

Springer Series in Materials Science 225

Turab Lookman
Francis Alexander
Krishna Rajan *Editors*

Information Science for Materials Discovery and Design

 Springer

Springer Series in Materials Science

Volume 225

Series editors

Robert Hull, Charlottesville, USA

Chennupati Jagadish, Canberra, Australia

Richard M. Osgood, New York, USA

Jürgen Parisi, Oldenburg, Germany

Tae-Yeon Seong, Seoul, Korea, Republic of (South Korea)

Shin-ichi Uchida, Tokyo, Japan

Zhiming M. Wang, Chengdu, China

The Springer Series in Materials Science covers the complete spectrum of materials physics, including fundamental principles, physical properties, materials theory and design. Recognizing the increasing importance of materials science in future device technologies, the book titles in this series reflect the state-of-the-art in understanding and controlling the structure and properties of all important classes of materials.

More information about this series at <http://www.springer.com/series/856>

Turab Lookman · Francis J. Alexander
Krishna Rajan
Editors

Information Science for Materials Discovery and Design

 Springer

Editors

Turab Lookman
Theoretical Division
Los Alamos National Laboratory
Los Alamos, NM
USA

Francis J. Alexander
Computer and Computational Sciences
Division
Los Alamos National Laboratory
Los Alamos, NM
USA

Krishna Rajan
Department of Materials Design
and Innovation
University at Buffalo—The State University
of New York
Buffalo, NY
USA

ISSN 0933-033X

ISSN 2196-2812 (electronic)

Springer Series in Materials Science

ISBN 978-3-319-23870-8

ISBN 978-3-319-23871-5 (eBook)

DOI 10.1007/978-3-319-23871-5

Library of Congress Control Number: 2015952059

Springer Cham Heidelberg New York Dordrecht London

© Springer International Publishing Switzerland 2016

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, 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.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

Springer International Publishing AG Switzerland is part of Springer Science+Business Media
(www.springer.com)

Preface

Accelerating materials discovery has been the theme of a number reports from the Department of Energy's Office (DOE) of Basic Energy Science (BES), the National Science Foundation (NSF), the National Academies and other government agencies and professional societies. As a driver for accelerating materials discovery, the Materials Genome Initiative, announced by the President, is part of a bold plan to boost US manufacturing over the next few decades by halving the time it takes to discover and design new materials. In this plan, accelerating discovery relies on using in the material sciences large databases, computation, mathematics, and information science in a manner similar to the way that they were used to make the Human Genome Initiative a success for the biological sciences. Novel approaches are therefore being called for that can explore the enormous phase space presented by complex materials and processes. If we are to achieve the desired performance gains, then we must have a predictive capability that can guide experiments and computations in the most fruitful directions by reducing the possibilities that need to be tried.

Despite advances in computational and experimental techniques to generate large volumes of data to screen the vast search space, it is clear that the outstanding challenge remains to integrate information-theoretic tools and materials knowledge, in the form of constraints imposed by theory, to develop, robust, predictive tools for materials design and discovery. The rapidly emerging field of materials informatics provides the critical methodology that enables the discovery, identification, and harnessing of the materials "genes" for accelerated materials discovery and design.

We provide in this book a collection of articles in this nascent field which integrates contributions from the information sciences and materials communities. The collection is partly derived from a workshop held at Santa Fe, New Mexico, February 4–7, 2014 that was organized by the editors and sponsored with support from the Centers for Nonlinear Studies and Information Science and Technology at Los Alamos National Laboratory and National Science Foundation (Grant #: 13-07811). It outlines challenges and opportunities in the use of information-theoretic tools and

evaluates the state of the art on a number of materials-motivated problems. Presented are contrasting but complementary approaches, such as those based on high-throughput calculations or experiments, as well as data-driven discovery, together with the merits and challenges of machine-learning and statistical inference methods to accommodate searches within a high dimensional feature space.

The book is organized into *three* parts. In the *first* part, following a perspective of the state of the art in materials design and discovery, Chaps. 2–6 focus largely on information-theoretic tools and how they apply to specific materials problems. Chaps. 2 and 3 discuss how aspects of decision theory within a Bayesian framework can be used for optimal experimental design. In particular, Chap. 2 discusses how to decide on the best pair of experiments for inferring the parameters of a given model, as well as how to choose an experiment to distinguish between competing models. Chapter 3 discusses strategies based on methods for global optimization for choosing the next experiment to find a material with a desired property. Proceeding from problems involving regression to those requiring classification, Chap. 4 focuses on Bayesian methods for classifying objects, especially in the limit of small samples where classifier design procedures, which work well with large samples, can have problems when data is limited. The first part of this monograph is concluded with Chaps. 5 and 6 which deal with different aspects of clustering. Chapter 5 considers the effectiveness of data visualization algorithms that look for groupings of features and materials. Chapter 6 discusses how community detection, studied in statistical physics, can be used to partition a complex system into decoupled subsets at different spatial and temporal scales.

The focus of the *second* part of the book, Chaps. 7–12, is the application of informatics tools to materials science problems. Chapter 7 discusses how parameters in the additive manufacturing process may be constrained by combining simulations and experiments using feature selection and data-driven models. Learning from high-throughput data generated from electronic structure calculations is the emphasis of Chaps. 8–11. Techniques such as principal component analysis (PCA), support vector regression (SVR), partial least squares, and Kriging using Gaussian process modeling suggest new features and materials with specified properties. Chapter 8 shows how suitable dopants in an oxide may be identified for increasing water-splitting processes. Applications in Chap. 9 include the discovery of cathode materials for lithium-ion batteries and thermoelectrics. Chapter 10 focuses on the layered compounds known as MAX phases, and Chap. 11 discusses *ab initio* methods and applied crystallography tools for descriptor development to establish structure–property relationships. Chapter 12 describes hybrid methods that integrate statistical learning techniques, to extract features from the density of states for predicting elastic properties, such as bulk modulus and yet unexplored chemistries.

The *third* and final part, Chaps. 13 and 14, discusses high-throughput experiments, which generate large amounts of data. With appropriate characterization tools, the idea is to quickly identify the subspace of the large parameter space where a new compound with desired properties may be found. Such experiments, together with informatics tools, provide opportunities for “combinatorial materials science.”

Chap. 13 provides a review in the context of multifunctional materials, and Chap. 14 incorporates aspects of informatics with a focus on solar fuel applications and multicomponent oxide catalysts.

The book is aimed at an interdisciplinary audience as the subject spans aspects of statistics, computer science, and materials science and will be of timely appeal to those interested in learning about this emerging field. We are grateful to all the authors for their articles as well as their support of the editorial process.

Los Alamos, USA

Los Alamos, USA

Buffalo, NY

Turab Lookman

Francis J. Alexander

Krishna Rajan

Contents

Part I Data Analytics and Optimal Learning

1 A Perspective on Materials Informatics: State-of-the-Art and Challenges	3
T. Lookman, P.V. Balachandran, D. Xue, G. Pilania, T. Shearman, J. Theiler, J.E. Gubernatis, J. Hogden, K. Barros, E. BenNaim and F.J. Alexander	
1.1 Introduction	4
1.2 Statistical Inference and Design: Towards Accelerated Materials Discovery	5
1.3 Progress and Concluding Remarks	9
References	11
2 Information-Driven Experimental Design in Materials Science	13
R. Aggarwal, M.J. Demkowicz and Y.M. Marzouk	
2.1 Introduction	13
2.2 The Tools of Optimal Experimental Design	15
2.2.1 Bayesian Inference	15
2.2.2 Information Theoretic Objectives	16
2.2.3 Computational Considerations	18
2.3 Examples of Optimal Experimental Design	20
2.3.1 Film-Substrate Systems: Design for Parameter Inference	21
2.3.2 Heterophase Interfaces: Design for Model Discrimination	29
2.4 Outlook	37
References	39

3	Bayesian Optimization for Materials Design	45
	Peter I. Frazier and Jialei Wang	
3.1	Introduction	45
3.2	Bayesian Optimization	47
3.3	Gaussian Process Regression.	47
3.3.1	Choice of Covariance Function.	49
3.3.2	Choice of Mean Function.	51
3.3.3	Inference	51
3.3.4	Inference with Just One Observation	53
3.3.5	Inference with Noisy Observations	54
3.3.6	Parameter Estimation.	57
3.3.7	Diagnostics	58
3.3.8	Predicting at More Than One Point.	60
3.3.9	Avoiding Matrix Inversion.	61
3.4	Choosing Where to Sample	61
3.4.1	Expected Improvement	62
3.4.2	Knowledge Gradient	65
3.4.3	Going Beyond One-Step Analyses, and Other Methods	68
3.5	Software.	69
3.6	Conclusion	69
	References	73
4	Small-Sample Classification.	77
	Lori A. Dalton and Edward R. Dougherty	
4.1	Introduction	77
4.2	Classification	78
4.3	Error Estimation	82
4.4	Validity	84
4.5	MMSE Error Estimation.	87
4.6	Optimal Bayesian Classification	90
4.7	The Gaussian Model	91
4.8	Optimal Bayesian Classifier in the Gaussian Model	94
4.9	Concluding Remarks	97
	References	99
5	Data Visualization and Structure Identification	103
	J.E. Gubernatis	
5.1	Introduction	103
5.2	Theory	104
5.3	Results.	106
5.3.1	The Piezo Data.	107
5.3.2	The Pls Data	108
5.3.3	The Tree Data	108

5.4	Concluding Remarks	109
	References	113
6	Inference of Hidden Structures in Complex Physical Systems by Multi-scale Clustering	115
	Z. Nussinov, P. Ronhovde, Dandan Hu, S. Chakrabarty, Bo Sun, Nicholas A. Mauro and Kisor K. Sahu	
6.1	The General Problem	116
6.2	Ensemble Minimization	117
6.3	Community Detection and Data Mining	118
6.4	Multi-scale Community Detection	121
6.5	Image Segmentation	123
6.6	Community Detection Phase Diagram	126
6.7	Casting Complex Materials and Physical Systems as Networks	128
6.8	Summary	133
	References	135
 Part II Materials Prediction with Data, Simulations and High-throughput Calculations		
7	On the Use of Data Mining Techniques to Build High-Density, Additively-Manufactured Parts	141
	Chandrika Kamath	
7.1	Introduction	141
	7.1.1 Additive Manufacturing Using Laser Powder-Bed Fusion	142
7.2	Optimizing AM Parts for Density: The Current Approach	142
7.3	A Data Mining Approach Combining Experiments and Simulations	144
	7.3.1 Using Simple Simulations to Identify Viable Parameters	145
	7.3.2 Using Simple Experiments to Evaluate Simulation Results	150
	7.3.3 Determining Density by Building Small Pillars	152
7.4	Experimental Results	154
7.5	Summary	154
	References	154
8	Optimal Dopant Selection for Water Splitting with Cerium Oxides: Mining and Screening First Principles Data.	157
	V. Botu, A.B. Mhadeshwar, S.L. Suib and R. Ramprasad	
8.1	Introduction	158
8.2	Screening Framework	160

8.3	First Principles Studies	160
8.3.1	Methods and Models	160
8.3.2	Enforcing the 3-Step Criteria	161
8.4	Data Analysis	164
8.4.1	Principal Component Analysis	165
8.4.2	Random Forest	166
8.5	Summary and Outlook	168
	References	168
9	Toward Materials Discovery with First-Principles Datasets and Learning Methods	173
	Isao Tanaka and Atsuto Seko	
9.1	Introduction	173
9.2	High Throughput Screening of DFT Data—Cathode Materials of Lithium ion Batteries	175
9.3	Combination of DFT Data and Machine Learning I—Melting Temperatures	177
9.4	Combination of DFT Data and Machine Learning II—Lithium ion Conducting Oxides	182
9.5	Combination of DFT Data and Machine Learning III—Thermoelectric Materials	185
	References	186
10	Materials Informatics Using <i>Ab initio</i> Data: Application to MAX Phases	187
	Wai-Yim Ching	
10.1	Introduction	187
10.2	MAX Phases: A Unique Class of Material	189
10.3	Applications of Materials Informatics to MAX Phases	191
10.3.1	Initial Screening and Construction of the MAX Database	191
10.3.2	Representative Results on Mechanical Properties and Electronic Structure of MAX	192
10.3.3	Classification of Descriptors from the Database and Correlation Among Them	197
10.3.4	Verification of the Efficacy of the Materials Informatics Tools	198
10.4	Further Applications of MAX Data	201
10.4.1	Lattice Thermal Conductivity at High Temperature	201
10.4.2	Universal Elastic Anisotropy in MAX Phases	203
10.5	Extension to Other Materials Systems	205
10.5.1	MAX-Related Systems, MXenes, MAX Solid Solutions, and Similar Layered Structures	205
10.5.2	CSH-Cement Crystals	206

10.5.3	Extension to Other Materials Systems: Bulk Metallic Glasses and High Entropy Alloys	209
10.6	Conclusions	210
	References	211
11	Symmetry-Adapted Distortion Modes as Descriptors for Materials Informatics	213
	Prasanna V. Balachandran, Nicole A. Benedek and James M. Rondinelli	
11.1	Introduction	213
11.2	Distortion Modes as Descriptors	214
11.3	Perovskite Nickelates	216
11.3.1	Statistical Correlation Analysis	217
11.3.2	Principal Component Analysis (PCA)	218
11.4	Summary	220
	References	221
12	Discovering Electronic Signatures for Phase Stability of Intermetallics via Machine Learning	223
	Scott R. Broderick and Krishna Rajan	
12.1	Introduction	223
12.2	Informatics Background and Data Processing	224
12.3	Informatics-Based Parameterization of the DOS Spectra	228
12.4	Identifying the Bulk Modulus Fingerprint	233
12.5	Summary	237
	References	237
Part III Combinatorial Materials Science with High-throughput Measurements and Analysis		
13	Combinatorial Materials Science, and a Perspective on Challenges in Data Acquisition, Analysis and Presentation	241
	Robert C. Pullar	
13.1	Combinatorial Materials Science—20 Years of Progress?	242
13.2	Combinatorial Materials Synthesis	247
13.3	High-Throughput Measurement and Analysis	254
13.4	Data Analysis and Presentation	261
	References	267
14	High Throughput Combinatorial Experimentation + Informatics = Combinatorial Science	271
	Santosh K. Suram, Meyer Z. Pesenson and John M. Gregoire	
14.1	Tailoring Material Function Through Material Complexity: The Utility of High Throughput and Combinatorial Methods	272
14.2	Materials Datasets as an Instance of Big Data	272

14.3	High Throughput Experimental Pipelines: The Example of Solar Fuels Materials Discovery	275
14.4	An Illustrative Dataset: Ni-Fe-Co-Ce Oxide Electrocatalysts for the Oxygen Evolution Reaction	276
14.5	Automating Sample Down-Selection for Maximal Information Retention: Clustering by Composition-Property Relationships	277
14.5.1	Down-Selection for Maximal Information Content	279
14.5.2	Information-Theoretic Approach	280
14.5.3	Genetic Programming Based Clustering	282
14.5.4	Calculating Membership	283
14.5.5	Application to a Synthetic Library.	284
14.5.6	Experimental Dataset.	285
14.6	The Simplex Sample Space and Statistical Analysis of Compositional Data	286
14.6.1	The Closure Effects—Induced Correlation	288
14.6.2	Illustrative Example	289
14.6.3	Sub-Compositional Coherence	290
14.6.4	Principled Analysis of Compositional Data.	290
14.6.5	Composition Spread and Distances	292
14.6.6	Interpolation of Compositional Data: Composition Profiles from Sputtering.	294
14.7	Summary and Conclusions	296
	References	297
	Index	301

Contributors

Raghav Aggarwal Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA

F.J. Alexander CCS Division, Los Alamos National Laboratory, Los Alamos, USA

Prasanna V. Balachandran Theoretical Division, Los Alamos National Laboratory, Los Alamos, USA

K. Barros Theoretical Division, T-1, Los Alamos National Laboratory, Los Alamos, USA

E. BenNaim Theoretical Division, Los Alamos National Laboratory, Los Alamos, USA

Nicole A. Benedek Department of Materials Science and Engineering, Cornell University, Ithaca, USA

V. Botu Department of Chemical and Biomolecular Engineering, University of Connecticut, Storrs, CT, USA

Scott R. Broderick Department of Materials Design and Innovation, University at Buffalo—The State University of New York, Buffalo, NY, USA

S. Chakrabarty Department of Physics, Indian Institute of Science, Bangalore, India

Wai-Yim Ching Curators Professor of Physics, Kansas City, MO, USA

Lori A. Dalton The Ohio State University, Columbus, OH, USA

M.J. Demkowicz Department of Materials Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA

Edward R. Dougherty Texas A&M University, College Station, TX, USA

Peter I. Frazier School of Operations Research & Information Engineering, Cornell University, Ithaca, NY, USA

John M. Gregoire Joint Center for Artificial Photosynthesis, California Institute of Technology, Pasadena, CA, USA

J.E. Gubernatis Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM, USA

J. Hogden CCS Division, CCS-3, Los Alamos National Laboratory, Los Alamos, USA

Dandan Hu Washington University in St. Louis, St. Louis, MO, USA

Chandrika Kamath Lawrence Livermore National Laboratory, Livermore, CA, USA

T. Lookman Theoretical Division, Los Alamos National Laboratory, Los Alamos, USA

Y.M. Marzouk Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA, USA

Nicholas A. Mauro North Central College, Naperville, IL, USA

A.B. Mhadeshwar Center for Clean Energy and Engineering, University of Connecticut, Storrs, CT, USA; Present Address: ExxonMobil Research and Engineering, Annandale, NJ, USA

Z. Nussinov Washington University in St. Louis, St. Louis, MO, USA; Department of Condensed Matter Physics, Weizmann Institute of Science, Rehovot, Israel

Meyer Z. Pesenson Joint Center for Artificial Photosynthesis, California Institute of Technology, Pasadena, CA, USA

G. Pilania Materials Science Division, Los Alamos National Laboratory, Los Alamos, USA

Robert C. Pullar Departamento de Engenharia de Materiais e Cerâmica/CICECO—Aveiro Institute of Materials, Universidade de Aveiro, Campus Universitário de Santiago, Aveiro, Portugal

Krishna Rajan Department of Materials Design and Innovation, University at Buffalo—The State University of New York, Buffalo, NY, USA

R. Ramprasad Institute of Materials Science, University of Connecticut, Storrs, CT, USA; Department of Materials Science and Engineering, University of Connecticut, Storrs, CT, USA

James M. Rondinelli Department of Materials Science and Engineering, Northwestern University, Evanston, USA

P. Ronhovde Findlay University, Findlay, OH, USA

Kisor K. Sahu School of Minerals, Metallurgical and Materials Engineering, Indian Institute of Technology, Bhubaneswar, India

Atsuto Seko Department of Materials Science and Engineering, Kyoto University, Kyoto, Japan

T. Shearman Program in Applied Mathematics, University of Arizona, Tucson, USA

S.L. Suib Department of Chemistry, University of Connecticut, Storrs, CT, USA; Institute of Materials Science, University of Connecticut, Storrs, CT, USA

Bo Sun Washington University in St. Louis, St. Louis, MO, USA

Santosh K. Suram Joint Center for Artificial Photosynthesis, California Institute of Technology, Pasadena, CA, USA

Isao Tanaka Department of Materials Science and Engineering, Kyoto University, Kyoto, Japan

J. Theiler ISR Division, Los Alamos National Laboratory, Los Alamos, USA

Jialei Wang School of Operations Research and Information Engineering, Cornell University, Ithaca, NY, USA

D. Xue Theoretical Division, Los Alamos National Laboratory, Los Alamos, USA

Part I
Data Analytics and Optimal Learning

Chapter 1

A Perspective on Materials Informatics: State-of-the-Art and Challenges

**T. Lookman, P.V. Balachandran, D. Xue, G. Pilania, T. Shearman,
J. Theiler, J.E. Gubernatis, J. Hogden, K. Barros, E. BenNaim
and F.J. Alexander**

Abstract We review how classification and regression methods have been used on materials problems and outline a design loop that serves as a basis for adaptively finding materials with targeted properties.

T. Lookman (✉) · P.V. Balachandran · D. Xue · J.E. Gubernatis · E. BenNaim
Theoretical Division, T-4, Los Alamos National Laboratory, Los Alamos 87545, USA
e-mail: txl@lanl.gov

P.V. Balachandran
e-mail: pbalachandran@lanl.gov

D. Xue
e-mail: xdz@lanl.gov

J.E. Gubernatis
e-mail: jg@lanl.gov

E. BenNaim
e-mail: ebn@lanl.gov

G. Pilania
Materials Science Division, MST-8, Los Alamos National Laboratory,
Los Alamos 87545, USA
e-mail: gpilania@lanl.gov

T. Shearman
Program in Applied Mathematics, University of Arizona, Tucson 85721, USA
e-mail: toby.shearman@gmail.com

J. Theiler
ISR Division, Los Alamos National Laboratory, Los Alamos 87545, USA
e-mail: jt@lanl.gov

J. Hogden
CCS Division, CCS-3, Los Alamos National Laboratory, Los Alamos 87545, USA
e-mail: hogden@lanl.gov

K. Barros
Theoretical Division, T-1, Los Alamos National Laboratory, Los Alamos 87545, USA
e-mail: kbarros@lanl.gov

F.J. Alexander
CCS Division, Los Alamos National Laboratory, Los Alamos 87545, USA
fja@lanl.gov

© Springer International Publishing Switzerland 2016
T. Lookman et al. (eds.), *Information Science for Materials
Discovery and Design*, Springer Series in Materials Science 225,
DOI 10.1007/978-3-319-23871-5_1

1.1 Introduction

There has been considerable interest over the last few years in accelerating the process of materials design and discovery. The Materials Genome Initiative (MGI) [1], Integrated Computational Materials Engineering (ICME) [2] and Advanced Manufacturing [3] initiatives have spurred considerable activity and brought new researchers into the nascent field of materials informatics which includes the accelerated design and discovery of new materials. The activity has also highlighted some of the open questions in this emerging area and our objective here is to provide a perspective of the field in terms of general problems and information science methods that have been used to study classes of materials, and point to some of the outstanding challenges that need to be addressed. We are guided here by our own recent work at the Los Alamos National Laboratory (LANL).

