-unique design, 9
- nugget effect, 149, 174
- numerical integration, 67

- one-compartment model, 74
- optimal allocation proportion, 43
- optimum weights, 213
- Ornstein-Uhlenbeck process, 175
- orthogonal array, 106
- orthogonal projection, 190
- outlier, 137

- paired comparison, 213
- pairwise multiple comparison, 49
- parabolic design, 176
- parametric (ϕ_p) optimality, 199
- parametric weights, 27
- patient recruitment, 8
- penalty function, 67, 168
- pharmacokinetic model, 74, 129
- Phi (Φ) optimality, 18
- Plackett-Burman design, 107
- Poisson-gamma model, 1
- pollutant, 221
- population model, 76
- positive trajectory, 75
- power, 4, 15, 82

- quadratic effect, 110
- quadratic model, 10

- radial basis, 121
- radial scanning statistic, 190
- random allocation, 82
- random effect, 129, 205
- random field, 146
- randomization, 1
 - permuted block, 3
 - rule, 17
 - stratified, 4
- Rasch model, 98
- resolution III design, 105
- response surface design, 25, 106
- response-adaptive design, 81, 85
- restricted parameter space, 143
- risk ratio, 46
- Robbins-Monro, 171

- robust design, 25
- robustness, 137
- rule-based test, 97
- Runge phenomenon, 122

- sample size, 4, 7, 53
- sampling rate, 168
- score function, 199
- screening design, 105, 157
- semi-Bayesian spirit, 102
- sequential design, 81, 118, 123, 186, 222
- sequential interpolation, 122
- sequential test, 87
- simultaneous optimization, 30
- space-filling design, 176, 189
- spacings, 192
- spline, 121
- split-plot design, 25
- splits, 77
- standardized response, 26
- stochastic differential equation, 73
- stratum, 1
- surrogate model, 221

- T-optimality, 9
- three + three (3 + 3) design, 181
- toxicity, 66, 113, 182
- transform both sides, 58
- treatment arm, 1
- trend test, 159
- trigonometric reparameterisation, 11
- two-factor interaction, 105
- two-stage design, 68, 69, 230

- uniform design
 - continuous, 33
 - exact, 33
 - random, 190
- universal kernel, 226
- urn, 82
- utility function, 67

- virtual noise, 146
 - zero, 147

- weak-heredity principle, 15
- Wiener process, 74
- within-subject variability, 73