One of the earliest-studied problems in modern materials informatics relates to the classification of AB solids into their stable crystal structures, based on key attributes of the chemistry and properties of the individual A and B constituents. The emphasis was on finding features that can give rise to easily visualized two-dimensional structural maps by “drawing” boundaries between classes. The problem was first studied in the 1960s [4] but Chelikowski and Phillips [5], studying the same problem in 1978, recognized the connections to information science. Realizing that energy differences between structures were rather small, they observed that “from the point of view of information theory, ...the available structural data already contain a great deal of information: about 120 bits, in the case of the AB octet compounds. Thus one can reverse the problem, and attempt to extract from the available data quantitative rules for chemical bonding in solids.” They realized that suitable combinations of orbital radii of the individual A and B atoms were appropriate features for predicting the crystal structure of the AB solids. Over the last few years, this problem has been revisited with a variety of machine learning methods (decision trees, support vector machines, gradient boosting, etc.) [6–8] and there have been a number of studies that have classified different materials classes, such as perovskites [9]. Feature selection from data remains a fundamental exercise and here principal component analysis and correlation maps have been widely employed. Recently, high-throughput approaches have been utilized to form combinations of features from a given set and then certain key combinations are down-selected [6].

The problem of materials design is about predicting the composition and processing of materials with a more desired targeted property and therefore involves regression that leads to an inference model from training data. For example, for ferroelectrics one may wish to discover lead-based or lead-free piezoelectrics with a high transition temperature or high piezoelectric coefficient. For shape memory alloys, one may seek compounds with reduced dissipation or low hysteresis. Typically, such materials are usually found in an Edisonian fashion using intuition and time-consuming trial and error. In recent years, theory has become powerful enough to predict very accurately some material characteristics, for example, *ab initio* calculations predict elastic constants, inter-atomic distances, crystal structure, polarization, etc. However, the parameter space is just too large and there are too many

possibilities, and even if nature rules out many of the possible combinations, the numbers are still staggering. Moreover, physical and chemical constraints make the realization of many theoretically possible materials impossible. Thus, one needs to successively improve or learn from available data candidate materials for further experiments and calculations. Recently, a number of studies have utilized regression methods to predict materials with given properties. However, most research in materials design has been based on high throughput approaches using electronic structure calculations. Typically, a large database is assembled with calculated properties and this is successively screened for materials with desired properties. High-throughput experiments have also been undertaken more recently to screen for candidate materials for further experiments [10, 11]. When it comes to multicomponent alloys or solid solutions, these methods have limitations. Moreover, very few studies have combined statistical inference with the high-throughput approach.

1.2 Statistical Inference and Design: Towards Accelerated Materials Discovery

Figure 1.1 illustrates our vision for the overall materials informatics/design problem. This shows a feedback loop that starts with the available assembled data (box 5), which may be obtained from multiple sources, including experiments or calculations. Materials knowledge (box 1) is then key in selecting the features and prescribing the constraints amongst them. Our aim is to train a statistical inference model that estimates the property (regression) or classification label with associated uncertainties (box 2). Classification models answer categorical questions: Is a compound stable?

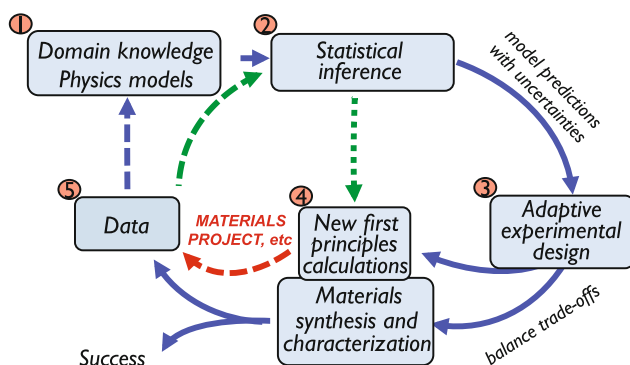


Fig. 1.1 Statistical Inference and design: A feedback loop to find a material with a desired targeted property. Prior or domain knowledge, including features, provide input to an inference model that predicts a label or a property with uncertainty. An experimental design or decision making module balances trade off between exploiting information or further exploring the high dimensional search space where the desired material may be found. A material is suggested for experimentation or calculation and the process repeats itself incorporating updated information

Is it a piezoelectric? What is its crystal symmetry? Regression models produce numerical estimates: What is the material's piezoelectric coefficient? What is its transition temperature? Because there usually is a limited quantity of training data, and because the space of possibilities is so high-dimensional, incorporation of domain knowledge is of potentially great value. Here explicitly Bayesian approaches, in which this knowledge is coded into prior probability distributions, and more traditional machine learning algorithms (such as support vector machines) in which case the domain knowledge could be incorporated as constraints or folded into the kernel design, become important [12].

Much existing work is essentially based on going from box 1 to box 4 in Fig. 1.1. A case in point are projects such as the Materials Project [13] and AFLOWLIB [14] focused on establishing databases using electronic structure calculations to make predictions. However, there are a few studies that use inference to make predictions. Examples include predictions of melting temperature [7, 8, 15] or piezoelectrics with high transition temperatures [16]. The search for piezoelectrics serves as a good example to contrast the two approaches. Extensive *ab initio* calculations were performed on a chemical space represented by $63^2 = 3969$ possible perovskite ABO_3 (up to Bi but excluding a few such as H and inert gases) end structures [17]. The number of possibilities were filtered down to 49 by discarding compounds that are nonmetallic or whose structures have small energy barriers to distortions across the morphotropic phase boundary (MPB) according to preset values. Almost no optimization or learning tools are used other than what may be involved in seeking an optimal minimum energy solution at zero temperature. All the physics is contained in this first-principles calculation, and we are not aware if any of this group's predictions of piezoelectricity have been verified experimentally. On the other hand, the approach of Balachandran et al. [16] on the same type of problem was to focus on a given subclass of piezoelectrics (e.g. Bi based) with known crystallographic and experimental data and use off-the-shelf inference tools to obtain candidates with high transition temperatures and that were formable. The tools included principal component analysis (PCA) for dimensionality reduction, partial least squares (PLS) regression for predicting transition temperatures and recursive partitioning (or decision trees) with a metric such as Shannon entropy for classification. The training data sets for PCA or regression studies were rather small (about 20 data points, 30 features) but data sets with 350 data points were also used to identify stable/formable perovskite compounds. Two new compounds were predicted, of which one has been synthesized [18], with the predicted transition temperature differing by 30–40%. However, a key element lacking is the issue of uncertainties in predictions.

In Fig. 1.2, we demonstrate using an example, where we have used bootstrap methods (i.e. sampling with replacement) to estimate prediction uncertainties. Here, we took the same Bi-based piezoelectrics data set as that utilized in the work of Balachandran et al. [16] We generated a large number of bootstrapped samples (as opposed to using just one in the earlier work of Balachandran et al.) and utilized support vector regression (SVR) for predicting the Curie temperature (T_C). Our results with uncertainties are shown in Fig. 1.2. On average, we obtained a standard deviation of 37°C from the mean value of predicted T_C . More importantly, we also

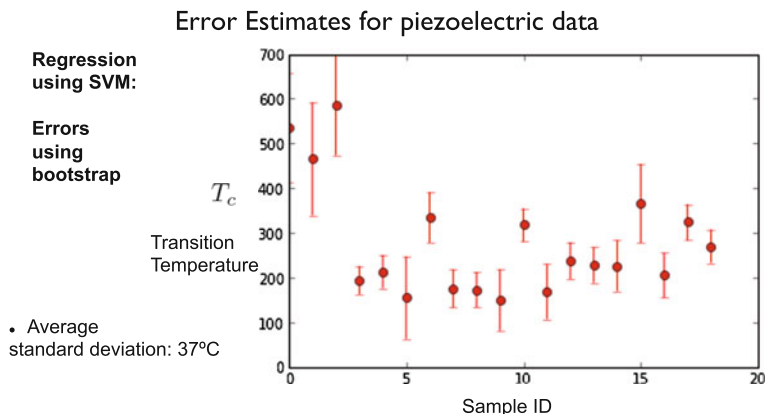


Fig. 1.2 Predictions using support vector regression (SVR) with uncertainties from bootstrap method. The piezoelectric data set of Bi-based PbTiO_3 solid solutions was used for machine learning. T_C (in $^\circ\text{C}$, y-axis) is the predicted ferroelectric Curie temperature at the morphotropic phase boundary (MPB). We use the SVR model and predicted T_C for two new compounds, $\text{BiLuO}_3\text{-PbTiO}_3$ and $\text{BiTmO}_3\text{-PbTiO}_3$, to be 552.5 ± 79 and 564.2 ± 97 $^\circ\text{C}$, respectively. Experimentally, T_C for $\text{BiLuO}_3\text{-PbTiO}_3$ was measured as 565 $^\circ\text{C}$ [18]

predicted the T_C for two new compounds, $\text{BiLuO}_3\text{-PbTiO}_3$ and $\text{BiTmO}_3\text{-PbTiO}_3$, to be 552.5 ± 79 and 564.2 ± 97 $^\circ\text{C}$, respectively with 95 % confidence. Experimentally, T_C for $\text{BiLuO}_3\text{-PbTiO}_3$ was measured as 565 $^\circ\text{C}$ [18], in close agreement with the current results from SVR. On the other hand, PLS predicted the T_C for $\text{BiLuO}_3\text{-PbTiO}_3$ to be 705 $^\circ\text{C}$. The merit of this example is that it shows in a rather modest manner that the informatics approach, even if manual and piecemeal, is potentially capable of predicting new materials.

A key aspect of our design loop is the uncertainty associated with the properties predicted from inference (box 2). These play a role in the adaptive experimental design (box 3) which suggests the next material to be chosen for further experiments or calculation (box 4) by balancing the tradeoffs between “exploration and exploitation”. That is, at any given stage a number of samples may be predicted to have given properties with uncertainties. The tradeoff is between exploiting the results by choosing to perform the next experiment on the material predicted to have the largest property or further improving the model by performing the experiment or calculation on a material where the predictions have the largest uncertainties. By choosing the latter, the uncertainty in the property is expected to (given the model, statistics) decrease, the model will probably improve and this will influence the results of the next iteration in the loop. While there is a considerable literature on error estimation methodologies, accurate and reliable error estimation with limited data is harder than simple prediction, and there is an even a stronger case for incorporating domain knowledge. [8, 19, 20]

Extracting measures of confidence, while at the same time encoding prior knowledge, is not an easy task but recent research in cancer genomics has demonstrated that increasing confidence in classification analysis built on small databases benefits

significantly from using prior knowledge [21, 22]. Prior domain knowledge constrains statistical outcomes by producing classifiers that are superior to those designed from data alone. How to use prior knowledge in classification and regression is a problem not only for materials and cancer genomics but for machine learning generally. Developing ways of constructing and using prior domain knowledge will distinguish the materials machine learning approach to classification and regression. The lesson learned from high-throughput genomics concerning classification is that, in high dimensional, small-sample settings, model-free classification is virtually impossible. The reason is that the salient property of any classifier is its error rate because the error rate quantifies its predictive capacity, which is the essential issue pertaining to scientific validity. Since the error rate must be estimated, there must be an estimation procedure and, with small samples, this procedure must be applied to the same data as that used for designing the classifier. In cancer genomics, Dalton and Dougherty [19, 20] addressed the problem by formulating error estimation as an optimization problem in a model-based framework and leads to a minimum-mean-square-error (MMSE) estimate of the classifier error. They formulate a prior probability distribution over a class of possible distributional models governing the features to be measured and the possible decisions to be made, each such model being known as a feature-label distribution. They then design a classifier from the data and an optimal MMSE error estimate is derived from the data. How well this approach will work for materials problems remains an open question.

In Figs. 1.3 and 1.4 we provide more details of our loop. Figure 1.3 shows how the loop would actually work in practice, and some of the algorithms that may be

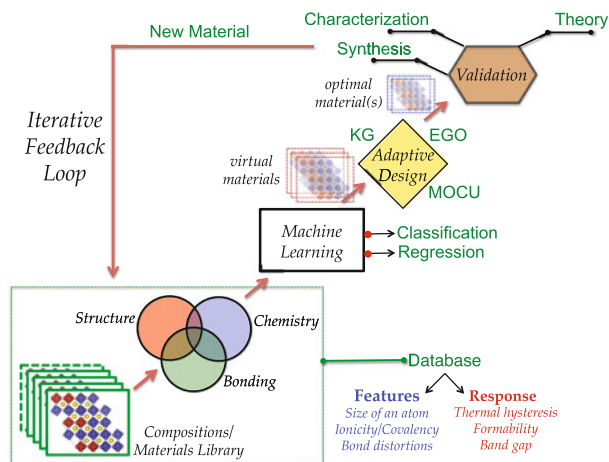


Fig. 1.3 The design loop in practice showing different stages of machine learning and adaptive design strategies with an iterative feedback loop. For completeness, we have also included experiments (synthesis and characterization), which are vital for validation and feedback. KG, EGO and MOCU stand for knowledge gradient, effective global optimization and mean objective cost of uncertainty, respectively

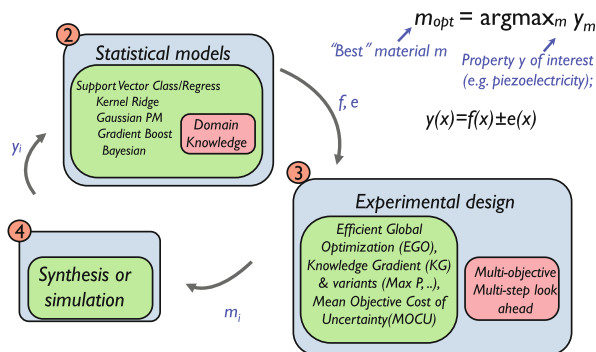


Fig. 1.4 A sub-component of our adaptive design loop showing the synergy between statistical models (box 2), experimental design (box 3) and validation (typically via experimental synthesis or simulation as shown in box 4). Statistical models use the available data to fit a regression model (f) along with an uncertainty measure (e). The experimental design component then evaluates the tradeoff between exploitation and exploration and suggests the “best” material (y_i) for validation. Here the term “best” need not correspond to a material with the optimal response. Alternatively, it refers to the choice of a material that would reduce the overall uncertainty in our model. Different statistical learning (including Bayesian learning) and adaptive design methods are given

used as part of the statistical inference and design tools, are shown in greater detail in Fig. 1.4. The green emphasize algorithms that can be utilized today and the red represent areas requiring further study and development. Design algorithms include well known exploitation-exploration strategies such as efficient global optimization (EGO) [23], and the closely related knowledge gradient(KG) [24] based on single-step look ahead.

1.3 Progress and Concluding Remarks

Our work at LANL has involved studying a number of materials problems along the lines of the approach described. These include problems involving classification learning and regression, which essentially involve an inner loop of Fig. 1.1 with boxes 2, 4 and 5. We have examined the role of features in classifying AB octet solids [8] and perovskites [9], as well as predicting new ductile RM intermetallics, where R and M are rare earth and transition metal elements, respectively [25]. These studies have suggested new features that led to better classification as well as new materials. In the case of RM intermetallics, we have shown that machine learning methods naturally uncover the functional forms that mimic most frequently used features in the literature, thereby providing a mathematical basis for feature set construction without *a priori* assumptions [25]. Our classification models (Fig. 1.5) that use orbital radii as features predicted that ScCo, ScIr, and YCd should be ductile, whereas each was previously proposed to be brittle. These results show it is possible to design

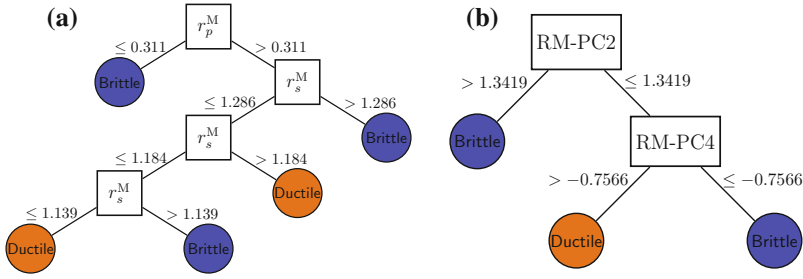


Fig. 1.5 Classification learning using decision trees to predict whether a given RM intermetallic, where R and M are rare earth and transition metal elements, respectively, is brittle and ductile. **a** Decision tree that uses the orbital radii as features and (b) Decision tree that uses the principal components (RM-PC2 and RM-PC4) that automatically extracts features in the form of the linear combinations of orbital radii. For example, RM-PC2 is defined as $-0.70r_p^M + 0.08r_s^M - 0.71r_d^M$. Features r_p^M , r_s^M and r_d^M are the p -, s - and d -orbital radii of atom-M, respectively

targeted mechanical properties in intermetallic compounds, which has significant implications for next-generation multi-component alloy discovery.

Our on-going work on multi-objective regression includes predicting functional polymers with large band gaps, as well as large dielectric constants for energy storage applications. Similarly, we are also performing high-throughput density functional theory (DFT) calculations to generate large data sets, which are subsequently mined using machine learning methods to identify new and previously unexplored candidate water splitting compounds for catalysis.

In the area of adaptive design, our focus has been on demonstrating the feedback loop of Figs. 1.1 or 1.3 with tight coupling to an “oracle”, which can be experiments (synthesis and characterization) or calculations. Specific materials studies include discovering new low thermal dissipation shape memory alloys, as well as Pb-free piezoelectric solid-solutions starting from experimental data on specific multicomponent systems. The search spaces can be well defined, for example, they can be a factor of 10^5 greater than the size of the training data. In addition, extensive databases from *ab initio* calculations become invaluable in benchmarking the various algorithms. For example, elastic moduli data for the hexagonal layered M_2AX phases consist of a library of 240 compounds. The *ab initio* data of the elastic constants and moduli were taken from the literature [26] with results well calibrated to experiments. In the M_2AX phases, X-atoms reside in the edge-connected M octahedral cages and the A atoms reside in slightly larger right prisms [27]. These M_2AX phases represent a unique family of materials with layered crystal structure and both metallic- and ceramic-like properties. We used orbital radii of M, A, and X atoms from the Waber-Cromer scale [28] as features, which include the s -, p -, and d -orbital radii for M, while the s - and p -orbital radii were used for A and X atoms. With the M_2AX data, we benchmarked our adaptive design strategy, i.e. explored different training set sizes, regressors, regressor/optimization combinations, etc., and uncovered invaluable guidelines that were eventually useful for real materials design problems.

Implementing the loop using simulation codes allows us to optimize the use of these codes in seeking a well defined set of parameters or constraints for given targeted outcomes. For example, an industry standard code for simulating semiconducting materials is APSYS (Advanced Physical Models of Semiconductor Devices). It is based on 2D/3D finite element analysis of electrical, optical and thermal properties of compound semiconductor devices, with silicon as a special case with an emphasis on band structure engineering and quantum mechanical effects. Inclusion of various optical modules allows one to configure applications involving photosensitive or light emitting diodes (LEDs). We have been recently using APSYS to investigate how to optimize the LED structure (number of quantum wells, indium concentration) of GaAs based systems for highest internal quantum efficiencies at high currents.

In summary, the use of classification and regression methods, in combination with optimization strategies, has the potential to impact discovery and design in materials science. What is needed is to establish how these tools perform on an array of materials classes with differing physics in order to distill some guiding principles for use by the materials community at large.

Acknowledgments We acknowledge funding support from a Laboratory Directed Research and Development (LDRD) DR (#20140013DR) at the Los Alamos National Laboratory (LANL).

References

1. *Materials Genome Initiative for Global Competitiveness* (2011)
2. S.R. Kalidindi, M. De Graef, Materials data science: current status and future outlook. *Ann. Rev. Mater. Res.* **45**(1), 171–193 (2015)
3. T.D. Wall, J.M. Corbett, C.W. Clegg, P.R. Jackson, R. Martin, Advanced manufacturing technology and work design: towards a theoretical framework. *J. Organ. Behav.* **11**(3), 201–219 (1990)
4. E. Mooser, W.B. Pearson, On the crystal chemistry of normal valence compounds. *Acta Crystallogr.* **12**, 1015–1022 (1959)
5. J.R. Chelikowsky, J.C. Phillips, Quantum-defect theory of heats of formation and structural transition energies of liquid and solid simple metal alloys and compounds. *Phys. Rev. B* **17**, 2453–2477 (1978)
6. L.M. Ghiringhelli, J. Vybiral, S.V. Levchenko, C. Draxl, M. Scheffler, Big data of materials science: critical role of the descriptor. *Phys. Rev. Lett.* **114**, 105503 (2015)
7. Y. Saad, D. Gao, T. Ngo, S. Bobbitt, J.R. Chelikowsky, W. Andreoni, Data mining for materials: computational experiments with AB compounds. *Phys. Rev. B* **85**, 104104 (2012)
8. G. Pilania, J.E. Gubernatis, T. Lookman, Structure classification and melting temperature prediction of octet AB solids via machine learning. *Phys. Rev. B* **91**, 124301 (2015)
9. G. Pilania, P.V. Balachandran, J.E. Gubernatis, T. Lookman, Predicting the formability of ABO₃ perovskite solids: a machine learning study. *Acta Crystallogr. B* **71**, 507–513 (2015)
10. S.M. Senkan, High-throughput screening of solid-state catalyst libraries. *Nature* **394** (6691), 350–353, 07 (1998)
11. H. Koinuma, I. Takeuchi, Combinatorial solid-state chemistry of inorganic materials. *Nat. Mater.* **3**, 429–438 (2004)
12. T. Hastie, R. Tibshirani, J. Friedman, *The Elements of Statistical Learning* (Springer, New York, 2008)

13. A. Jain, S.P. Ong, G. Hautier, W. Chen, W.D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, K.A. Persson, Commentary: the materials project: a materials genome approach to accelerating materials innovation. *APL Mater.* **1**(1) (2013)
14. S. Curtarolo, W. Setyawan, S. Wang, J. Xue, K. Yang, R.H. Taylor, L.J. Nelson, G.L. Hart, S. Sanvito, M. Buongiorno-Nardelli, N. Mingo, O. Levy, AFLOWLIB.ORG: a distributed materials property repository from high-throughput ab initio calculations. *Comput. Mater. Sci.* **58**, 227–235 (2012)
15. A. Seko, T. Maekawa, K. Tsuda, I. Tanaka, Machine learning with systematic density-functional theory calculations: application to melting temperatures of single- and binary-component solids. *Phys. Rev. B* **89**, 054303 (2014)
16. P.V. Balachandran, S.R. Broderick, K. Rajan, Identifying the inorganic gene for high-temperature piezoelectric perovskites through statistical learning. *Proc. R. Soc. A: Math. Phys. Eng. Sci.* **467**(2132), 2271–2290 (2011)
17. R. Armiento, B. Kozinsky, M. Fornari, G. Ceder, Screening for high-performance piezoelectrics using high-throughput density functional theory. *Phys. Rev. B* **84**, 014103 (2011)
18. W. Hu, Experimental search for high Curie temperature piezoelectric ceramics with combinatorial approaches. Ph.D. dissertation, Iowa State University (2011)
19. L.A. Dalton, E.R. Dougherty, Optimal classifiers with minimum expected error within a Bayesian framework—Part I: discrete and Gaussian models. *Pattern Recognit.* **46**(5), 1301–1314 (2013)
20. L.A. Dalton, E.R. Dougherty, Optimal classifiers with minimum expected error within a Bayesian framework—Part II: properties and performance analysis. *Pattern Recognit.* **46**(5), 1288–1300 (2013)
21. K.E. Lee, N. Sha, E.R. Dougherty, M. Vannucci, B.K. Mallick, Gene selection: a Bayesian variable selection approach. *Bioinformatics* **19**(1), 90–97 (2003)
22. E.R. Dougherty, A. Zollanvari, U.M. Braga-Neto, The illusion of distribution-free small-sample classification in genomics. *Curr genomics* **12**(5), 333–341 (2011)
23. D.R. Jones, M. Schonlau, W.J. Welch, Efficient global optimization of expensive black-box functions. *J. Glob. Optim.* **13**(4), 455–492 (1998)
24. W. Powell, I. Ryzhov, *Optimal Learning*, Wiley Series in Probability and Statistics (Wiley, Hoboken, 2013)
25. P.V. Balachandran, J. Theiler, J. M. Rondinelli, T. Lookman, Materials Prediction via Classification Learning *Sci. Rep.* **5**, 13285 (2015)
26. M.F. Cover, O. Warschkow, M.M.M. Bilek, D.R. McKenzie, A comprehensive survey of M_2AX phase elastic properties. *J. Phys.: Condens. Matter* **21**(30), 305403 (2009)
27. M.W. Barsoum, M. Radovic, Elastic and mechanical properties of the MAX phases. *Ann. Rev. Mater. Res.* **41**, 195–227 (2011)
28. J.T. Waber, D.T. Cromer, Orbital radii of atoms and ions. *J. Chem. Phys.* **42**(12), 4116–4123 (1965)

Chapter 2

Information-Driven Experimental Design in Materials Science

R. Aggarwal, M.J. Demkowicz and Y.M. Marzouk

Abstract Optimal experimental design (OED) aims to maximize the value of experiments and the data they produce. OED ensures efficient allocation of limited resources, especially when numerous repeated experiments cannot be performed. This chapter presents a fully Bayesian and decision theoretic approach to OED—accounting for uncertainties in models, model parameters, and experimental outcomes, and allowing optimality to be defined according to a range of possible experimental goals. We demonstrate this approach on two illustrative problems in materials research. The first example is a parameter inference problem. Its goal is to determine a substrate property from the behavior of a film deposited thereon. We design experiments to yield maximal information about the substrate property using only two measurements. The second example is a model selection problem. We design an experiment that optimally distinguishes between two models for helium trapping at interfaces. In both instances, we provide model-based justifications for why the selected experiments are optimal. Moreover, both examples illustrate the utility of reduced-order or surrogate models in optimal experimental design.

2.1 Introduction

Experiments are essential prerequisites of all scientific research. They are the basis for developing and refining mathematical models of physical reality. Experimental data are used to infer model parameters, to improve the accuracy of model-based predictions, to discriminate among competing models, to assess model validity, and

R. Aggarwal · M.J. Demkowicz (✉)
Department of Materials Science and Engineering, Massachusetts Institute of Technology,
77 Massachusetts Avenue, Cambridge, MA 02139, USA
e-mail: demkowicz@mit.edu

Y.M. Marzouk
Department of Aeronautics and Astronautics, Room 37-451, Massachusetts Institute
of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA
e-mail: ymarz@mit.edu

© Springer International Publishing Switzerland 2016
T. Lookman et al. (eds.), *Information Science for Materials
Discovery and Design*, Springer Series in Materials Science 225,
DOI 10.1007/978-3-319-23871-5_2

to improve design and decision-making under uncertainty. Yet experimental observations can be difficult, time-consuming, and expensive to acquire. Maximizing the value of experimental observations—i.e., designing experiments to be *optimal* by some appropriate measure—is therefore a critical task. Experimental design encompasses questions of where and when to measure, which variables to interrogate, and what experimental conditions to employ.

Conventional experimental design methods, such as factorial and composite designs, are largely used as heuristics for exploring the relationship between input factors and response variables. By contrast, *optimal* experimental design uses a concrete hypothesis—expressed as a quantitative model—to guide the choice of experiments for a particular purpose, such as parameter inference, prediction, or model discrimination. Optimal design has seen extensive development for linear models (where the measured quantities depend linearly on the model parameters) endowed with Gaussian distributions [5]. Extensions to nonlinear models are often based on linearization and Gaussian approximations [15, 21, 36], as analytical results are otherwise impractical or impossible to obtain. With advances in computational power, however, optimal experimental design for nonlinear systems can now be tackled directly using numerical simulation [48, 49, 64, 65, 84, 89, 93, 96].

This chapter will present an overview of model-based optimal experimental design, connecting this approach to illustrative applications in materials science—a field replete with potential applications for optimal experimentation. We will take a fully Bayesian and decision-theoretic approach. In this formulation, one first defines the *utility* of an experiment and then, taking into account uncertainties in both the parameter values and the observations, chooses experiments by maximizing an *expected utility*. We will define these utilities according to information theoretic considerations, reflecting the particular experimental goals at hand.

The evaluation and optimization of information theoretic design criteria, in particular those that invoke complex physics-based models, requires the synthesis of several computational tools. These include: (1) statistical estimators of expected information gain; (2) efficient optimization methods for stochastic or noisy objectives (since expected utilities are typically evaluated with Monte Carlo methods); and (3) reduced-order or surrogate models that can accelerate the estimation of information gain. For a simple film-substrate system, we will present an example of such a reduced-order model, derived from physical scaling principles and an “offline” set of detailed/full model simulations. This is but one example; reduced-order models constructed through a variety of techniques have practical use in a wide range of optimal experimental design applications [48].

The rest of this chapter is organized as follows. Section 2.2 will present the foundational tools of optimal experimental design, beginning with Bayesian inference and proceeding to discuss several information theoretic design criteria. It will also discuss the computational challenges presented by this formulation. Section 2.3 will illustrate the information-driven approach with two examples: optimal design for *parameter inference*, in the context of a film-substrate system; and optimal design for *model selection*, in the context of heterophase interfaces in layered metal composites. Section 2.4 will discuss open questions and topics of ongoing research.

2.2 The Tools of Optimal Experimental Design

We will formulate our experimental design criteria in a Bayesian setting. Bayesian statistics offers a foundation for inference from noisy, indirect, and incomplete data; a mechanism for incorporating multiple heterogeneous sources of information; and a complete assessment of uncertainty in parameters, models, and predictions. The Bayesian approach also provides natural links to decision theory, which we will exploit below.

2.2.1 Bayesian Inference

The essence of the Bayesian paradigm is to describe uncertainty or lack of knowledge *probabilistically*. This idea applies to model parameters, to observations, and even to competing models. For simplicity, we first describe the case of parameter inference.

Let $\theta \in \Theta \subseteq \mathbb{R}^n$ represent the parameters of a given model. We describe our state of knowledge about these parameters with a prior probability density $p(\theta)$. (For the remainder of this article, we assume that all parameter and data probability distributions have densities with respect to Lebesgue measure.) We would like to update our knowledge about θ by performing an experiment at conditions $\eta \in \mathcal{H} \subseteq \mathbb{R}^d$. η is therefore our vector of experimental design parameters. This experiment will yield observations $y \in \mathcal{Y} \subseteq \mathbb{R}^m$. The relationship between the model parameters, experimental conditions, and observations is captured by the likelihood function $p(y|\theta, \eta)$, i.e., the probability density of the observations given a *particular* choice of θ, η . The likelihood naturally incorporates a physical model of the experiment. For instance, one often has a computational model $G(\theta, \eta)$ that predicts the quantity being measured by a proposed experiment. This prediction may be imperfect, and is almost always corrupted by some observational errors. A simple likelihood then results from the additive model $y = G(\theta, \eta) + \epsilon$, where ϵ is a random variable representing measurement and model errors. If ϵ is Gaussian with mean zero and variance σ^2 , and independent of θ and η , then we have the Gaussian likelihood $p(y|\theta, \eta) \sim N(G(\theta, \eta), \sigma^2)$. More complex likelihoods describe signal-dependent noise, or include more sophisticated representations of model error (e.g., the discrepancy models of [53]).

Putting these ingredients together via Bayes' rule, we obtain the posterior probability density $p(\theta|y, \eta)$ of the parameters:

$$p(\theta|y, \eta) = \frac{p(y|\theta, \eta)p(\theta)}{p(y|\eta)}, \quad (2.1)$$

where we have assumed (quite reasonably) that the prior knowledge on the parameters is independent of the experimental design. The posterior density describes the state of knowledge about the parameters θ *after* conditioning on the result of the experiment.

The design criteria described below will formalize the intuitive idea of choosing values of η to make the posterior distribution of θ as “informed” as possible.

Many problems, whether in materials science or other domains, do not have parameter inference as an end goal. Rather than learning about parameters that appear in a single fixed model of interest, one may wish to collect data that help discriminate among competing models. For instance, different hypothesized physical mechanisms may lead to different models of a phenomenon. In this context, the Bayesian approach involves characterizing a posterior probability distribution over models. Let the model space \mathcal{M} consist of an enumerable number of competing models M_i , $i \in \{1, 2, \dots\}$. Let each model M_i be endowed with parameters $\theta_i \in \Theta_i \subseteq \mathbb{R}^{n_i}$. Then Bayes’ rule writes the posterior probability of a model M_i as:

$$P(M_i|y, \eta) = \frac{p(y|M_i, \eta)P(M_i)}{p(y|\eta)}, \quad (2.2)$$

where the *marginal likelihood* of each model (i.e., $p(y|M_i, \eta)$ for the i th model) is obtained by averaging the likelihood over the prior distribution on the model’s parameters:

$$p(y|M_i, \eta) = \int_{\Theta_i} p(y|\theta_i, \eta, M_i)p(\theta_i|M_i)d\theta_i. \quad (2.3)$$

Each model has its own parameters θ_i and its own prior $p(\theta_i|M_i)$. The marginal likelihood incorporates an automatic Occam’s razor that penalizes unnecessary model complexity [8, 66]. The effective use of the posterior distribution over models $P(M_i|y, \eta)$ can then depend on the goals at hand. For instance, one may wish to know which model is best supported by the data; in this case, one simply selects the model with the highest posterior probability, thus performing Bayesian model selection. Alternatively, if the end goal is to make a prediction that accounts for model uncertainty, one can perform Bayesian model averaging [46] by taking a linear combination of predictions from each model, weighed according to the posterior model probabilities.

2.2.2 Information Theoretic Objectives

Following a decision theoretic approach, Lindley [63] suggests that an objective for experimental design should have the following general form:

$$U(\eta) = \int_y \int_{\Theta} u(\eta, y, \theta) p(\theta, y|\eta) d\theta dy, \quad (2.4)$$

where $u(\eta, y, \theta)$ is a *utility function* and $U(\eta)$ is the *expected utility*. The utility function u should be chosen to reflect the usefulness of an experiment at conditions η , given a particular value of the parameters θ and a particular outcome y . Since

we do not know the precise value of θ and we cannot know the outcome of the experiment before it is performed, we obtain U by taking the expectation of u over the joint distribution of θ and y ; hence the name ‘expected’ utility.

The choice of utility function u reflects the purpose of the experiment. To accommodate nonlinear models and avoid restrictive distributional assumptions on the parameters or model predictions, we advocate the use of utility functions that reflect the gain in *Shannon information* in quantities of interest [42]. For instance, if the object of the experiment is *parameter inference*, then a useful utility function is the relative entropy or Kullback-Leibler (KL) divergence from the posterior to the prior:

$$\begin{aligned} u(\eta, y, \theta) = u(\eta, y) &= D_{\text{KL}}(p(\theta|y, \eta) \parallel p(\theta)) \\ &= \int_{\theta} p(\theta|y, \eta) \log \frac{p(\theta|y, \eta)}{p(\theta)} d\theta. \end{aligned} \quad (2.5)$$

Taking the expectation of this quantity over the prior predictive of the data, as in (2.4), yields a U equal to the *expected information gain* in θ . This quantity is equivalent to the mutual information [26] between the data and the parameters, $I(y; \theta)$.

Inferring parameters may not be the true object of an experiment, however. For many experiments, the goal is to improve *predictions* of some quantity Q . This quantity may depend strongly on some model parameters and weakly on others. Moreover, some model parameters might simply be “knobs” without a strict physical interpretation or meaning. In this setting, we can put $u(\eta, y, \theta) = u(\eta, y)$ equal to the Kullback-Leibler divergence evaluated from the posterior predictive distribution, $p(Q|y, \eta) = \int p(Q|\theta)p(\theta|y, \eta)d\theta$, to the prior predictive distribution, $p(Q) = \int p(Q|\theta)p(\theta)d\theta$. Taking the expectation of this utility function over the data yields $U(\eta) = I(Q; y|\eta)$, that is, the conditional mutual information between data and predictions. This quantity implicitly incorporates an information theoretic “forward” sensitivity analysis, as the experiments that are most informative about Q will automatically constrain the directions in the parameter space that strongly influence Q .

As mentioned above, another common experimental goal is *model discrimination*. From the Bayesian perspective, we wish to maximize the relative entropy between the posterior and prior distributions over models:

$$u(\eta, y) = \sum_i P(M_i|y, \eta) \log \frac{P(M_i|y, \eta)}{P(M_i)}. \quad (2.6)$$

Moving from this utility to an expected utility requires integrating over the prior predictive distribution of the data, as specified in (2.4). Since the utility function u here does not depend on the parameters θ , we simply have $U(\eta) = \int_y u(\eta, y)p(y|\eta)dy$. Because we are now considering multiple competing models, however, the prior predictive distribution is itself a mixture of the prior predictive distribution of each model:

$$p(y|\eta) = \sum_i P(M_i) p(y|M_i, \eta) = \sum_i P(M_i) \int_{\Theta_i} p(y|\theta_i, \eta, M_i) p(\theta_i|M_i) d\theta_i. \quad (2.7)$$

The resulting expected information gain in model space favors designs that are expected to focus the posterior distribution onto fewer models [75]. In more intuitive terms, we will be driven to test where we know the least and where we also expect to learn the most.

2.2.3 Computational Considerations

Evaluating expected information gain. Except in special cases (e.g., linear-Gaussian models), the expected utilities described above cannot be evaluated in closed form. Instead, the integrals in these expressions must be approximated numerically. Note that, even in the simplest case of parameter inference—with utility given by (2.5)—evaluating the posterior density of the parameters requires calculating the posterior normalizing constant, which (like the posterior distribution itself) is a function of the data y and the design parameters η . In this situation, it is convenient to rewrite the expected information gain in the parameters θ as follows:

$$\begin{aligned} U(\eta) &= \int_{\mathcal{Y}} \int_{\Theta} p(\theta|y, \eta) \log \left(\frac{p(\theta|y, \eta)}{p(\theta)} \right) d\theta p(y|\eta) dy \\ &= \int_{\mathcal{Y}} \int_{\Theta} \log \left(\frac{p(y|\theta, \eta)}{p(y|\eta)} \right) p(y|\theta, \eta) p(\theta) d\theta dy \\ &= \int_{\mathcal{Y}} \int_{\Theta} \{\log p(y|\theta, \eta) - \log p(y|\eta)\} p(y|\theta, \eta) p(\theta) d\theta dy, \quad (2.8) \end{aligned}$$

where the second equality is due to the application of Bayes' rule to the quantities both inside and outside the logarithm. Introducing Monte Carlo approximations of the evidence $p(y|\eta)$ and the outer integrals, we obtain the nested Monte Carlo estimator proposed by Ryan [84]:

$$\begin{aligned} U(\eta) &\approx \hat{U}_{N,M}(\eta) \\ &:= \frac{1}{N} \sum_{i=1}^N \left\{ \log(p(y^{(i)}|\theta^{(i)}, \eta)) - \log \left(\frac{1}{M} \sum_{j=1}^M p(y^{(i)}|\tilde{\theta}^{(i,j)}, \eta) \right) \right\}. \quad (2.9) \end{aligned}$$

Here $\{\theta^{(i)}\}$ and $\{\tilde{\theta}^{(i,j)}\}$, $i = 1 \dots N$, $j = 1 \dots M$, are independent samples from the prior $p(\theta)$, and each $y^{(i)}$ is an independent sample from the likelihood $p(y|\theta^{(i)}, \eta)$, for $i = 1 \dots N$. The variance of this estimator is approximately $A(\eta)/N + B(\eta)/(NM)$, and its bias is (to leading order) $C(\eta)/M$ [84], where A , B , and C are terms that depend only on the distributions at hand. The estimator $\hat{U}_{N,M}$ is thus biased for finite M , but asymptotically unbiased.

Analogous, though more complex, Monte Carlo estimators can be derived for the expected information gain in some predictions Q , or for the expected information gain in the model indicator M_i .

Optimization approaches. Regardless of the particular utility function u used to define U , selecting an optimal experimental design requires solving an optimization problem of the form:

$$\max_{\eta \in \mathcal{H}} U(\eta). \quad (2.10)$$

Using the Monte Carlo approaches described above, only noisy estimates (e.g., $\hat{U}_{N,M}$) of the objective function U are available. Hence, the optimal design problem becomes a stochastic optimization problem, typically over a continuous design space \mathcal{H} . Many algorithms have been devised to solve continuous optimization problems with stochastic objectives. While some do not require the direct evaluation of gradients (e.g., Nelder-Mead [76], Kiefer-Wolfowitz [54], and simultaneous perturbation stochastic approximation [90]), other algorithms can use gradient evaluations to great advantage. Broadly, these algorithms involve either stochastic approximation (SA) [56] or sample average approximation (SAA) [87], where the latter approach must also invoke a gradient-based deterministic optimization algorithm. SA requires an unbiased estimator of the gradient of the objective, computed anew at each optimization iteration. SAA approaches, on the other hand, “freeze” the randomness in the objective and solve the resulting deterministic optimization problem, the solution of which yields an *estimate* of the solution of (2.10) [6]. Hybrids of the two approaches are possible as well. [49] presents a systematic comparison of SA and SAA approaches in the context of optimal experimental design, where SAA is coupled with a BFGS quasi-Newton method for deterministic optimization.

An alternative approach to the optimization problem (2.10) involves constructing and optimizing Gaussian process models of $U(\eta)$, again from noisy evaluations. As presented in [96], this approach generalizes the EGO (efficient global optimization) algorithm of [51] by choosing successive evaluation points η according to an expected quantile improvement criterion [80].

Surrogate models. An efficient optimization approach is only one part of the computational toolbox for optimal experimental design. Evaluating estimators such as $\hat{U}_{N,M}(\eta)$ (2.9) for even a single value of η can be computationally taxing when the likelihood $p(y|\theta, \eta)$ contains a computationally intensive model $G(\theta, \eta)$ —a situation that occurs very often in physical systems, including in materials science. As a result, considerable effort has gone into the development of reduced-order or “surrogate” models, designed to serve as computationally inexpensive replacements for G .

Useful surrogate models can take many different forms. [34] categorizes surrogates into three different classes: data-fit models, reduced-order models, and hierarchical models. Data-fit models are typically generated using interpolation or regression of the input-output relationship induced by the high-fidelity model $G(\theta, \eta)$, based on evaluations of G at selected input values $(\theta^{(i)}, \eta^{(i)})$. This class includes

polynomial chaos expansions that are constructed non-intrusively [41, 57, 100] and, more broadly, interpolation or pseudospectral approximation with standard basis functions on (adaptive) sparse grids [24, 40, 101]. Gaussian process emulators [53, 99], widely used in the statistics community, fall into this category as well. Indeed, the systematic and efficient construction of data-fit surrogates, particularly for high-dimensional input spaces, has been the focus of a vast body of work in computational mathematics and statistics over the past decade. While many of these methods are used in forward uncertainty propagation (e.g., the solution of PDEs with random input data), recent work [48] has employed sparse grid polynomial surrogates specifically for the case of optimal Bayesian experimental design.

Reduced-order models are commonly derived using a projection framework; that is, the governing equations of the forward model are projected onto a subspace of reduced dimension. This reduced subspace is defined via a set of basis vectors, which, for general nonlinear problems, can be calculated via the proper orthogonal decomposition (POD) [47, 81, 88] or with reduced basis methods [43, 77]. For both approaches, the empirical basis is pre-constructed using full forward problem simulations or “snapshots.” Systematic projection-based model reduction schemes for parameter-dependent models have also seen extensive development in recent years [17, 22]. To our knowledge, such reduction schemes have not yet been used for optimal experimental design, but in principle they are directly applicable.

Hierarchical surrogate models span a range of physics-based models of lower accuracy and reduced computational cost. Hierarchical surrogates are derived from higher-fidelity models using approaches such as simplifying physics assumptions, coarser grids, alternative basis expansions, and looser residual tolerances. These approaches may not be particularly systematic, in that their success and applicability are strongly problem-dependent, but they can be quite powerful in certain cases. One of the examples in the next section will use a reduced order model derived from a combination of simplifying physics assumptions and fits to simulation data from a higher-fidelity model.

2.3 Examples of Optimal Experimental Design

In this section, we present two examples of Bayesian experimental design in materials-related applications. The first illustrates experimental design for parameter estimation in a simple substrate-film model. This example also demonstrates the usefulness of reduced-order models in accelerating the design process. The second example is concerned with experimental design for model selection. It will illustrate this process using competing models of impurity precipitation at heterophase interfaces.

2.3.1 *Film-Substrate Systems: Design for Parameter Inference*

A classical application of Bayesian methods to physical modeling involves inferring the properties of the interior of an object from observations of its surface, e.g., of the mantle or core of the Earth from observations at the Earth’s crust [16, 44]. In the context of materials science, similar problems arise when observing the surface of a material and trying to infer the subsurface properties. One example of such a problem involves observing a thin film deposited on a heterogeneous substrate. The heterogeneity of the substrate—e.g., in temperature [58], local chemistry [3], or topography [14]—induces some corresponding heterogeneity in the film—e.g., melting [58], condensation [3], or buckling [14]. The goal is to deduce information about the substrate from the behavior of the film.

We have recently developed a convenient model for studying the inference of substrate properties from film behavior [2]. Figure 2.1 shows a film deposited on a substrate. Though the substrate is not directly observable, we would like to infer its properties from the behavior of the film deposited above. In the present example, we will use this simple model to demonstrate aspects of Bayesian experimental design. Our objective will be to choose experiments that provide maximal information about a parameter of interest for a fixed number of allowed experiments.

2.3.1.1 Physical Background

In our model problem, the substrate is described by a non-uniform scalar field $\mathcal{T}(x, y)$ on a two-dimensional spatial domain, $(x, y) \in \Omega := [0, L_D] \times [0, L_D]$. In other words, $\mathcal{T}(x, y)$ describes the variation of the substrate property \mathcal{T} over a square domain. Realizations of the substrate are random, and hence we model $\mathcal{T}(x, y)$ as a




Fig. 2.1 A film deposited on top of a substrate. The substrate is not directly observable, but some of its properties may be inferred from the behavior of the film

zero-mean Gaussian random field with a squared exponential covariance kernel [82]. One of the key parameters of this covariance kernel is the characteristic length scale ℓ_s , which describes the scale over which spatial variations in the random field occur. When ℓ_s is large, realizations of the substrate field have a relatively coarse structure, while smaller values of ℓ_s produce realizations with more fine-scale variation. The film deposited on the substrate is a two-component mixture represented by an order parameter field $c(x, y, t)$. The order parameter takes values in the range $[-1, 1]$, where $c = -1$ and $c = 1$ represent single-component phases and $c = 0$ represents a uniformly mixed phase.

The behavior of the film is modeled by the Cahn-Hilliard equation [18]:

$$\frac{\partial c}{\partial t} = \Delta \left(\frac{\partial g}{\partial c} - \epsilon^2 \Delta c \right), \quad (2.11)$$

where

$$g(c, \mathcal{T}(x, y)) = \frac{c^4}{4} + \mathcal{T}(x, y) \frac{c^2}{2} \quad (2.12)$$

is a substrate-dependent energy potential function. The two components of the film separate in regions of the substrate where $\mathcal{T}(x, y) < 0$ and mix in regions where $\mathcal{T}(x, y) > 0$. Hence, the substrate field can be thought of as a difference from some critical temperature, where temperatures above the critical value promote phase mixing while those below the critical value promote phase separation. The parameter ϵ in (2.11) governs the thickness of the interface between separated phases. Films with larger values of ϵ have thicker interfaces between their phase-separated regions than films with lower values of ϵ .

We model the time evolution of an initially uniform film $c(x, y, t = 0) = 0$ deposited on a substrate by solving the Cahn-Hilliard equation using Eyre's method for time discretization [35]. We find that the order parameter field c converges to a static configuration in the long time limit for any combination of ℓ_s and ϵ . A detailed description of the model implementation and analysis of the time-dependence of c is given in [2]. For the purpose of the example presented here, it suffices to know that the converged order parameter field has a characteristic length scale of its own, which we call Λ_∞ .

Figure 2.2 illustrates converged order parameter fields of films with two different values of ϵ ($\epsilon = 0.02$ and $\epsilon = 0.04$) deposited on substrates with two different values of ℓ_s ($\ell_s = 0.77$ and $\ell_s = 0.13$). For both substrates, we observe that increasing the value of ϵ increases the value of Λ_∞ . Yet the behavior of the film on the two substrates is qualitatively different. For the substrate with $\ell_s = 0.77$, the thickness of interfaces between phase-separated parts of the film is sufficiently small for fluctuations in c to be correlated with fluctuations in \mathcal{T} . By contrast, no direct correlation of this sort exists for the substrate with $\ell_s = 0.13$, because its characteristic length is smaller than the thickness of interfaces between phase-separated parts of the films in Fig. 2.2e, f. Instead, the fluctuations in c for these films reflects a local spatial average of T over a length scale that depends on ϵ .




Fig. 2.2 Substrate fields with (a) $l_s = 0.77$ and (d) $l_s = 0.13$. Plots b and c show converged order parameter distributions for films deposited on the substrate in a with $\epsilon = 0.02$ and $\epsilon = 0.04$ respectively. Similarly, plots e and f show converged order parameter distributions for films deposited on the substrate in d with $\epsilon = 0.02$ and $\epsilon = 0.04$ respectively. The converged length scale Λ_∞ is indicated for each film

The value of ϵ determines how Λ_∞ changes with l_s . For example, in films with $\epsilon = 0.02$, reducing the value of l_s from 0.77 to 0.13 reduces Λ_∞ from 1.05 to 0.83. However, the opposite effect is observed for $\epsilon = 0.04$, where reducing l_s from 0.77 to 0.13 increases Λ_∞ from 1.35 to 2.93. These observations show that Λ_∞ , l_s , and ϵ are related, albeit in a non-trivial way.

Our goal is to infer the substrate length scale l_s from the value of Λ_∞ of a film of known ϵ , deposited on the substrate. In this context, Λ_∞ is the data obtained from an experiment, l_s is the value to be inferred, and ϵ is a parameter of the experiment that we control (e.g., by manipulating the chemical composition of the film). In previous work, we showed how to perform this inference and how to improve it by performing multiple measurements of Λ_∞ using films with different ϵ values [2]. In the experimental design problem described here, we would like to choose optimal values of ϵ that lead to the most efficient inference of l_s .

For any given l_s and ϵ , Λ_∞ may be obtained by solving the Cahn-Hilliard equation for the time evolution of the film on the substrate. This calculation does not call for extraordinary computational resources; indeed, it can be performed in roughly 100s on a modern workstation. In Bayesian experimental design, however, this calculation would have to be carried out many millions of times. The potential computational

effort of this approach is compounded by the stochasticity of $\mathcal{T}(x, y)$; to evaluate the likelihood function for any given value of ℓ_s , we must account for many possible substrate field realizations. Therefore, to make optimal experimental design tractable, we construct a ‘‘reduced order model’’ (ROM) relating Λ_∞ , ℓ_s , and ϵ . We use a relation of the form

$$\Lambda_\infty = \overbrace{f(\epsilon, \ell_s) + \gamma(\epsilon, \ell_s)}^{\text{reduced order model}}. \quad (2.13)$$

deterministic term
random term

The deterministic term captures the average response of the film/substrate system, and the random term captures the inherent stochasticity of the film/substrate system and any systematic error in the deterministic term. The stochasticity of the film/substrate system is due to the random nature of the substrate field and the initial condition of the Cahn-Hilliard equation, among other factors [2].

The proposed ROM can be simplified using the Buckingham Pi theorem [102]. Since ϵ , ℓ_s , and Λ_∞ all have dimensions of length, we can form two Pi groups: $(\Lambda_\infty/\epsilon)$ and (ℓ_s/ϵ) . The ROM may then be simplified to

$$\frac{\Lambda_\infty}{\epsilon} = F\left(\frac{\ell_s}{\epsilon}\right) + \Gamma\left(\frac{\ell_s}{\epsilon}\right). \quad (2.14)$$

To obtain the form of $F(\ell_s/\epsilon)$ and $\Gamma(\ell_s/\epsilon)$, we carried out multiple runs of the Cahn-Hilliard model, with values of ℓ_s sampled over $[0.1, 1]$ and values of ϵ sampled over $[0.01, 0.1]$. Figure 2.3a plots Λ_∞/ϵ as a function of ℓ_s/ϵ , confirming that these quantities lie on a single curve, on average. However, there is a spread about this curve as well. This is caused by the stochastic nature of the relation between Λ_∞/ϵ and ℓ_s/ϵ , and justifies the random term in the ROM. The exact forms of $F(\ell_s/\epsilon)$ and $\Gamma(\ell_s/\epsilon)$ are then:

$$F\left(\frac{\ell_s}{\epsilon}\right) = \frac{\ell_s}{\epsilon} \left(a + \frac{b}{(\ell_s/\epsilon - 1)^c} \right) \quad (2.15)$$

$$\Gamma\left(\frac{\ell_s}{\epsilon}\right) \sim N\left(0, \sigma^2 \left(\frac{\ell_s}{\epsilon}\right)\right) \quad (2.16)$$

with parameters of the mean term F obtained by least squares fitting:

$$a = 1.05 \quad b = 79.51 \quad c = 1.54.$$

The dependence of σ^2 on (ℓ_s/ϵ) is captured nonparametrically using Gaussian process regression [82], as shown in Fig. 2.3b. Details of the derivation of the ROM can be found in [2].

To perform inference, we use the Cahn-Hilliard model as a proxy for a physical experiment. We generate multiple realizations of substrates with the same value of ℓ_s . Then, using each substrate as an input, we run the Cahn-Hilliard model, which also requires ϵ as a parameter. Given one or more choices for ϵ and the values of Λ_∞




Fig. 2.3 **a** A plot of Λ_∞/ϵ against ℓ_s/ϵ . **b** A plot of the non-stationary variance of the random term $\Gamma(\ell_s/\epsilon)$

thus obtained, we infer the value of ℓ_s using the ROM. Inference may be conducted using one or multiple $(\Lambda_\infty, \epsilon)$ pairs.

To infer ℓ_s in a Bayesian setting, we need to calculate the likelihood $p(\Lambda_\infty, \epsilon)$. This can be done using the ROM as follows:

$$p(\Lambda_\infty|\ell_s, \epsilon) = \frac{1}{\sqrt{2\pi}\sigma(\ell_s/\epsilon)} \exp\left(-\frac{(\Lambda_\infty/\ell_s - F(\ell_s/\epsilon))^2}{2\sigma^2(\ell_s/\epsilon)}\right). \quad (2.17)$$

Since runs of the Cahn-Hilliard equation are conditionally independent given ℓ_s and ϵ , the likelihood for multiple $(\Lambda_\infty, \epsilon)$ pairs can be found using the product rule

$$p(\Lambda_{\infty,1:n}|\ell_s, \epsilon_{1:n}) = \prod_i p(\Lambda_{\infty,i}|\ell_s, \epsilon_i). \quad (2.18)$$

Finally, the posterior density is calculated using Bayes' rule

$$p(\ell_s|\Lambda_{\infty,1:n}, \epsilon_{1:n}) = \frac{p(\Lambda_{\infty,1:n}|\ell_s, \epsilon_{1:n})p(\ell_s)}{\int p(\Lambda_{\infty,1:n}|\ell_s, \epsilon_{1:n})p(\ell_s)d\ell_s}. \quad (2.19)$$

We use a truncated Jeffreys prior [50] for ℓ_s

$$p(\ell_s) \propto \ln(1/\ell_s), \quad \ell_s \in [0.1, 1]. \quad (2.20)$$

The prior density is set to zero outside the range $[0.1, 1]$. This restriction is imposed for reasons of computational convenience and may easily be relaxed.

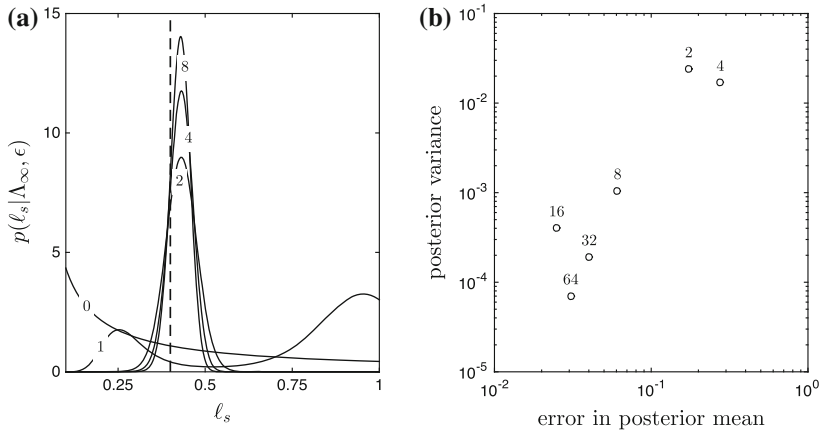


Fig. 2.4 **a** Posterior probability densities for different numbers of $(\Lambda_\infty, \epsilon)$ pairs. With the inclusion of ever more data, uncertainty in the posterior on ℓ_s decreases steadily. **b** Posterior variance and error in posterior mean for different numbers of $(\Lambda_\infty, \epsilon)$ pairs. Both error and variance decrease with increasing numbers of data points

The results of an iterative inference process that incorporates successive $(\Lambda_\infty, \epsilon)$ pairs are shown in Fig. 2.4a. Here, the true value of the substrate length scale (i.e., the value used to generate the data) is $\ell_s = 0.4$. Values of ϵ are selected by sampling uniformly in log-space over the interval $[0.01, 0.1]$. The probability density marked ‘0’ (i.e., with zero data points) is the prior. The posterior probability density with one data point (marked ‘1’) is bimodal, but the bimodality of the posterior vanishes with two or more data points. As additional $(\Lambda_\infty, \epsilon)$ pairs are introduced, the peak in the posterior moves towards the true value of $\ell_s = 0.4$. Any number of point estimates of ℓ_s may be calculated from the posterior, such as the mean, median, or mode, but the posterior probability density itself gives a full characterization of the uncertainty in ℓ_s . As an example, we have plotted in Fig. 2.4b both the posterior variance (a measure of uncertainty) and the absolute difference between the posterior mean and the true value of ℓ_s (a measure of error) for different numbers of data points. As more data are used in the inference problem, both the posterior variance and the error in the posterior mean decrease. Note that the ultimate convergence of the posterior mean towards the true value of ℓ_s , as the number of data points approaches infinity, is a more subtle issue; it is related to the frequentist properties of this Bayesian estimator, here in the presence of model error. For a fuller discussion of this topic, see [2].

2.3.1.2 Bayesian Experimental Design

Thus far, we have described a model problem wherein the characteristic length scale ℓ_s of a substrate is inferred from the behavior of films with known values of ϵ , deposited on the substrate. In the preceding calculations, we chose ϵ randomly from




Fig. 2.5 Map of expected information gain $U(\epsilon_1, \epsilon_2)$ in the substrate length scale parameter ℓ_s , as a function of experimental design parameters ϵ_1 and ϵ_2 . The three experiments discussed in the text are marked with *red squares*

a distribution. Since ϵ is in fact an experimental parameter that we can control, this choice is equivalent to performing experiments at random. Now we would like to consider a more focused experimental campaign, choosing values of ϵ to maximize the information gained with each experiment. In the language of Sect. 2.2, we will take our utility function u to be the Kullback-Leibler (KL) divergence from the posterior to the prior (2.5). The expected utility (2.4) will represent expected information gain in the parameters ℓ_s . To connect the present problem to the general formulation of Sect. 2.2, note that Λ_∞ is the experimental data y , ℓ_s is the parameter θ to be inferred, and ϵ is the experimental parameter η over which we will optimize the expected utility.

The expected KL divergence from posterior to prior is estimated via the Monte Carlo estimator in (2.9). To perform the calculation, we need to be able to sample $\ell_s^{(i)}$ and $\tilde{\ell}_s^{(i,j)}$ from the prior $p(\ell_s)$, and $\Lambda_\infty^{(i)}$ from the likelihood $p(\Lambda_\infty | \ell_s^{(i)}, \epsilon)$. The length scales ℓ_s can be sampled from the truncated Jeffreys prior using a standard inverse CDF transformation [83]. The observation Λ_∞ is Gaussian given ϵ and ℓ_s , and can be sampled by evaluating (2.14) with distributional information given in (2.15)–(2.16).

We will use this formulation to design an optimal experiment consisting of two measurements. In other words, two films with independently controlled values of ϵ will be deposited on substrates with the same value of ℓ_s , and the two values of Λ_∞ generated will be used for inference. The values of ϵ will be restricted to the design range $[0.01, 0.095]$. As before, this restriction is not essential and is easily relaxed. Figure 2.10 shows the resulting Monte Carlo estimates of expected information gain $U(\epsilon_1, \epsilon_2)$.

Because the ordering of the experiments is immaterial, the map of the expected information gain is symmetric about the $\epsilon_1 = \epsilon_2$ line, aside from Monte Carlo estimation error. We draw attention to three points marked by squares in Fig. 2.10. The first is at $(\epsilon_1, \epsilon_2) = (0.025, 0.025)$, where $U(\epsilon_1, \epsilon_2) = 0.49$; it is near the minimum of the expected utility function. This point corresponds to the least useful pair of experiments. The second is at $(\epsilon_1, \epsilon_2) = (0.01, 0.095)$, with $U(\epsilon_1, \epsilon_2) = 2.9$; it is the maximum of the expected utility map and is expected to yield the most informative experiments. The point $(\epsilon_1, \epsilon_2) = (0.08, 0.08)$, where $U(\epsilon_1, \epsilon_2) = 2.0$, lies midway between these extremes: it is expected to be more informative than the first design but less informative than the second.

To illustrate how the three (ϵ_1, ϵ_2) pairs highlighted above yield different expected utilities, we carry out the corresponding inferences of ℓ_s following the procedure described in Sect. 2.3.1.1. To simulate each experiment, we fix ℓ_s and the desired value of ϵ , then generate a converged order parameter length scale Λ_∞ by generating a realization of the substrate and simulating the Cahn-Hilliard equation. Given the data $\Lambda_{\infty,1}$ and $\Lambda_{\infty,2}$ corresponding to (ϵ_1, ϵ_2) , we evaluate the corresponding posterior density and calculate the actual KL divergence from posterior to prior, $D_{KL}(p(\ell_s | \Lambda_{\infty,1:2}, \epsilon_{1:2}) \| p(\ell_s))$. The results of these three experiments are summarized in Fig. 2.6a. As expected, the second experiment, performed at




Fig. 2.6 **a** Experiments corresponding to the three (ϵ_1, ϵ_2) pairs indicated in Fig. 2.5. The posterior densities from the three experiments are marked #1, #2, and #3. **b** $\Lambda_\infty = \epsilon F(\ell_s/\epsilon)$ versus ℓ_s for $\epsilon = 0.01, 0.025, 0.095$

$(\epsilon_1, \epsilon_2) = (0.01, 0.095)$ is the most informative, and has a large information gain of $D_{KL} = 2.10$ nats.¹ The first experiment is the least informative, with a small information gain of $D_{KL} = 0.99$ nats. The third experiment, with $D_{KL} = 1.44$ nats, lies in between. The actual values of D_{KL} are different from their expected values because the expected information gains are calculated by averaging over all possible prior values of ℓ_s and all possible experimental outcomes, whereas the actual values are calculated only for particular ℓ_s and Λ_∞ values, given ϵ . However, the values of D_{KL} follow the same trend as their expectations.

To better understand why these experiments produce different values of the information gain, Fig. 2.6b plots $\Lambda_\infty = \epsilon F(\ell_s/\epsilon)$ as a function of ℓ_s for $\epsilon = 0.01, 0.025,$ and 0.095 . We observe that Λ_∞ is not very sensitive to variations in ℓ_s for $\epsilon = 0.025$. This explains why an experiment with $(\epsilon_1, \epsilon_2) = (0.025, 0.025)$ is not particularly informative. On the other hand, Λ_∞ is sensitive to variations in ℓ_s for $\epsilon = 0.095$ and $\epsilon = 0.01$. Additionally, Λ_∞ is a decreasing function of ℓ_s for $\epsilon = 0.095$, and an increasing function for $\epsilon = 0.01$. The complementarity of these trends makes the experiment $(\epsilon_1, \epsilon_2) = (0.01, 0.095)$ especially useful.

We can also compare the optimal experiment to the random experiments shown in Fig. 2.4. The information gained in the optimal experiment ($D_{KL} = 2.10$ nats), with two values of ϵ , is comparable to the information gained from the experiment with *eight* randomly selected values of ϵ ($D_{KL} = 2.29$ nats). Hence by using optimal Bayesian experimental design in this example, we are able to reduce the experimental effort over a random strategy by roughly a factor of four! This reduction is especially valuable when experiments are difficult or expensive to conduct.

2.3.2 *Heterophase Interfaces: design for model discrimination*

As noted in Sect. 2.2.1, experiments often yield data that may be explained by multiple models. Additional measurements may then be required to determine which of many possible models is best supported by the data. In such situations, it is desirable to determine which further experiments are likely to distinguish between alternative models most efficiently. Naturally, this guidance is needed before the additional work is actually carried out. Determining which experiments are most informative for distinguishing between alternative models is the goal of Bayesian experimental design for model selection [63, 75]. This capability is especially useful when the experiments are very resource-intensive and brute force data acquisition over a wide parameter range is not feasible. This section will illustrate Bayesian experimental design for model selection on an example taken from investigations of heterophase interfaces in layered metal composites.

¹A nat is a unit of information, analogous to a bit, but with a natural logarithm rather than a base two logarithm in (2.5).




Fig. 2.7 A Cu-Nb multilayer composite synthesized by PVD [31]

2.3.2.1 Physical Background

Figure 2.7 shows a multilayer composite of two transition metals—copper (Cu) and niobium (Nb)—created by physical vapor deposition (PVD) [74]. In this synthesis technique, atoms impinge upon a flat substrate, adhere to it, and aggregate into crystalline layers. By alternating the elements being deposited—e.g., first Cu, then Nb, then Cu again, and so on—layered composites such as the one in Fig. 2.7 may be made. The thickness of each layer may be controlled by changing the total deposition time for each element. Many multilayer compositions besides Cu/Nb have been synthesized this way, including Cu/V (V: vanadium) [37, 105], Cu/Mo (Mo: molybdenum) [62], Ag/V (Ag: Silver) [97, 98], Al/Nb (Al: Aluminum) [37, 61, 62], Fe/W (Fe: Iron, W: Tungsten) [60], and others [11, 12, 86].

Layered composites are ideal for studying the properties of heterophase interfaces. In Fig. 2.7, each pair of adjacent Cu and Nb layers forms one Cu-Nb interface. The total amount of interface area per unit volume of the material may be changed by adjusting the thickness of the layers. For composites where all the individual layers have identical thickness l , the volume of material corresponding to interface area A is $V = A \times l$. Thus, the interface area per unit volume is $A/V = 1/l$: as the layers are made thinner, A/V rises and the influence of interfaces on the physical properties of the composite as a whole increases. For l in the nanometer range, i.e., $l \lesssim 10$ nm, interfaces dominate the behavior of the multilayer composite, leading to enhanced strength [73], resistance to radiation [72], and increased fatigue life [95]. In multilayer composites, all of these desirable properties are due to the influence

of interfaces. Thus, considerable effort continues to be invested into elucidating the structure and properties of individual interfaces [10, 19, 71].

The present example will consider the relationship between the structure of metal-metal heterophase interfaces and trapping of helium (He) impurities. Implanted He is a major concern for the performance of materials in nuclear energy applications [92, 106]. Trapping and stable storage of these impurities at interfaces is one way of mitigating the deleterious effects of implanted He [32, 78]. The influence of interfaces on He may be clearly seen in multilayer composites. Experiments carried out on Cu-Nb [29], Cu-Mo [62], and Cu-V [38] multilayers synthesized by PVD show that implanted He is preferentially trapped at the interfaces. Moreover, not all interfaces are equally effective at trapping He: the maximum concentration c of interfacial He—expressed as the number of He atoms per unit interface area—that may be trapped at an interface before detectable He precipitates form differs from interface to interface [32].

Figure 2.8 plots c for Cu-Nb, Cu-Mo, and Cu-V interfaces as a function of one parameter: the interface “misfit” m . Cu, Nb, Mo, and V all have cubic crystal structures: Cu is face-centered cubic (fcc) while Nb, Mo, and V are body-centered cubic (bcc). Thus, all three composites used in Fig. 2.8 are made up of alternating fcc (Cu) and bcc (Nb, Mo, or V) layers. The edge length of a single cubic unit cell in fcc or bcc crystals is the lattice parameter, a_{fcc} or a_{bcc} . The misfit m is defined as $m = a_{\text{bcc}}/a_{\text{fcc}}$. Intuitively, m measures the mismatch in inter-atomic spacing in the adjacent crystals that make up an interface.

According to Fig. 2.8, the ability of interfaces to trap He, as measured by c , depends on the misfit: $c = c(m)$. A simple model that may be proposed based on this data is that the relationship between c and m is linear: $c = \alpha_0 + \alpha_1 m$. Indeed, a linear fit represents the available data reasonably well. Its most apparent drawback




Fig. 2.8 Maximum interfacial He impurity concentration, c , plotted as a function of misfit, m

is that it predicts negative c values for $m \lesssim 0.83$. Thus, a better model might be $c = \alpha|m - m_0|$. This model predicts that c drops to zero as m decreases to m_0 and begins to rise again as m is further reduced below m_0 . Physically, this model may be rationalized by stating that at some special value of misfit, m_0 , the atomic matching between adjacent layers is especially good, leading to few sites at the interface where He impurities may be trapped. As m departs from m_0 in either direction, the atomic matching becomes worse, giving rise to more He trapping sites and therefore higher c .

The structure of fcc/bcc interfaces—including the degree of atomic matching—may be investigated in detail by constructing atomic-level models [30, 33]. Figure 2.9 shows such a model of the terminal atomic planes of Cu and Nb found at PVD Cu-Nb interfaces. The pattern of overlapping atoms from the two planes contains sites where a Cu atom is nearly on top of a Nb atom. Such sites are thought to be preferential locations for trapping of He impurities [52]. They arise from the geometrical interference of the overlapping atomic arrangements in the adjacent crystalline layers and, in that sense, may be viewed as analogous to features in a Moiré pattern.

The density and distribution of He trapping sites of the type shown in Fig. 2.9 may be computed directly for any given fcc/bcc interface as a function of the geometry of the interface using the well-known O-lattice theory [13]. In PVD Cu/Nb, Cu/Mo, and Cu/V composites, the relative orientation of the adjacent crystals is identical. Thus, differences in the geometry Cu-Nb, Cu-Mo, and Cu-V interfaces arise solely from differences in the lattice parameters of the adjacent crystals, as described by the misfit parameter, m . The areal density of He trapping sites for these interfaces is therefore only a function of m and may be written as

$$f(m) = \frac{|(4m - 3)(\sqrt{3}m - \sqrt{2})|}{3a_{\text{Cu}}^2 m^2}. \quad (2.21)$$




Fig. 2.9 *Left* A Cu-Nb bilayer. *Right* The terminal Cu and Nb planes that meet at the Cu-Nb interface. He trapping occurs at sites where a Cu atom is nearly on *top* of a Nb atom. One such site is shown with the *dashed circle*

Using this expression, we propose a second model for the dependence of c on m , namely: $c = \beta f(m)$. Here, the proportionality constant β determines the number of He atoms that may be trapped at a single site of the type illustrated in Fig. 2.9. The best fit for this second model is plotted in Fig. 2.8. Both this model and the previously described linear model fit the available experimental data reasonably well. Moreover, both predict c values of zero for $m \approx 0.82$ – 0.83 . However, unlike the linear model, $c = \beta f(m)$ predicts that c is also zero at $m \approx 0.75$.

We wish to determine what additional experimental data will help distinguish between the two models described above. However, since measuring even a single value of c requires considerable resources, our goal is to limit the additional data to just one (c, m) pair. In an experiment, we may select m by choosing to synthesize a fcc/bcc multilayer composite of specific composition. In other words, we control m . However, we do not know c in advance. In this context, our goal is to determine what one value of m is most likely to distinguish between the two models, regardless of the c value actually found in the subsequent experiment. In the following section, we will apply Bayesian experimental design to address this challenge. In addition to the two models described above, we will also consider a third model encapsulating the hypothesis that c does not depend on m at all: $c = \gamma = \text{constant}$.

2.3.2.2 Bayesian Experimental Design for Model Selection

As described in Sect. 2.2, the goal of experimental design is to maximize the expectation of some measure of information. In the present example, we will maximize the expected KL divergence, as applied to model discrimination, described in (2.6) and (2.7). In this context, m is the experimental parameter η that we control; c is the observed data y ; M_1 , M_2 , and M_3 are the competing models; (α, m_0) are the parameters θ_1 of model M_1 ; β is the parameter θ_2 of model M_2 ; and γ is the parameter θ_3 of model M_3 .

The expected KL divergence can be computed by combining (2.6) and (2.7); this requires knowing both the prior $p(M_i)$ and the posterior $p(M_i|c, m)$ for each model. We use a flat or “indifference” prior over models, $p(M_i) = 1/3$. The posterior model probabilities are calculated from Bayes rule as given in (2.2). Evaluating Bayes’ rule in this setting requires that we calculate the marginal likelihood for each model and proposed experiment, $p(c|M_i, m)$, as shown in (2.3). We now detail this procedure.

The previous section identified three models connecting c and m . They are:

$$M_1 : c = \alpha|m - m_0| + \epsilon_1 \quad (2.22)$$

$$M_2 : c = \beta f(m) + \epsilon_2 \quad (2.23)$$

$$M_3 : c = \gamma + \epsilon_3 \quad (2.24)$$

where $\epsilon_i \sim N(0, \sigma_\epsilon^2)$. In addition to specifying the functional form of each model, each expression above also contains an additive noise term ϵ_i . This term is a random variable that describes uncertainty in the measured c , i.e., due to the observational

process itself. For simplicity, we assume the observational error variance σ_ϵ^2 to be known. The model parameters α , m_0 , β , and γ are endowed with priors that reflect our state of knowledge after performing the three experiments shown in Fig. 2.8, before beginning the current experimental design problem. These priors are taken to be Gaussian. In other words, we suppose that they are the result of Bayesian linear regression with Gaussian priors or improper uniform priors; the posterior following the three previous experiments becomes the prior for the current experimental design problem. We denote the current prior means by $\bar{\alpha}$, \bar{m}_0 , $\bar{\beta}$, and $\bar{\gamma}$, and the current prior standard deviations as σ_α , σ_{m_0} , σ_β , and σ_γ .

Given these assumptions, we can express the probability density of the observation c for each parameterized model as:

$$p(c|m, \alpha, m_0, M_1) = \frac{1}{\sqrt{2\pi}\sigma_\epsilon} \exp\left(-\frac{(c - \alpha|m - m_0|)^2}{2\sigma_\epsilon^2}\right) \quad (2.25)$$

$$p(c|m, \beta, M_2) = \frac{1}{\sqrt{2\pi}\sigma_\epsilon} \exp\left(-\frac{(c - \beta f(m))^2}{2\sigma_\epsilon^2}\right) \quad (2.26)$$

$$p(c|m, \gamma, M_3) = \frac{1}{\sqrt{2\pi}\sigma_\epsilon} \exp\left(-\frac{(c - \gamma)^2}{2\sigma_\epsilon^2}\right). \quad (2.27)$$

Each of these densities is normal with mean given by the model and variance σ_ϵ^2 . For fixed m and c , these densities can be viewed as the *likelihood functions* for the corresponding model parameters, i.e., α and m_0 for model 1, β for model 2, and γ for model 3. To obtain the *marginal likelihoods* $p(c|m, M_i)$, we marginalize out these parametric dependencies as follows:

$$p(c|m, M_1) = \int_{-\infty}^{\infty} p(c|m, \alpha, m_0, M_1) p(\alpha) p(m_0) d\alpha dm_0 \quad (2.28)$$

$$p(c|m, M_2) = \int_{-\infty}^{\infty} p(c|m, \beta, M_2) p(\beta) d\beta \quad (2.29)$$

$$p(c|m, M_3) = \int_{-\infty}^{\infty} p(c|m, \gamma, M_1) p(\gamma) d\gamma \quad (2.30)$$

Here, $p(\alpha)$, $p(m_0)$, $p(\beta)$, and $p(\gamma)$ denote the Gaussian prior probability densities described above, e.g.,

$$p(\alpha) = \frac{1}{\sqrt{2\pi}\sigma_\alpha} \exp\left(-\frac{(\alpha - \bar{\alpha})^2}{2\sigma_\alpha^2}\right), \text{ etc.} \quad (2.31)$$

In the expressions for $p(c|m, M_i)$, integration over α , β , and γ can be performed analytically, e.g.,

$$p(c|m, M_2) = \frac{1}{\sigma_\beta\sigma_\epsilon\sqrt{\frac{2\pi}{\sigma_\beta^2} + \frac{2\pi f(m)^2}{\sigma_\epsilon^2}}} \exp\left(-\frac{(c - \bar{\beta}f(m))^2}{2(f(m)^2\sigma_\beta^2 + \sigma_\epsilon^2)}\right). \quad (2.32)$$

Table 2.1 Prior model parameters

Model	Parameter	Standard deviation
M_1	$\bar{\alpha} \approx 94/\text{nm}^2$	$\sigma_\alpha \approx 0.49/\text{nm}^2$
	$\bar{m}_0 \approx 0.83$	$\sigma_{m_0} \approx 0.62$
M_2	$\bar{\beta} \approx 26/\text{nm}^2$	$\sigma_\beta \approx 4.2/\text{nm}^2$
M_3	$\bar{\gamma} \approx 4.5/\text{nm}^2$	$\sigma_\gamma = 2.0/\text{nm}^2$

The integral over m_0 in the expression for $p(c|m, M_1)$ must be found numerically, however. In the present example, this integral is easily computed using standard numerical quadrature. If the integral had been too high dimensional, however, then a Monte Carlo scheme might be used instead [83]. We carry out these calculations using prior parameters listed in Table 2.1. The experimental uncertainty was set to $\sigma_\epsilon = 2.5/\text{nm}^2$, following [29].

To calculate the expected information gain U in the model indicator, as a function of the m value for a single additional experiment, we first substitute the prior and posterior model probabilities calculated above into (2.6). Then we take the expectation of this utility over the prior predictive distribution, as in (2.7), by integrating over the data c . More explicitly, we calculate:

$$U(m) = \int u(m, c) p(c|m) dc, \quad (2.33)$$

where the utility is

$$u(m, c) = \sum_{i=1}^3 P(M_i|c, m) \log \frac{P(M_i|c, m)}{P(M_i)},$$

and the design-dependent prior predictive probability density is

$$p(c|m) = \sum_{i=1}^3 P(M_i) p(c|m, M_i).$$

The integral in (2.33) formally is taken over $(-\infty, \infty)$, since this is the range of the prior predictive. Negative values of c are not physical, of course, but they are exceedingly rare: the mean predictions of models 1 and 2 are necessarily positive, and the Gaussian prior on γ in model 3 is almost entirely supported above zero. The Gaussian measurement noise ϵ can also lead to negative c values, but it too has a relatively small variance.

Figure 2.10 plots $U(m)$ computed using all three models. For comparison, the figure also shows $U(m)$ found using only models 1 and 2, i.e., excluding the constant model $c = \gamma$. Values of m that maximize $U(m)$ are the best choices for an experiment to distinguish between models. When all three models are considered,




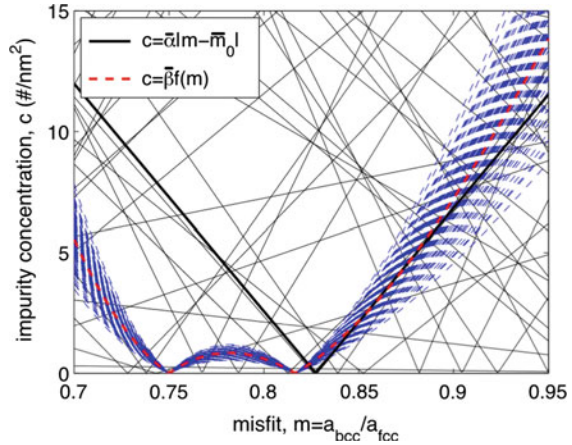
Fig. 2.10 Expected information gain for model discrimination $U(m)$

$U(m)$ is greatest for high misfit, i.e., $m \approx 0.95$. By contrast, when only models 1 and 2 are considered, $U(m)$ is least in the high m limit. The reason for this difference is clear from comparing Fig. 2.10 with Fig. 2.8: models 1 and 2 predict comparable c at high m while model 3 predicts a markedly lower c . Thus, when all three models are considered, the value of $U(m)$ is high for $m \approx 0.95$ because a measurement at that m value makes it possible to distinguish models 1 and 2 from model 3. By contrast, when only models 1 and 2 are considered, $U(m)$ is least at high m because a measurement in that m range has limited value for distinguishing between models 1 and 2.

Putting aside model 3, $U(m)$ predicts greatest utility for an experiment carried out in the range $0.74 < m < 0.84$, i.e., in the vicinity of the minima of function $f(m)$. To understand the reason for the significance of this m range, it is important to realize that the plots in Fig. 2.8 only show a single realization of models 1 and 2, namely those corresponding to $\alpha = \bar{\alpha}$, $m_0 = \bar{m}_0$, and $\beta = \bar{\beta}$ (the prior means on the parameters). Since we assume that α , m_0 , and β are normally distributed, many other realizations of these models are possible. Figure 2.11 shows 100 different realizations of models 1 and 2 obtained by drawing α , m_0 , and β at random from their prior distributions.

Figure 2.11 makes clear an important distinction between models 1 and 2 that is not apparent in Fig. 2.8: model 1 exhibits an extreme sensitivity to its fitting parameters within the range of uncertainty of those parameters. In particular, the minimum in c predicted by model 1 may occur at many different m values. By contrast, model 2 is relatively less sensitive to its fitting parameters, especially for $0.74 < m \lesssim 0.84$. Unlike model 1, the locations of its minima are fixed. Thus, measuring a low c value for $0.74 < m \lesssim 0.84$ has the potential to exclude a large number of realizations of

Fig. 2.11 100 different realizations of models 1 ($c = \alpha|m - m_0|$) and 2 ($c = \beta f(m)$) obtained by drawing parameters α , m_0 , and β at random from their prior distributions. The thick lines plot the realizations of the models at the prior mean values of these parameters



model 1, while measuring a high c value in that range essentially excludes model 2. Bayesian methods naturally capture this subtle aspect of experimental design without any special prior analysis of the competing models.

2.4 Outlook

The examples presented here demonstrate how the formalism of optimal Bayesian experimental design, coupled with information theoretic objectives, can be adapted to different questions of optimal data collection in materials science. In one example, we seek the best pair of experiments for inferring the parameters of a given model. In another example, we seek the single experiment that can best distinguish between competing models, where each model has a distinct form and a distinct set of uncertain parameters. Though simple, these examples also demonstrate the use of key computational and analytical tools, including the Monte Carlo estimation of expected information gain (in the nonlinear and non-Gaussian setting of our first example) and the use of reduced-order models (also in the first example).

Reduced order models (ROMs) are increasingly being recognized as crucial to materials science, especially computational materials design [39, 70, 94, 104]. The reason for their utility is clear: strictly speaking, the complete set of degrees of freedom describing a material is the complete set of positions and types of all its constituent atoms. This set defines a design space far too vast to explore. Even if mesoscale entities such as crystal defects (e.g., dislocations [45] or interfaces [91]) or microstructure [1] are used to define the degrees of freedom for design, the resulting design space may nevertheless remain too vast to examine comprehensively.

Therefore, it is crucial to identify only those degrees of freedom that significantly affect properties of interest (e.g., those affecting performance metrics in a design) and create a ROM to connect the two. Yet while formal methods of model order

reduction are well established in many fields of science and engineering (as reviewed in Sect. 2.2.3), the automated and systematic construction of ROMs in materials contexts is in an early stage of development. In practice, most ROMs in materials science are constructed “manually.” The inherently collective and multiscale character of many materials-related phenomena calls for the development and validation of new methods of automatic model order reduction to address materials-specific challenges.

Surrogate or reduced-order models are also essential to making Bayesian inference computationally tractable, particularly inference with computationally intensive physics-based models. Indeed, the past several years have seen a steady stream of developments in model reduction for Bayesian inference, mostly in the applied mathematics, computational science, and statistics communities. These include many types of prior-based ROMs [67–69], posterior-focused approximations [59] and projection-based reduced-order models [27], hierarchical surrogates [23], and numerous other approaches [25, 85]. The utility of ROMs is increasingly being recognized in materials-related inference problems as well. For instance, the model film-substrate problem described in Sect. 2.3.1 relies on a ROM to circumvent computationally expensive forward problem evaluations, thereby making rapid inference of substrate properties tractable [2]. Earlier examples of this approach in materials science problems include [28], which, starting with what is effectively a ROM for the energy-temperature relation, inferred the melting point of Ti_2GaN . Data-fit surrogates constructed based on existing literature may also serve an analogous purpose to a ROM. Using such a surrogate, [103] modeled the creep rupture life of Ni-base superalloys.

Similar to reduced-order modeling, the usefulness of Bayesian approaches is now becoming better recognized within the materials community. They can be applied to parameter inference and model inference, as demonstrated here, but also to problems involving prediction under uncertainty. For example, [55] used Bayesian inference to assess the uncertainty of cluster expansion methods for computing the internal energies of alloys. These authors point out that cluster expansions are also a kind of surrogate model—i.e., a ROM—and that uncertainty quantification should, among other goals, assess how well the surrogate reproduces the output of a more computationally expensive reference model.

Despite growing interest in Bayesian methods within the materials community, there are fewer examples of their application to experimental design. An early (yet very recent) effort is [4], which applies information-theoretic criteria and Bayesian methods to stress-strain response and texture evolution in polycrystalline solids. Nevertheless, opportunities for expanded application of optimal Bayesian experimental design abound in materials-related work. In particular, detailed and resource-intensive experiments such as those described in Sect. 2.3.2 are poised to benefit from it immensely. One potential hurdle to widespread adoption is the up-front investment of effort currently needed to understand and implement the associated mathematical formalism. Thus, expanded availability of user-friendly, well-documented, and multi-functional software [79] is likely to accelerate the adoption and integration of Bayesian experimental design into mainstream materials research.

Finally, we emphasize that optimal experimental design itself—not limited to the materials science context—is the topic of much current research. This research

focuses on both formulational issues and on computational methodology. Examples of the latter include developing reduced-order or multi-fidelity models tailored to the needs of stochastic optimization, or devising more efficient estimators of expected information gain, using importance sampling, high-dimensional kernel density estimators, and other approaches. An interesting foundational challenge, on the other hand, involves understanding and accounting for model error or misspecification in optimal design. If the model relating parameters of interest to experimental observables is incomplete or under-resolved, how useful—or close to optimal—are the experiments designed according to this relationship? When a convergent sequence of models of differing fidelity is available (as in the ROM setting), then this question is more tractable. But if all available models are inadequate, many questions remain open. One promising approach to this challenge uses nonparametric statistical models, perhaps formulated in a hierarchical Bayesian manner, to account for interactions and inputs missing from the current model of the experiment. Sequential experimental design is also useful in this context, as successive batches of experiments can help uncover the unmodeled mismatch between a model and physical reality.

Sequential experimental design is useful much more broadly as well. Recall that in all the examples of this chapter, we designed a single batch of experiments all-at-once: even if the batch contained multiple experiments, we chose the design parameters before performing any of the experiments. Sequential design, in contrast, allows information from each experiment to influence the design of the next. The most widely used sequential approaches are greedy, where one designs the next batch of experiments as if it were the final batch—using the current state of knowledge as the prior distribution, with design criteria similar to those used here. But greedy approaches are sub-optimal in general, as they do not account for the information to be gained from *future* experiments. An optimal approach can instead be obtained by formulating sequential experimental design as a problem of dynamic programming [7, 9, 20]. Making this dynamic programming approach computationally tractable, outside of specialized settings, remains a significant challenge.

References

1. B.L. Adams, S.R. Kalidindi, D.T. Fullwood, *Microstructure Sensitive Design for Performance Optimization* (Butterworth-Heinemann, Newton, 2012)
2. R. Aggarwal, M. Demkowicz, Y. Marzouk, Bayesian inference of substrate properties from film behavior. *Model. Simul. Mater. Sci. Eng.* **23**, 015009 (2015)
3. J. Aizenberg, A. Black, G. Whitesides, Controlling local disorder in self-assembled monolayers by patterning the topography of their metallic supports. *Nature* **394**, 868–871 (1998)
4. S. Atamturktur, J. Hegenderfer, B. Williams, C. Unal, Selection criterion based on an exploration-exploitation approach for optimal design of experiments. *J. Eng. Mech.* **141** (2014)
5. A.C. Atkinson, A.N. Donev, *Optimum Experimental Designs*, Oxford Statistical Science Series (Oxford University Press, Oxford, 1992)
6. G. Bayraksan, D.P. Morton, Assessing solution quality in stochastic programs via sampling. *INFORMS Tutor. Oper. Res.* **5**, 102–122 (2009)

7. I. Ben-Gal, M. Caramanis, Sequential doe via dynamic programming. *IIE Trans.* **34**, 1087–1100 (2002)
8. J. Berger, L. Pericchi, Objective Bayesian methods for model selection: introduction and comparison, in *Model Selection*, IMS Lecture Notes—Monograph Series, ed. by P. Lahiri (2001), pp. 135–207
9. D.P. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd edn. (Athena Scientific, Belmont, 2007)
10. I. Beyerlein, M. Demkowicz, A. Misra, B. Uberuaga, Defect-interface interactions. *Progr. Mater. Sci.* (2015)
11. D. Bhattacharyya, N. Mara, P. Dickerson, R. Hoagland, A. Misra, Transmission electron microscopy study of the deformation behavior of Cu/Nb and Cu/Ni nanoscale multilayers during nanoindentation. *J. Mater. Res.* **24**, 1291–1302 (2009)
12. D. Bhattacharyya, N. Mara, P. Dickerson, R. Hoagland, A. Misra, Compressive flow behavior of Al-TiN multilayers at nanometer scale layer thickness. *Acta Mater.* **59**, 3804–3816 (2011)
13. W. Bollmann, *Crystal Defects and Crystalline Interfaces* (Springer, Berlin, 1970)
14. N. Bowden, S. Brittain, A. Evans, J. Hutchinson, G. Whitesides, Spontaneous formation of ordered structures in thin films of metals supported on an elastomeric polymer. *Nature* **393**, 146–149 (1998)
15. G.E.P. Box, H.L. Lucas, Design of experiments in non-linear situations. *Biometrika* **46**, 77–90 (1959)
16. T. Bui-Thanh, O. Ghattas, J. Martin, G. Stadler, A computational framework for infinite-dimensional Bayesian inverse problems part I: the linearized case, with application to global seismic inversion. *SIAM J. Sci. Comput.* **35**, A2494–A2523 (2013)
17. T. Bui-Thanh, K. Willcox, O. Ghattas, Model reduction for large-scale systems with high-dimensional parametric input space. *SIAM J. Sci. Comput.* **30**, 3270–3288 (2008)
18. J. Cahn, J. Hilliard, Free energy of a nonuniform system. I. Interfacial free energy. *J. Chem. Phys.* **28**, 258–267 (1958)
19. P.R. Cantwell, M. Tang, S.J. Dillon, J. Luo, G.S. Rohrer, M.P. Harmer, Grain boundary complexions. *Acta Mater.* **62**, 1–48 (2014)
20. B.P. Carlin, J.B. Kadane, A.E. Gelfand, Approaches for optimal sequential decision analysis in clinical trials. *Biometrics*, pp. 964–975 (1998)
21. K. Chaloner, I. Verdinielli, Bayesian experimental design: a review. *Stat. Sci.* **10**, 273–304 (1995)
22. S. Chaturantabut, D.C. Sorensen, Nonlinear model reduction via discrete empirical interpolation. *SIAM J. Sci. Comput.* **32**, 2737–2764 (2010)
23. J.A. Christen, C. Fox, MCMC using an approximation. *J. Comput. Graph. Stat.* **14**, 795–810 (2005)
24. P. Conrad, Y.M. Marzouk, Adaptive Smolyak pseudospectral approximations. *SIAM J. Sci. Comput.* **35**, A2643–A2670 (2013)
25. P. Conrad, Y.M. Marzouk, N. Pillai, A. Smith, Accelerating asymptotically exact MCMC for computationally intensive models via local approximations. *J. Am. Stat. Assoc.* submitted (2014). [arXiv:1402.1694](https://arxiv.org/abs/1402.1694)
26. T.M. Cover, J.A. Thomas, *Elements of Information Theory*, 2nd edn. (Wiley, Hoboken, 2006)
27. T. Cui, Y.M. Marzouk, K. Willcox, Data-driven model reduction for the Bayesian solution of inverse problems. *Int. J. Numer. Methods Eng.* **102**, 966–990 (2015)
28. S. Davis et al., Bayesian inference as a tool for analysis of first-principles calculations of complex materials: an application to the melting point of ti2gan. *Model. Simul. Mater. Sci. Eng.* **21**, 075001 (2013)
29. M. Demkowicz, D. Bhattacharyya, I. Usov, Y. Wang, M. Nastasi, A. Misra, The effect of excess atomic volume on he bubble formation at fcc-bcc interfaces. *Appl. Phys. Lett.* **97**, 161903–161903 (2010)
30. M. Demkowicz, R. Hoagland, Structure of kurdjumov-sachs interfaces in simulations of a copper-niobium bilayer. *J. Nucl. Mater.* **372**, 45–52 (2008)

31. M. Demkowicz, R. Hoagland, B. Uberuaga, A. Misra, Influence of interface sink strength on the reduction of radiation-induced defect concentrations and fluxes in materials with large interface area per unit volume. *Phys. Rev. B* **84**, 104102 (2011)
32. M. Demkowicz, A. Misra, A. Caro, The role of interface structure in controlling high Helium concentrations. *Current Opin. Solid State Mater. Sci.* **16**, 101–108 (2012)
33. M.J. Demkowicz, J. Wang, R.G. Hoagland, Interfaces between dissimilar crystalline solids. *Dislocat. Solids* **14**, 141–205 (2008)
34. M. Eldred, S. Giunta, S. Collis, Second-order corrections for surrogate-based optimization with model hierarchies, in *AIAA Paper 2004-4457, Proceedings of the 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference* (2004)
35. D. Eyre, *An unconditionally stable one-step scheme for gradient systems*. Unpublished manuscript, University of Utah, Salk Lake City, June (1998)
36. I. Ford, D.M. Titterton, K. Christos, Recent advances in nonlinear experimental design. *Technometrics* **31**, 49–60 (1989)
37. E. Fu, N. Li, A. Misra, R. Hoagland, H. Wang, X. Zhang, Mechanical properties of sputtered cu/v and al/nb multilayer films. *Mater. Sci. Eng.: A*, **493**, 283–287 (2008). Mechanical Behavior of Nanostructured Materials, a Symposium Held in Honor of Carl Koch at the TMS Annual Meeting 2007, Orlando, Florida
38. E. Fu, A. Misra, H. Wang, L. Shao, X. Zhang, Interface enabled defects reduction in helium ion irradiated Cu/V nanolayers. *J. Nucl. Mater.* **407**, 178–188 (2010)
39. L.D. Gabbay, S. Senturia, Computer-aided generation of nonlinear reduced-order dynamic macromodels. I. Non-stress-stiffened case, *Microelectromech. Syst. J.* **9**, 262–269 (2000)
40. T. Gerstner, M. Griebel, Dimension-adaptive tensor-product quadrature. *Computing* **71**, 65–87 (2003)
41. R. Ghanem, P. Spanos, *Stochastic Finite Elements: A Spectral Approach* (Springer, Berlin, 1991)
42. J. Ginebra, On the measure of the information in a statistical experiment. *Bayesian Anal.* **2**, 167–212 (2007)
43. M. Grepl, Y. Maday, N. Nguyen, A. Patera, *Efficient reduced-basis treatment of nonaffine and nonlinear partial differential equations*. *Math. Model. Numer. Anal. (M2AN)* **41**, 575–605 (2007)
44. G.E. Hilley, R. Bürgmann, P.-Z. Zhang, P. Molnar, Bayesian inference of plastosphere viscosities near the kunlun fault, northern tibet, *Geophys. Res. Lett.* **32**, n/a–n/a (2005)
45. J. Hirth, J. Lothe, *Theory of Dislocations* (Wiley, New York, 1992)
46. J.A. Hoeting, D. Madigan, A.E. Raftery, C.T. Volinsky, Bayesian model averaging: a tutorial. *Stat. Sci.* **14**, 382–417 (1999)
47. P. Holmes, J. Lumley, G. Berkooz, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry* (Cambridge University Press, Cambridge, 1996)
48. X. Huan, Y.M. Marzouk, Simulation-based optimal Bayesian experimental design for nonlinear systems. *J. Comput. Phys.* **232**, 288–317 (2013)
49. X. Huan, Y.M. Marzouk, Gradient-based stochastic optimization methods in Bayesian experimental design. *Int. J. Uncertain. Quantif.* **4**, 479–510 (2014)
50. H. Jeffreys, *An invariant form for the prior probability in estimation problems*, in *Proceedings of the Royal Society* (1946)
51. D.R. Jones, M. Schonlau, W.J. Welch, Efficient global optimization of expensive black-box functions. *J. Global Optim.* **13**, 455–492 (1998)
52. A. Kashinath, A. Misra, M. Demkowicz, Stable storage of helium in nanoscale platelets at semicoherent interfaces. *Phys. Rev. Lett.* **110**, 086101 (2013)
53. M.C. Kennedy, A. O’Hagan, Bayesian calibration of computer models. *J. R. Stat. Soc. Ser. B (Stat. Methodol.)* **63**, 425–464 (2001)

54. J. Kiefer, J. Wolfowitz, Stochastic estimation of the maximum of a regression function. *Ann. Math. Stat.* **23**, 462–466 (1952)
55. J. Kristensen, N.J. Zabarar, Bayesian uncertainty quantification in the evaluation of alloy properties with the cluster expansion method. *Comput. Phys. Commun.* **185**, 2885–2892 (2014)
56. H. Kushner, G. Yin, *Stochastic Approximation and Recursive Algorithms and Applications*, Applications of mathematics (Springer, Berlin, 2003)
57. O.P. Le Maître, O.M. Knio, *Spectral Methods for Uncertainty Quantification: With Applications to Computational Fluid Dynamics* (Springer, Berlin, 2010)
58. J. Lewandowski, A. Greer, Temperature rise at shear bands in metallic glasses. *Nat. Mater.* **5**, 15–18 (2006)
59. J. Li, Y.M. Marzouk, Adaptive construction of surrogates for the Bayesian solution of inverse problems. *SIAM J. Sci. Comput.* **36**, A1163–A1186 (2014)
60. N. Li, E. Fu, H. Wang, J. Carter, L. Shao, S. Maloy, A. Misra, X. Zhang, He ion irradiation damage in Fe/W nanolayer films. *J. Nucl. Mater.* **389**, 233–238 (2009)
61. N. Li, M. Martin, O. Anderoglu, A. Misra, L. Shao, H. Wang, X. Zhang, He ion irradiation damage in Al/Nb multilayers. *J. Appl. Phys.* **105**, 123522 (2009)
62. N. Li, J. Wang, J. Huang, A. Misra, X. Zhang, In situ tem observations of room temperature dislocation climb at interfaces in nanolayered Al/Nb composites. *Scripta Materialia* **63**, 363–366 (2010)
63. D.V. Lindley, *Bayesian Statistics, A Review* (Society for Industrial and Applied Mathematics (SIAM), Philadelphia, 1972)
64. T.J. Loredo, Rotating stars and revolving planets: Bayesian exploration of the pulsating sky, in *Bayesian Statistics 9: Proceedings of the Ninth Valencia International Meeting*, Oxford University Press (2010), pp. 361–392
65. T.J. Loredo, D.F. Chernoff, *Bayesian adaptive exploration, in Statistical Challenges of Astronomy* (Springer, Berlin, 2003), pp. 57–69
66. D.J. MacKay, *Information theory, inference, and learning algorithms*, vol. 7, Citeseer (2003)
67. Y.M. Marzouk, H.N. Najm, Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems. *J. Comput. Phys.* **228**, 1862–1902 (2009)
68. Y.M. Marzouk, H.N. Najm, L.A. Rahn, Stochastic spectral methods for efficient Bayesian solution of inverse problems. *J. Comput. Phys.* **224**, 560–586 (2007)
69. Y.M. Marzouk, D. Xiu, A stochastic collocation approach to Bayesian inference in inverse problems. *Commun. Comput. Phys.* **6**, 826–847 (2009)
70. J.E. Mehner, L.D. Gabbay, S.D. Senturia, *Computer-aided generation of nonlinear reduced-order dynamic macromodels. II. Stress-stiffened case*. *Microelectromech. Syst. J.* **9**, 270–278 (2000)
71. Y. Mishin, M. Asta, J. Li, Atomistic modeling of interfaces and their impact on microstructure and properties. *Acta Materialia* **58**, 1117–1151 (2010)
72. A. Misra, M. Demkowicz, X. Zhang, R. Hoagland, The radiation damage tolerance of ultra-high strength nanolayered composites. *Jom* **59**, 62–65 (2007)
73. A. Misra, J. Hirth, R. Hoagland, Length-scale-dependent deformation mechanisms in incoherent metallic multilayered composites. *Acta Materialia* **53**, 4817–4824 (2005)
74. T.E. Mitchell, Y.C. Lu, A.J.G. Jr, M. Nastasi, H. Kung, Structure and mechanical properties of copper/niobium multilayers. *J. Am. Ceram. Soc.* **80**, 1673–1676 (1997)
75. J.I. Myung, M.A. Pitt, Optimal experimental design for model discrimination. *Psychol. Rev.* **116**, 499–518 (2009)
76. J.A. Nelder, R. Mead, A simplex method for function minimization. *Comput. J.* **7**, 308–313 (1965)
77. A. Noor, J. Peters, Reduced basis technique for nonlinear analysis of structures. *AIAA J.* **18**, 455–462 (1980)

78. G. Odette, M. Alinger, B. Wirth, Recent developments in irradiation-resistant steels. *Annu. Rev. Mater. Res.* **38**, 471–503 (2008)
79. M. Parno, P. Conrad, A. Davis, Y.M. Marzouk, *MIT uncertainty quantification (MUQ) library*. <http://bitbucket.org/mituq/muq>
80. V. Picheny, D. Ginsbourger, Y. Richet, G. Caplin, Quantile-based optimization of noisy computer experiments with tunable precision. *Technometrics* **55**, 2–13 (2013)
81. Z.-Q. Qu, *Model Order Reduction Techniques with Applications in Finite Element Analysis: With Applications in Finite Element Analysis* (Springer Science & Business Media, Berlin, 2004)
82. C. Rasmussen, C. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, Cambridge, 2006)
83. C.P. Robert, G. Casella, *Monte Carlo Statistical Methods* (Springer, Berlin, 2004)
84. K.J. Ryan, Estimating expected information gains for experimental designs with application to the random fatigue-limit model. *J. Comput. Graph. Stat.* **12**, 585–603 (2003)
85. C. Schwab, A.M. Stuart, Sparse deterministic approximation of Bayesian inverse problems. *Inv. Prob.* **28**, 045003 (2012)
86. S. Shao, H. Zbib, I. Mastorakos, D. Bahr, The void nucleation strengths of the cu-ni-nb-based nanoscale metallic multilayers under high strain rate tensile loadings. *Comput. Mater. Sci.* **82**, 435–441 (2014)
87. A. Shapiro, Asymptotic analysis of stochastic programs. *Ann. Oper. Res.* **30**, 169–186 (1991)
88. L. Sirovich, Turbulence and the dynamics of coherent structures. Part 1: coherent structures. *Q. Appl. Math.* **45**, 561–571 (1987)
89. A. Solonen, H. Haario, M. Laine, Simulation-based optimal design using a response variance criterion. *J. Comput. Graph. Stat.* **21**, 234–252 (2012)
90. J.C. Spall, An overview of the simultaneous perturbation method for efficient optimization. *Johns Hopkins APL Tech. Dig.* **19**, 482–492 (1998)
91. A. Sutton, R. Balluffi, *Interfaces in Crystalline Materials*, Monographs on the Physics and Chemistry of Materials (Clarendon Press, Oxford, 1995)
92. H. Ullmaier, The influence of helium on the bulk properties of fusion reactor structural materials. *Nucl. Fus.* **24**, 1039 (1984)
93. J. van den Berg, A. Curtis, J. Trampert, Optimal nonlinear Bayesian experimental design: an application to amplitude versus offset experiments. *Geophy. J. Int.* **155**, 411–421 (2003)
94. A. Vattré, N. Abdollahi, K. Kolluri, M. Demkowicz, Computational design of patterned interfaces using reduced order models. *Sci. Rep.* **4** (2014)
95. Y.-C. Wang, A. Misra, R. Hoagland, Fatigue properties of nanoscale Cu/Nb multilayers. *Scripta Materialia* **54**, 1593–1598 (2006)
96. B.P. Weaver, B.J. Williams, C.M. Anderson-Cook, D.M. Higdon, Computational enhancements to Bayesian design of experiments using Gaussian processes. *Bayesian Analysis* (2015)
97. Q. Wei, N. Li, N. Mara, M. Nastasi, A. Misra, Suppression of irradiation hardening in nanoscale V/Ag multilayers. *Acta Materialia* **59**, 6331–6340 (2011)
98. Q. Wei, A. Misra, Transmission electron microscopy study of the microstructure and crystallographic orientation relationships in v/ag multilayers. *Acta Materialia* **58**, 4871–4882 (2010)
99. B. Williams, D. Higdon, J. Gattiker, L. Moore, M. McKay, S. Keller-McNulty, Combining experimental data and computer simulations, with an application to flyer plate experiments. *Bayesian Anal.* **1**, 765–792 (2006)
100. D. Xiu, Efficient collocational approach for parametric uncertainty analysis. *Commun. Comput. Phys.* **2**, 293–309 (2007)
101. D. Xiu, J.S. Hesthaven, High-order collocation methods for differential equations with random inputs. *SIAM J. Sci. Comput.* **27**, 1118–1139 (2005)
102. L. Yarin, *The Pi-Theorem: Applications to Fluid Mechanics and Heat and Mass Transfer*, vol. 1 (Springer, Berlin, 2012)
103. Y. Yoo, C. Jo, C. Jones, Compositional prediction of creep rupture life of single crystal ni base superalloy by Bayesian neural network, *Materials Science and Engineering*, pp. 22–29 (2001)

104. D. Yuryev, M. Demkowicz, Computational design of solid-state interfaces using o-lattice theory: an application to mitigating helium-induced damage. *Appl. Phys. Lett.* **105**, 221601 (2014)
105. X. Zhang, E. Fu, A. Misra, M. Demkowicz, Interface-enabled defect reduction in he ion irradiated metallic multilayers. *Jom* **62**, 75–78 (2010)
106. S. Zinkle, N. Ghoniem, Operating temperature windows for fusion reactor structural materials. *Fusion Eng. Des.* **51**, 55–71 (2000)

Chapter 3

Bayesian Optimization for Materials Design

Peter I. Frazier and Jialei Wang

Abstract We introduce Bayesian optimization, a technique developed for optimizing time-consuming engineering simulations and for fitting machine learning models on large datasets. Bayesian optimization guides the choice of experiments during materials design and discovery to find good material designs in as few experiments as possible. We focus on the case when materials designs are parameterized by a low-dimensional vector. Bayesian optimization is built on a statistical technique called Gaussian process regression, which allows predicting the performance of a new design based on previously tested designs. After providing a detailed introduction to Gaussian process regression, we describe two Bayesian optimization methods: expected improvement, for design problems with noise-free evaluations; and the knowledge-gradient method, which generalizes expected improvement and may be used in design problems with noisy evaluations. Both methods are derived using a value-of-information analysis, and enjoy one-step Bayes-optimality.

3.1 Introduction

In materials design and discovery, we face the problem of choosing the chemical structure, composition, or processing conditions of a material to meet design criteria. The traditional approach is to use iterative trial and error, in which we (1) choose some material design that we think will work well based on intuition, past experience, or theoretical knowledge; (2) synthesize and test the material in physical experiments; and (3) use what we learn from these experiments in choosing the material design to try next. This iterative process is repeated until some combination of success and exhaustion is achieved.

P.I. Frazier (✉) · J. Wang
School of Operations Research & Information Engineering, Cornell University,
Ithaca, NY 14853, USA
e-mail: pf98@cornell.edu

© Springer International Publishing Switzerland 2016
T. Lookman et al. (eds.), *Information Science for Materials
Discovery and Design*, Springer Series in Materials Science 225,
DOI 10.1007/978-3-319-23871-5_3

While trial and error has been extremely successful, we believe that mathematics and computation together promise to accelerate the pace of materials discovery, not by changing the fundamental iterative nature of materials design, but by improving the choices that we make about which material designs to test, and by improving our ability to learn from previous experimental results.

In this chapter, we describe a collection of mathematical techniques, based on Bayesian statistics and decision theory, for augmenting and enhancing the trial and error process. We focus on one class of techniques, called Bayesian optimization (BO), or Bayesian global optimization (BGO), which use machine learning to build a predictive model of the underlying relationship between the design parameters of a material and its properties, and then use decision theory to suggest which design or designs would be most valuable to try next. The most well-developed Bayesian optimization methods assume that (1) the material is described by a vector of continuous variables, as is the case, e.g., when choosing ratios of constituent compounds, or choosing a combination of temperature and pressure to use during manufacture; (2) we have a single measure of quality that we wish to make as large as possible; and (3) the constraints on feasible materials designs are all known, so that any unknown constraints are incorporated into the quality measure. There is also a smaller body of work on problems that go beyond these assumptions, either by considering discrete design decisions (such as small molecule design), multiple competing objectives, or by explicitly allowing unknown constraints.

Bayesian optimization was pioneered by [1], with early development through the 1970s and 1980s by Mockus and Zilinskas [2, 3]. Development in the 1990s was marked by the popularization of Bayesian optimization by Jones, Schonlau, and Welch, who, building on previous work by Mockus, introduced the Efficient Global Optimization (EGO) method [4]. This method became quite popular and well-known in engineering, where it has been adopted for design applications involving time-consuming computer experiments, within a broader set of methods designed for optimization of expensive functions [5]. In the 2000s, development of Bayesian optimization continued in statistics and engineering, and the 2010s have seen additional development from the machine learning community, where Bayesian optimization is used for tuning hyperparameters of computationally expensive machine learning models [6]. Other introductions to Bayesian optimization may be found in the tutorial article [7] and textbooks [8, 9], and an overview of the history of the field may be found in [10].

We begin in Sect. 3.2 by introducing the precise problem considered by Bayesian Optimization. We then describe in Sect. 3.3 the predictive technique used by Bayesian Optimization, which is called Gaussian Process (GP) regression. We then show, in Sect. 3.4, how Bayesian Optimization recommends which experiments to perform. In Sect. 3.5 we provide an overview of software packages, both freely available and commercial, that implement the Bayesian Optimization methods described in this chapter. We offer closing remarks in Sect. 3.6.

3.2 Bayesian Optimization

Bayesian optimization considers materials designs parameterized by a d -dimensional vector x . We suppose that the space of materials designs in which x takes values is a known set $A \subseteq \mathbb{R}^d$.

For example, $x = (x(1), \dots, x(d))$ could give the ratio of each of d different constituents mixed together to create some aggregate material. In this case, we would choose A to be the set $A = \{x : \sum_{i=1}^d x(i) = 1\}$. As another example, setting $d = 2$, $x = (x(1), x(2))$ could give the temperature ($x(1)$) and pressure ($x(2)$) used in material processing. In this case, we would choose A to be the rectangle bounded by the experimental setup's minimum and maximal achievable temperature on one axis, T_{\min} and T_{\max} , and the minimum and maximum achievable pressure on the other. As a final example, we could let $x = (x(1), \dots, x(d))$ be the temperatures used in some annealing schedule, assumed to be decreasing over time. In this case, we would set A to be the set $\{x : T_{\max} \geq x(1) \geq \dots \geq x(d) \geq T_{\min}\}$.

Let $f(x)$ be the quality of the material with design parameter x . The function f is unknown, and observing $f(x)$ requires synthesizing material design x and observing its quality in a physical experiment. We would like to find a design x for which $f(x)$ is large. That is, we would like to solve

$$\max_{x \in A} f(x). \tag{3.1}$$

This is challenging because evaluating $f(x)$ is typically expensive and time-consuming. While the time and expense depends on the setting, synthesizing and testing a new material design could easily take days or weeks of effort and thousands of dollars of materials.

In Bayesian optimization, we use mathematics to build a predictive model for the function f based on observations of previous materials designs, and then use this predictive model to recommend a materials design that would be most valuable to test next. We first describe this predictive model in Sect. 3.3, which is performed using a machine learning technique called Gaussian process regression. We then describe, in Sect. 3.4, how this predictive model is used to recommend which design to test next.

3.3 Gaussian Process Regression

The predictive piece of Bayesian optimization is based on a machine learning technique called Gaussian process regression. This technique is a Bayesian version of a frequentist technique called kriging, introduced in the geostatistics literature by South-African mining engineer Daniel Krige [11], and popularized later by Matheron and colleagues [12], as described in [13]. A modern monograph on

Gaussian process regression is [14], and a list of software implementing Gaussian process regression may be found at [15].

In Gaussian process regression, we seek to predict $f(x)$ based on observations at previously evaluated points, call them x_1, \dots, x_n . We first treat the case where $f(x)$ can be observed exactly, without noise, and then later treat noise in Sect. 3.3.5. In this noise-free case, our observations are $y_i = f(x_i)$ for $i = 1, \dots, n$.

Gaussian process regression is a Bayesian statistical method, and in Bayesian statistics we perform inference by placing a so-called *prior probability distribution* on unknown quantities of interest. The prior probability distribution is often called, more simply, the *prior distribution* or, even more simply, the *prior*. This prior distribution is meant to encode our intuition or domain expertise regarding which values for the unknown quantity of interest are most likely. We then use Bayes rule, together with any data observed, to calculate a *posterior probability distribution* on these unknowns. For a broader introduction to Bayesian statistics, see the textbook [16] or the research monograph [17].

In Gaussian process regression, if we wish to predict the value of f at a single candidate point x^* , it is sufficient to consider our unknowns to be the values of f at the previously evaluated points, x_1, \dots, x_n , and the new point x^* at which we wish to predict. That is, we take our unknown quantity of interest to be the vector $(f(x_1), \dots, f(x_n), f(x^*))$. We then take our data, which is $f(x_1), \dots, f(x_n)$, and use Bayes rule to calculate a posterior probability distribution on the full vector of interest, $(f(x_1), \dots, f(x_n), f(x^*))$, or, more simply, just on $f(x^*)$.

To calculate the posterior, we must first specify the prior, which Gaussian process regression assumes to be multivariate normal. It calculates the mean vector of this multivariate normal prior distribution using a function, called the *mean function* and written here as $\mu_0(\cdot)$, which takes a single x as an argument. It applies this mean function to each of the points x_1, \dots, x_n, x^* to create an $n + 1$ -dimensional column vector. Gaussian process regression creates the covariance matrix of the multivariate normal prior distribution using another function, called the *covariance function* or *covariance kernel* and written here as $\Sigma_0(\cdot, \cdot)$, which takes a pair of points x, x' as arguments. It applies this covariance function to every pair of points in x_1, \dots, x_n, x to create an $(n + 1) \times (n + 1)$ matrix.

Thus, Gaussian process regression sets the prior probability distribution to,

$$\begin{bmatrix} f(x_1) \\ \vdots \\ f(x_n) \\ f(x^*) \end{bmatrix} \sim \text{Normal} \left(\begin{bmatrix} \mu_0(x_1) \\ \vdots \\ \mu_0(x_n) \\ \mu_0(x^*) \end{bmatrix}, \begin{bmatrix} \Sigma_0(x_1, x_1) & \cdots & \Sigma_0(x_1, x_n) & \Sigma_0(x_1, x^*) \\ \vdots & \ddots & \vdots & \vdots \\ \Sigma_0(x_n, x_1) & \cdots & \Sigma_0(x_n, x_n) & \Sigma_0(x_n, x^*) \\ \Sigma_0(x^*, x_1) & \cdots & \Sigma_0(x^*, x_n) & \Sigma_0(x^*, x^*) \end{bmatrix} \right) \quad (3.2)$$

The subscript “0” in μ_0 and Σ_0 indicate that these functions are relevant to the prior distribution, before any data has been collected.

We now discuss how the mean and covariance functions are chosen, focusing on the covariance function first because it tends to be more important in getting good results from Gaussian process regression.

3.3.1 Choice of Covariance Function

In choosing the covariance function $\Sigma_0(\cdot, \cdot)$, we wish to satisfy two requirements.

The first is that it should encode the belief that points x and x' near each other tend to have more similar values for $f(x)$ and $f(x')$. To accomplish this, we want the covariance matrix in (3.2) to have entries that are larger for pairs of points that are closer together, and closer to 0 for pairs of points that are further apart.

The second is that the covariance function should always produce positive semi-definite covariance matrices in the multivariate normal prior. That is, if Σ is the covariance matrix in (3.2), then we require that $a^T \Sigma a \geq 0$ for all column vectors a (where a is assumed to have the appropriate length, $n + 1$). This requirement is necessary to ensure that the multivariate normal prior distribution is a well-defined probability distribution, because if θ is multivariate normal with mean vector μ and covariance matrix Σ , then the variance of $a \cdot \theta$ is $a^T \Sigma a$, and we require variances to be non-negative.

Several covariance functions satisfy these two requirements. The most commonly used is called the *squared exponential*, or Gaussian kernel, and is given by,

$$\Sigma_0(x, x') = \alpha \exp\left(-\sum_{i=1}^d \beta_i (x_i - x'_i)^2\right). \quad (3.3)$$

This kernel is parameterized by $d + 1$ parameters: α , and β_1, \dots, β_d .

The parameter $\alpha > 0$ controls how much overall variability there is in the function f . We observe that under the prior, the variance of $f(x)$ is $\text{Var}(f(x)) = \text{Cov}(f(x), f(x)) = \alpha$. Thus, when α is large, we are encoding in our prior distribution that $f(x)$ is likely to take a larger range of values.

The parameters $\beta_i > 0$ controls how quickly the function f varies with x . For example, consider the relationship between some point x and another point $x' = x + [1, 0, \dots, 0]$. When β_1 is small (close to 0), the covariance between $f(x)$ and $f(x')$ is $\alpha \exp(-\beta_1) \approx \alpha$, giving a correlation between $f(x)$ and $f(x')$ of nearly 1. This reflects a belief that $f(x)$ and $f(x')$ are likely to be very similar, and that learning the value of $f(x)$ will also teach us a great deal about $f(x')$. In contrast, when β_1 is large, the covariance between $f(x)$ and $f(x')$ is nearly 0, given a correlation between $f(x)$ and $f(x')$ that is also nearly 0, reflecting a belief that $f(x)$ and $f(x')$ are unrelated to each other, and learning something about $f(x)$ will teach us little about (x') .

Going beyond the squared exponential kernel

There are several other possibilities for the covariance kernel beyond the squared exponential kernel, which encode different assumptions about the underlying behavior of the function f . One particularly useful generalization of the squared exponential covariance kernel is the Matérn covariance kernel, which allows more flexibility in modeling the smoothness of f .

Before describing this kernel, let $r = \sqrt{\sum_i \left(\frac{x_i - x'_i}{\beta_i}\right)^2}$ be the Euclidean distance between x and x' , but where we have altered the length scale in each dimension by some strictly positive parameter β_i . Then, the squared exponential covariance kernel can be written as, $\Sigma_0(x, x') = \alpha \exp(-r^2)$.

With this notation, the Matérn covariance kernel is,

$$\Sigma_0(x, x') = \alpha \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu r}\right)^\nu K_\nu\left(\sqrt{2\nu r}\right),$$

where K_ν is the modified Bessel function. If we take the limit as $\nu \rightarrow \infty$, we obtain the squared exponential kernel ([14], Sect. 4.2 p. 85).

The Matérn covariance kernel is useful because it allows modeling the smoothness of f in a more flexible way, as compared with the squared exponential kernel. Under the squared exponential covariance kernel, the function f is infinitely mean-square differentiable,¹ which may not be an appropriate assumption in many applications. In contrast, under the Matérn covariance kernel, f is k -times mean-square differentiable if and only if $\nu > k$. Thus, we can model a function that is twice differentiable but no more by choosing $\nu = 5/2$, and a function that is once differentiable but no more by choosing $\nu = 3/2$.

While the squared exponential and Matérn covariance kernels allow modeling a wide range of behaviors, and together represent a toolkit that will handle a wide variety of applications, there are other covariance kernels. For a thorough discussion of these, see Chap. 4 of [14].

Both the Matérn and squared exponential covariance kernel require choosing parameters. While it certainly is possible for one to choose the parameters α and β_i (and ν in the case of Matérn) based on one's intuition about f , and what kinds of variability f is likely to have in a particular application, it is more common to choose these parameters (especially α and β_i) adaptively, so as to best fit previously observed points. We discuss this more below in Sect. 3.3.6. First, however, we discuss the choice of the mean function.

¹Being “mean-square differentiable” at x in the direction given by the unit vector e_i means that the limit $\lim_{\delta \rightarrow 0} (f(x + \delta e_i) - f(x))/\delta$ exists in mean square. Being “ k -times mean-square differentiable” is defined analogously.

3.3.2 Choice of Mean Function

We now discuss choosing the mean function $\mu_0(\cdot)$. Perhaps the most common choice is to simply set the mean function equal to a constant, μ . This constant must be estimated, along with parameters of the covariance kernel such as α and β_i , and is discussed in Sect. 3.3.6.

Beyond this simple choice, if one believes that there will be trends in f that can be described in a parametric way, then it is useful to include trend terms into the mean function. This is accomplished by choosing

$$\mu_0(x) = \mu + \sum_{j=1}^J \gamma_j \Psi_j(x),$$

where $\Psi_j(\cdot)$ are known functions, and $\gamma_j \in \mathbb{R}$, along with $\mu \in \mathbb{R}$, are parameters that must be estimated.

A common choice for the Ψ_j , if one chooses to include them, are polynomials in x up to some small order. For example, if $d = 2$, so x is two-dimensional, then one might include all polynomials up to second order, $\Psi_1(x) = x_1$, $\Psi_2(x) = x_2$, $\Psi_3(x) = (x_1)^2$, $\Psi_4(x) = (x_2)^2$, $\Psi_5(x) = x_1 x_2$, setting $J = 5$. One recovers the constant mean function by setting $J = 0$.

3.3.3 Inference

Given the prior distribution (3.2) on $f(x_1), \dots, f(x_n), f(x^*)$, and given (noise-free) observations of $f(x_1), \dots, f(x_n)$, the critical step in Gaussian process regression is calculating the posterior distribution on $f(x^*)$. We rely on the following general result about conditional probabilities and multivariate normal distributions. Its proof, which may be found in the Derivations and Proofs section, relies on Bayes rule and algebraic manipulation of the probability density of the multivariate normal distribution.

Proposition 1 *Let θ be a k -dimensional multivariate normal random column vector, with mean vector μ and covariance matrix Σ . Let $k_1 \geq 1, k_2 \geq 1$ be two integers summing to k . Decompose θ, μ and Σ as*

$$\theta = \begin{bmatrix} \theta_{[1]} \\ \theta_{[2]} \end{bmatrix}, \quad \mu = \begin{bmatrix} \mu_{[1]} \\ \mu_{[2]} \end{bmatrix}, \quad \Sigma = \begin{bmatrix} \Sigma_{[1,1]} & \Sigma_{[1,2]} \\ \Sigma_{[2,1]} & \Sigma_{[2,2]} \end{bmatrix},$$

so that $\theta_{[i]}$ and $\mu_{[i]}$ are k_i -column vectors, and $\Sigma_{[i,j]}$ is a $k_i \times k_j$ matrix, for each $i, j = 1, 2$.

If $\Sigma_{1,1}$ and $\Sigma_{2,2}$ are invertible, then, for any $u \in \mathbb{R}^{k_1}$, the conditional distribution of $\theta_{[2]}$ given that $\theta_{[1]} = u$ is multivariate normal with mean

$$\mu_{[2]} + \Sigma_{[2,1]} \Sigma_{[1,1]}^{-1} (u - \mu_{[1]})$$

and covariance matrix

$$\Sigma_{[2,2]} - \Sigma_{[2,1]} \Sigma_{[1,1]}^{-1} \Sigma_{[1,2]}.$$

We use this proposition to calculate the posterior distribution on $f(x^*)$, given $f(x_1), \dots, f(x_n)$.

Before doing so, however, we first introduce some additional notation. We let $y_{1:n}$ indicate the column vector $[y_1, \dots, y_n]^T$, and we let $x_{1:n}$ indicate the sequence of vectors (x_1, \dots, x_n) . We let $f(x_{1:n}) = [f(x_1), \dots, f(x_n)]^T$, and similarly for other functions of x , such as $\mu_0(\cdot)$. We introduce similar additional notation for functions that take pairs of points x, x' , so that $\Sigma(x_{1:n}, x_{1:n})$ is the matrix $\begin{bmatrix} \Sigma_0(x_1, x_1) & \dots & \Sigma_0(x_1, x_n) \\ \vdots & \ddots & \vdots \\ \Sigma_0(x_n, x_1) & \dots & \Sigma_0(x_n, x_n) \end{bmatrix}$, $\Sigma_0(x^*, x_{1:n})$ is the row vector $[\Sigma_0(x^*, x_1), \dots, \Sigma_0(x^*, x_n)]$, and $\Sigma_0(x_{1:n}, x^*)$ is the column vector $[\Sigma_0(x_1, x^*), \dots, \Sigma_0(x_n, x^*)]^T$.

This notation allows us to rewrite (3.2) as

$$\begin{bmatrix} y_{1:n} \\ f(x^*) \end{bmatrix} = \text{Normal} \left(\begin{bmatrix} \mu_0(x_{1:n}) \\ \mu_0(x^*) \end{bmatrix}, \begin{bmatrix} \Sigma_0(x_{1:n}, x_{1:n}) & \Sigma_0(x_{1:n}, x^*) \\ \Sigma_0(x^*, x_{1:n}) & \Sigma_0(x^*, x^*) \end{bmatrix} \right). \quad (3.4)$$

We now examine this expression in the context of Proposition 1. We set $\theta_{[1]} = f(x_{1:n})$, $\theta_{[2]} = f(x^*)$, $\mu_{[1]} = \mu_0(x_{1:n})$, $\mu_{[2]} = \mu_0(x^*)$, $\Sigma_{[1,1]} = \Sigma_0(x_{1:n}, x_{1:n})$, $\Sigma_{[1,2]} = \Sigma_0(x_{1:n}, x^*)$, $\Sigma_{[2,1]} = \Sigma_0(x^*, x_{1:n})$, and $\Sigma_{[2,2]} = \Sigma_0(x^*, x^*)$.

Then, applying Proposition 1, we see that the posterior distribution on $f(x^*)$ given observations $y_i = f(x_i)$, $i = 1, \dots, n$ is normal, with a mean $\mu_n(x^*)$ and variance $\sigma_n^2(x^*)$ given by,

$$\mu_n(x^*) = \mu_0(x^*) + \Sigma_0(x^*, x_{1:n}) \Sigma_0(x_{1:n}, x_{1:n})^{-1} (f(x_{1:n}) - \mu_0(x_{1:n})), \quad (3.5)$$

$$\sigma_n^2(x^*) = \Sigma_0(x^*, x^*) - \Sigma_0(x^*, x_{1:n}) \Sigma_0(x_{1:n}, x_{1:n})^{-1} \Sigma_0(x_{1:n}, x^*). \quad (3.6)$$

The invertibility of $\Sigma_0(x_{1:n}, x_{1:n})$ (and also $\Sigma_0(x^*, x^*)$) required by Proposition 1 depends on the covariance kernel and its parameters (typically called hyperparameters), but this invertibility typically holds as long as these hyperparameters satisfy mild non-degeneracy conditions, and the $x_{1:n}$ are distinct, i.e., that we have not measured the same point more than once. For example, under the squared exponential covariance kernel, invertibility holds as long as $\alpha > 0$ and the $x_{1:n}$ are distinct. If we have measured a point multiple times, then we can safely drop all but one of the measurements, here where observations are noise-free. Below, we treat the case where observations are noisy, and in this case including multiple measurements of the same point is perfectly reasonable and does not cause issues.




Fig. 3.1 Illustration of Gaussian process regression with noise-free evaluations. The *circles* show previously evaluated points, $(x_i, f(x_i))$. The *solid line* shows the posterior mean, $\mu_n(x)$, as a function of x , which is an estimate $f(x)$, and the *dashed lines* show a Bayesian credible interval for each $f(x)$, calculated as $\mu_n(x) \pm 1.96\sigma_n(x)$. Although this example shows f taking a scalar input, Gaussian process regression can be used for functions with vector inputs

Figure 3.1 shows the output from Gaussian process regression. In the figure, circles show points $(x_i, f(x_i))$, the solid line shows $\mu_n(x^*)$ as a function of x^* , and the dashed lines are positioned at $\mu_n(x^*) \pm 1.96\sigma_n(x^*)$, forming a 95 % Bayesian credible interval for $f(x^*)$, i.e., an interval in which $f(x^*)$ lies with posterior probability 95 %. (A credible interval is the Bayesian version of a frequentist confidence interval.) Because observations are noise-free, the posterior mean $\mu_n(x^*)$ interpolates the observations $f(x^*)$.

3.3.4 Inference with Just One Observation

The expressions (3.5) and (3.6) are complex, and perhaps initially difficult to assimilate. To give more intuition about them, and also to support some additional analysis below in Sect. 3.4, it is useful to consider the simplest case, when we have just a single measurement, $n = 1$.

In this case, all matrices in (3.5) and (3.6) are scalars, $\Sigma_0(x^*, x_1) = \Sigma_0(x_1, x^*)$, and the expressions (3.5) and (3.6) can be rewritten as,

$$\mu_1(x^*) = \mu_0(x^*) + \frac{\Sigma_0(x^*, x_1)}{\Sigma_0(x_1, x_1)}(f(x_1) - \mu_0(x_1)), \quad (3.7)$$

$$\sigma_1^2(x^*) = \Sigma_0(x^*, x^*) - \frac{\Sigma_0(x^*, x_1)^2}{\Sigma_0(x_1, x_1)}. \quad (3.8)$$

Intuition about the expression for the posterior mean

We first examine (3.7). We see that the posterior mean of $f(x^*)$, $\mu_1(x^*)$, which we can think of as our estimate of $f(x^*)$ after observing $f(x_1)$, is obtained by taking our original estimate of $f(x^*)$, $\mu_0(x^*)$, and adding to it a correction term. This correction term is itself the product of two quantities: the error $f(x_1) - \mu_0(x_1)$ in our original estimate of $f(x_1)$, and the quantity $\frac{\Sigma_0(x^*, x_1)}{\Sigma_0(x_1, x_1)}$. Typically, $\Sigma_0(x^*, x_1)$ will be positive, and hence also $\frac{\Sigma_0(x^*, x_1)}{\Sigma_0(x_1, x_1)}$. (Recall, $\Sigma_0(x_1, x_1)$ is a variance, so is never negative.) Thus, if $f(x_1)$ is bigger than expected, $f(x_1) - \mu_0(x_1)$ will be positive, and our posterior mean $\mu_1(x^*)$ will be larger than our prior mean $\mu_0(x^*)$. In contrast, if $f(x_1)$ is smaller than expected, $f(x_1) - \mu_0(x_1)$ will be negative, and our posterior mean $\mu_1(x^*)$ will be smaller than our prior mean $\mu_0(x^*)$.

We can examine the quantity $\frac{\Sigma_0(x^*, x_1)}{\Sigma_0(x_1, x_1)}$ to understand the effect of the position of x^* relative to x_1 on the magnitude of the correction to the posterior mean. Notice that x^* only enters this expression through the numerator. If x^* is close to x_1 , then $\Sigma_0(x^*, x_1)$ will be large under the squared exponential and most other covariance kernels, and positive values for $f(x_1) - \mu_0(x_1)$ will also cause a strong positive change in $\mu_1(x^*)$ relative to $\mu_0(x^*)$. If x^* is far from x_1 , then $\Sigma_0(x^*, x_1)$ will be close to 0, and $f(x_1) - \mu_0(x_1)$ will have little effect on $\mu_1(x^*)$.

Intuition about the expression for the posterior variance

Now we examine (3.8). We see that the variance of our belief on $f(x^*)$ under the posterior, $\sigma_1^2(x^*)$, is smaller than its value under the prior, $\Sigma_0(x^*, x^*)$. Moreover, when x^* is close to x_1 , $\Sigma_0(x^*, x_1)$ will be large, and the reduction in variance from prior to posterior will also be large.

Conversely, when x^* is far from x_1 , $\Sigma_0(x^*, x_1)$ will be close to 0, and the variance under the posterior will be similar to its value under the prior.

As a final remark, we can also rewrite the expression (3.8) in terms of the squared correlation under the prior, $\text{Corr}(f(x^*), f(x_1))^2 = \Sigma_0(x^*, x_1)^2 / (\Sigma_0(x^*, x^*) \Sigma_0(x_1, x_1)) \in [0, 1]$, as

$$\sigma_1^2(x^*) = \Sigma_0(x^*, x^*) (1 - \text{Corr}(f(x^*), f(x_1))^2).$$

We thus see that the reduction in variance of the posterior distribution depends on the squared correlation under the prior, with larger squared correlation implying a larger reduction.

3.3.5 Inference with Noisy Observations

The previous section assumed that $f(x^*)$ can be observed exactly, without any error. When $f(x^*)$ is the outcome of a physical experiment, however, our observations are obscured by noise. Indeed, if we were to synthesize and test the same material design x^* multiple times, we might observe different results.

To model this situation, Gaussian process regression can be extended to allow observations of the form,

$$y(x_i) = f(x_i) + \varepsilon_i,$$

where we assume that the ε_i are normally distributed with mean 0 and constant variance, λ^2 , with independence across i . In general, the variance λ^2 is unknown, but we treat it as a known parameter of our model, and then estimate it along with all the other parameters of our model, as discussed below in Sect. 3.3.6.

These assumptions of constant variance (called *homoscedasticity*) and independence make the analysis significantly easier, although they are often violated in practice. Experimental conditions that tend to violate these assumptions are discussed below, as are versions of GP regression that can be used when they are violated.

Analysis of independent homoscedastic noise

To perform inference under independent homoscedastic noise, and calculate a posterior distribution on the value of the function $f(x_*)$ at a given point x_* , our first step is to write down the joint distribution of our observations y_1, \dots, y_n and the quantity we wish to predict, $f(x_*)$, under the prior. That is, we write down the distribution of the vector $[y_1, \dots, y_n, f(x_*)]$.

We first observe that $[y_1, \dots, y_n, f(x_*)]$ is the sum of $[f(x_1), \dots, f(x_n), f(x_*)]$ and another vector, $[\varepsilon_1, \dots, \varepsilon_n, 0]$. The first vector has a multivariate normal distribution given by (3.4). The second vector is independent of the first and is also multivariate normal, with a mean vector that is identically 0, and a covariance matrix $\text{diag}(\lambda^2, \dots, \lambda^2, 0)$. The sum of two independent multivariate normal random vectors is itself multivariate normal, with a mean vector and covariance matrix given, respectively, by the sums of the mean vectors and covariance matrices of the summands. This gives the distribution of $[y_1, \dots, y_n, f(x_*)]$ as

$$\begin{bmatrix} y_{1:n} \\ f(x^*) \end{bmatrix} \sim \text{Normal} \left(\begin{bmatrix} \mu_0(x_{1:n}) \\ \mu_0(x^*) \end{bmatrix}, \begin{bmatrix} \Sigma_0(x_{1:n}, x_{1:n}) + \lambda^2 I_n & \Sigma_0(x_{1:n}, x^*) \\ \Sigma_0(x^*, x_{1:n}) & \Sigma_0(x^*, x^*) \end{bmatrix} \right), \quad (3.9)$$

where I_n is the n -dimensional identity matrix.

As we did in Sect. 3.3.3, we can use Proposition 1 with the above expression to compute the posterior on $f(x^*)$ given $f(x_{1:n})$. We obtain,

$$\mu_n(x^*) = \mu_0(x^*) + \Sigma_0(x^*, x_{1:n}) \left[\Sigma_0(x_{1:n}, x_{1:n}) + \lambda^2 I_n \right]^{-1} (y_{1:n} - \mu_0(x_{1:n})) \quad (3.10)$$

$$\sigma_n^2(x^*) = \Sigma_0(x^*, x^*) - \Sigma_0(x^*, x_{1:n}) \left[\Sigma_0(x_{1:n}, x_{1:n}) + \lambda^2 I_n \right]^{-1} \Sigma_0(x_{1:n}, x^*). \quad (3.11)$$

If we set $\lambda^2 = 0$, so there is no noise, then we recover (3.5) and (3.6).




Fig. 3.2 Illustration of Gaussian process regression with noisy evaluations. As in Fig. 3.1, the *circles* show previously evaluated points, (x_i, y_i) , where y_i is $f(x_i)$ perturbed by constant-variance independent noise. The *solid line* shows the posterior mean, $\mu_n(x)$, as a function of x , which is an estimate of the underlying function f , and the *dashed lines* show a Bayesian credible interval for f , calculated as $\mu_n(x) \pm 1.96\sigma_n(x)$

Figure 3.2 shows an example of a posterior distribution calculated with Gaussian process regression with noisy observations. Notice that the posterior mean no longer interpolates the observations, and the credible interval has a strictly positive width at points where we have measured. Noise prevents us from observing function values exactly, and so we remain uncertain about the function value at points we have measured.

Going beyond homoscedastic independent noise

Constant variance is violated if the experimental noise differs across materials designs, which occurs most frequently when noise arises during the synthesis of the material itself, rather than during the evaluation of a material that has already been created. Some work has been done to extend Gaussian process regression to flexibly model heteroscedastic noise (i.e., noise whose variance changes) [18–21]. The main idea in much of this work is to use a second Gaussian process to model the changing variance across the input domain. Much of this work assumes that the noise is independent and Gaussian, though [21] considers non-Gaussian noise.

Independence is most typically violated, in the context of physical experiments, when the synthesis and evaluation of multiple materials designs is done together, and the variation in some shared component simultaneously influences these designs, e.g., through variation in the temperature while the designs are annealing together, or through variation in the quality of some constituent used in synthesis. We are aware of relatively little work modeling dependent noise in the context of Gaussian process regression and Bayesian optimization, with one exception being [22].

3.3.6 Parameter Estimation

The mean and covariance functions contain several parameters. For example, if we use the squared exponential kernel, a constant mean function, and observations have independent homoscedastic noise, then we must choose or estimate the parameters $\mu, \alpha, \beta_1, \dots, \beta_d, \lambda$. These parameters are typically called *hyperparameters* because they are parameters of the prior distribution. (λ^2 is actually a parameter of the likelihood function, but it is convenient to treat it together with the parameters of the prior.) While one may simply choose these hyperparameters directly, based on intuition about the problem, a more common approach is to choose them adaptively, based on data.

To accomplish this, we write down an expression for the probability of the observed data $y_{1:n}$ in terms of the hyperparameters, marginalizing over the uncertainty on $f(x_{1:n})$. Then, we optimize this expression over the hyperparameters to find settings that make the observed data as likely as possible. This approach to setting hyperparameters is often called *empirical Bayes*, and it can be seen as an approximation to full Bayesian inference.

We detail this approach for the squared exponential kernel with a constant mean function. Estimation for other kernels and mean functions is similar. Using the probability distribution of $y_{1:n}$ from (3.9), and neglecting constants, the natural logarithm of this probability, $\log p(y_{1:n} \mid x_{1:n})$ (called the “log marginal likelihood”), can be calculated as

$$-\frac{1}{2}(y_{1:n} - \mu)^T \left(\Sigma_0(x_{1:n}, x_{1:n}) + \lambda^2 I_n \right)^{-1} (y_{1:n} - \mu) - \frac{1}{2} \log |\Sigma_0(x_{1:n}, x_{1:n}) + \lambda^2 I_n|,$$

where $|\cdot|$ applied to a matrix indicates the determinant.

To find the hyperparameters that maximize this log marginal likelihood (the neglected constant does not affect the location of the maximizer), we will take partial derivatives with respect to each hyperparameter. We will then use them to find maximizers of μ and $\sigma^2 := \alpha + \lambda^2$ analytically, and then use gradient-based optimization to maximize the other hyperparameters.

Taking a partial derivative with respect to μ , setting it to zero, and solving for μ , we get that the value of μ that maximizes the marginal likelihood is

$$\hat{\mu} = \frac{\sum_{i=1}^n \left((\Sigma_0(x_{1:n}, x_{1:n}) + \lambda^2 I_n)^{-1} y_{1:n} \right)_i}{\sum_{i,j=1}^n \left(\Sigma_0(x_{1:n}, x_{1:n}) + \lambda^2 I_n \right)^{-1}_{ij}}.$$

Define R as the matrix with components

$$R_{ij} = \begin{cases} 1 & i = j, \\ g \exp \left(- \sum_{i=1}^d \beta_i (x_i - x_j)^2 \right) & i \neq j, \end{cases}$$

where $g = \frac{\alpha}{\sigma^2}$. Then $\Sigma_0(x_{1:n}, x_{1:n}) + \lambda^2 I_n = \sigma^2 R$ and $\hat{\mu}$ can be written in terms of R as $\hat{\mu} = \frac{\sum_{i=1}^n (R^{-1} y_{1:n})_i}{\sum_{i,j=1}^n R_{ij}^{-1}}$. The log marginal likelihood (still neglecting constants) becomes

$$\log p(y_{1:n} | x_{1:n}) \sim -\frac{1}{2} (y_{1:n} - \hat{\mu})^T (\sigma^2 R)^{-1} (y_{1:n} - \hat{\mu}) - \frac{1}{2} \log |\sigma^2 R|.$$

Taking the partial derivative with respect to σ^2 , and noting that $\hat{\mu}$ does not depend on σ^2 , we solve for σ^2 and obtain

$$\widehat{\sigma^2} = \frac{1}{n} (y_{1:n} - \hat{\mu})^T R^{-1} (y_{1:n} - \hat{\mu}).$$

Substituting this estimate, the log marginal likelihood becomes

$$\log p(y_{1:n} | x_{1:n}) \sim -\log \left(\frac{1}{n} |R|^{\frac{1}{n}} (y_{1:n} - \hat{\mu})^T R^{-1} (y_{1:n} - \hat{\mu}) \right). \quad (3.12)$$


The expression (3.12) cannot in general be optimized analytically. Instead, one typically optimizes it numerically using a first- or second-order optimization algorithm, such as Newton's method or gradient descent, obtaining estimates for β_1, \dots, β_d and g . These estimates are in turn substituted to provide an estimate of R , from which estimates $\hat{\mu}$ and $\widehat{\sigma^2}$ may be computed. Finally, using $\widehat{\sigma^2}$ and the estimated value of g , we may estimate α and λ .

3.3.7 Diagnostics

When using Gaussian process regression, or any other machine learning technique, it is advisable to check the quality of the predictions, and to assess whether the assumptions made by the method are met. One way to do this is illustrated by Fig. 3.3, which comes from a simulation of blood flow near the heart, based on [23], for which we get exact (not noisy) observations of $f(x)$.

This plot is created with a technique called leave-one-out cross validation. In this technique, we iterate through the datapoints $x_{1:n}, y_{1:n}$, and for each $i \in \{1, \dots, n\}$, we train a Gaussian process regression model on all of the data *except* x_i, y_i , and then use it, together with x_i , to predict what the value y_i should be. We obtain from this a posterior mean (the prediction), call it $\mu_{-i}(x_i)$, and also a posterior standard deviation, call it $\sigma_{-i}(x_i)$. When calculating these estimates, it is best to separately re-estimate the hyperparameters each time, leaving out the data (x_i, y_i) . We then calculate a 95% credible interval $\mu_{-i}(x_i) \pm 2\sigma_{-i}(x_i)$, and create Fig. 3.3 by plotting "Predicted" versus "Actual", where the "Actual" coordinate (on the x-axis) is y_i , and the "Predicted" value (on the y-axis) is pictured as an error bar centered at $\mu_{-i}(x_i)$ with half-width $2\sigma_{-i}(x_i)$.

Fig. 3.3 Diagnostic plot for Gaussian process regression, created with leave-one-out cross validation. For each point in our dataset, we hold that point (x_i, y_i) out, train on the remaining points, calculate a 95% credible interval for y_i , and plot this confidence interval as an error bar whose x-coordinate is the actual value y_i . If Gaussian process regression is working well, 95% of the *error bars* will intersect the diagonal line Predicted=Actual



If the uncertainty estimates outputted by Gaussian process regression are behaving as anticipated, then approximately 95% of the credible intervals will intersect the diagonal line Predicted=Actual. Moreover, if Gaussian process regression's predictive accuracy is high, then the credible intervals will be short, and their centers will be close to this same line Predicted=Actual.

This idea may be extended to noisy function evaluations, under the assumption of independent homoscedastic noise. To handle the fact that the same point may be sampled multiple times, let $m(x)$ be the number of times that a point $x \in \{x_1, \dots, x_n\}$ was sampled, and let $\bar{y}(x)$ be the average of the observed values at this point. Moreover, by holding out all $m(x)$ samples of x and training Gaussian process regression, we would obtain a normal posterior distribution on $f(x_i)$ that has mean $\mu_{-i}(x_i)$ and standard deviation $\sigma_{-i}(x_i)$.

Since $\bar{y}(x_i)$ is then the sum of $f(x_i)$ and some normally distributed noise with mean 0 and variance $\lambda^2/m(x_i)$, the resulting distribution of $\bar{y}(x_i)$ is normal with mean $\mu_{-i}(x_i)$ and standard deviation $\sqrt{\sigma_{-i}^2(x_i) + \lambda^2/m(x_i)}$.

From this, a 95% credible interval for $\bar{y}(x_i)$ is then $\mu_{-i}(x_i) \pm 2\sqrt{\sigma_{-i}^2(x_i) + \lambda^2/m(x_i)}$. We would plot Predicted versus Observed by putting this credible interval along the y-axis at x-coordinate $\bar{y}(x_i)$. If Gaussian process regression is working well, then approximately 95% of these credible intervals will intersect the line Predicted=Observed.

For Gaussian process regression to best support Bayesian optimization, it is typically most important to have good uncertainty estimates, and relatively less important to have high predictive accuracy. This is because Bayesian optimization uses Gaussian process regression as a guide for deciding where to sample, and so if Gaussian process regression reports that there is a great deal of uncertainty at a particular location and thus low predictive accuracy, Bayesian optimization can choose to sample at this location to improve accuracy. Thus, Bayesian optimization has

a recourse for dealing with low predictive accuracy, as long as the uncertainty is accurately reported. In contrast, if Gaussian process regression estimates poor performance at a location that actually has near-optimal performance, and also provides an inappropriately low error estimate, then Bayesian optimization may not sample there within a reasonable timeframe, and thus may never correct the error.

If either the uncertainty is incorrectly estimated, or the predictive accuracy is unsatisfactorily low, then the most common “fixes” employed are to adopt a different covariance kernel, or to transform the objective function f . If the objective function is known to be non-negative, then the transformations $\log(f)$ and \sqrt{f} are convenient for optimization because they are both strictly increasing, and so do not change the set of maximizers (or minimizers). If f is not non-negative, but is bounded below by some other known quantity a , then one may first shift f upward by a .

3.3.8 Predicting at More Than One Point

Below, to support the development of the knowledge-gradient method in Sects. 3.4.2 and 3.6, it will be useful to predict the value of f at multiple points, x_1^*, \dots, x_k^* , with noise. To do so, we could certainly apply (3.10) and (3.11) separately for each x_1^*, \dots, x_k^* , and this would provide us with both an estimate (the posterior mean) and a measure of the size of the error in this estimate (the posterior variance) associated with each $f(x_i^*)$. It would not, however, quantify the relationship between the errors at several different locations. For this, we must perform the estimation jointly.

As we did in Sect. 3.3.5, we begin with our prior on $[y_{1:n}, f(x_{1:k}^*)]$, which is,

$$\begin{bmatrix} y_{1:n} \\ f(x_{1:k}^*) \end{bmatrix} \sim \text{Normal} \left(\begin{bmatrix} \mu_0(x_{1:n}) \\ \mu_0(x_{1:k}^*) \end{bmatrix}, \begin{bmatrix} \Sigma_0(x_{1:n}, x_{1:n}) + \lambda^2 I_n & \Sigma_0(x_{1:n}, x_{1:k}^*) \\ \Sigma_0(x_{1:k}^*, x_{1:n}) & \Sigma_0(x_{1:k}^*, x_{1:k}^*) \end{bmatrix} \right),$$

We then use Proposition 1 to compute the posterior on $f(x_{1:k}^*)$ given $f(x_{1:n})$, which is multivariate normal with mean vector $\mu_n(x_{1:k}^*)$ and covariance matrix $\Sigma_n(x_{1:k}^*, x_{1:k}^*)$ given by,

$$\mu_n(x_{1:k}^*) = \mu_0(x_{1:k}^*) + \Sigma_0(x_{1:k}^*, x_{1:n}) \left[\Sigma_0(x_{1:n}, x_{1:n}) + \lambda^2 I_n \right]^{-1} (y_{1:n} - \mu_0(x_{1:n})), \quad (3.13)$$

$$\Sigma_n(x_{1:k}^*, x_{1:k}^*) = \Sigma_0(x_{1:k}^*, x_{1:k}^*) - \Sigma_0(x_{1:k}^*, x_{1:n}) \left[\Sigma_0(x_{1:n}, x_{1:n}) + \lambda^2 I_n \right]^{-1} \Sigma_0(x_{1:n}, x_{1:k}^*). \quad (3.14)$$

We see that setting $k = 1$ provides the expressions (3.10) and (3.11) from Sect. 3.3.5.

3.3.9 Avoiding Matrix Inversion

The expressions (3.10) and (3.11) for the posterior mean and variance in the noisy case, and also (3.7) and (3.8) in the noise-free case, include a matrix inversion term. Calculating this matrix inversion is slow and can be hard to accomplish accurately in practice, due to the finite precision of floating point implementations. Accuracy is especially an issue when Σ has terms that are close to 0, which arises when data points are close together.

In practice, rather than calculating a matrix inverse directly, it is typically faster and more accurate to use a mathematically equivalent algorithm, which performs a Cholesky decomposition and then solves a linear system. This algorithm is described below, and is adapted from Algorithm 2.1 in Sect. 2.3 of [14]. This algorithm also computes the log marginal likelihood required for estimating hyperparameters in Sect. 3.3.6.

Algorithm 1 Implementation using Cholesky decomposition

Require: $x_{1:n}$ (inputs), $y_{1:n}$ (responses), $\Sigma_0(x, x')$ (covariance function), λ^2 (variance of noise), x^* (test input).

- 1: $L = \text{Cholesky}(\Sigma_0(x_{1:n}, x_{1:n}) + \lambda^2 I_n)$
 - 2: $\delta = L^T \setminus (L \setminus (y_{1:n} - \mu_0(x_{1:n})))$
 - 3: $\mu_n(x^*) = \mu_0(x^*) + \Sigma_0(x^*, x_{1:n})\delta$
 - 4: $v = L \setminus \Sigma_0(x_{1:n}, x^*)$
 - 5: $\sigma_n^2(x^*) = \Sigma_0(x^*, x^*) - v^T v$
 - 6: $\log p(y_{1:n} | x_{1:n}) = -\frac{1}{2} (y_{1:n} - \mu_0(x_{1:n}))^T \alpha - \sum_i \log L_{ii} - \frac{n}{2} \log 2\pi$
 - 7: **return** $\mu_n(x^*)$ (mean), $\sigma_n^2(x^*)$ (variance), $\log p(y_{1:n} | x_{1:n})$ (log marginal likelihood).
-

3.4 Choosing Where to Sample

Being able to infer the value of the objective function $f(x)$ at unevaluated points based on past data $x_{1:n}, y_{1:n}$ is only one part of finding good designs. The other part is using this ability to make good decisions about where to direct future sampling.

Bayesian optimization methods address this by using a measure of the value of the information that would be gained by sampling at a point. Bayesian optimization methods then choose the point to sample next for which this value is largest. A number of different ways of measuring the value of information have been proposed. Here, we describe two in detail, expected improvement [2, 4], and the knowledge gradient [24, 25], and then survey a broader collection of design criteria.

3.4.1 Expected Improvement

Expected improvement, as it was first proposed, considered only the case where measurements are free from noise. In this setting, suppose we have taken n measurements at locations $x_{1:n}$ and observed $y_{1:n}$. Then

$$f_n^* = \max_{i=1,\dots,n} f(x_i)$$

is the best value observed so far. Suppose we are considering evaluating f at a new point x . After this evaluation, the best value observed will be

$$f_{n+1}^* = \max(f(x), f_n^*),$$

and the difference between these values, which is the *improvement* due to sampling, is

$$f_{n+1}^* - f_n^* = \max(f(x) - f_n^*, 0) = (f(x) - f_n^*)^+,$$

where $a^+ = \max(a, 0)$ indicates the positive part function.

Ideally, we would choose x to make this improvement as large as possible. Before actually evaluating $f(x)$, however, we do not know what this improvement will be, so we cannot implement this strategy. However, we do have a probability distribution on $f(x)$, from Gaussian process regression. The *expected improvement*, indicated $EI(x)$, is obtained by taking the expectation of this improvement with respect to the posterior distribution on $f(x)$ given $x_{1:n}, y_{1:n}$.

$$EI(x) = E_n[(f(x) - f_n^*)^+], \quad (3.15)$$

where $E_n[\cdot] = E[\cdot | x_{1:n}, y_{1:n}]$ indicates the expectation with respect to the posterior distribution.

The expectation in (3.15) can be computed more explicitly, in terms of the normal cumulative distribution function (cdf) $\Phi(\cdot)$, and the normal probability density function (pdf) $\varphi(\cdot)$. Recalling from Sect. 3.3.3 that $f(x) \sim \text{Normal}(\mu_n(x), \sigma_n^2(x))$, where $\mu_n(x)$ and $\sigma_n^2(x)$ are given by (3.5) and (3.6), and integrating with respect to the normal distribution (a derivation may be found in the Derivations and Proofs section), we obtain,

$$EI(x) = (\mu_n(x) - f_n^*)\Phi\left(\frac{\mu_n(x) - f_n^*}{\sigma_n(x)}\right) + \sigma_n(x)\varphi\left(\frac{\mu_n(x) - f_n^*}{\sigma_n(x)}\right). \quad (3.16)$$

Figure 3.4 plots this expected improvement for a problem with a one-dimensional input space. We can see from this plot that the expected improvement is largest at locations where both the posterior mean $\mu_n(x)$ is large, and also the posterior standard deviation $\sigma_n(x)$ is large. This is reasonable because those points that are most likely to provide large gains are those points that have a high predicted value, but that also





Fig. 3.4 *Upper panel* shows the posterior distribution in a problem with no noise and a one-dimensional input space, where the *circles* are previously measured points, the *solid line* is the posterior mean $\mu_n(x)$, and the *dashed lines* are at $\mu_n(x) \pm 2\sigma_n(x)$. *Lower panel* shows the expected improvement $EI(x)$ computed from this posterior distribution. An “x” is marked at the point with the largest expected improvement, which is where we would evaluate next

have significant uncertainty. Indeed, at points where we have already observed, and thus have no uncertainty, the expected improvement is 0. This is consistent with the idea that, in a problem without noise, there is no value to repeating an evaluation that has already been performed.

This idea of favoring points that, on the one hand, have a large predicted value, but, on the other hand, have a significant amount of uncertainty, is called the *exploration versus exploitation tradeoff*, and appears in areas beyond Bayesian optimization, especially in reinforcement learning [26, 27] and multi-armed bandit problems [28, 29]. In these problems, we are taking actions repeatedly over time whose payoffs are uncertain, and wish to simultaneously get good immediate rewards, while learning the reward distributions for all actions to allow us to get better rewards in the future. We emphasize, however, that the correct balance between exploration and exploitation is different in Bayesian optimization as compared with multi-armed bandits, and should more favor exploration: in optimization, the advantage of measuring where the predicted value is high is that these areas tend to give more useful information about where the optimum lies; in contrast, in problems where we must “learn while doing” like multi-armed bandits, evaluating an action with high predicted reward is good primarily because it tends to give a high immediate reward.

Fig. 3.5 Contour plot of the expected improvement, as a function of the difference in means $\Delta_n(x) := \mu_n(x) - f_n^*$ and the posterior standard deviation $\sigma_n(x)$. The expected improvement is larger when the difference in means is larger, and when the standard deviation is larger



We can also see the exploration versus exploitation tradeoff implicit in the expected improvement function in the contour plot, Fig. 3.5. This plot shows the contours of $EI(x)$ as a function of the posterior mean, expressed as a difference from the previous best, $\Delta_n(x) := \mu_n(x) - f_n^*$, and the posterior standard deviation $\sigma_n(x)$.

Given the expression (3.16), Bayesian optimization algorithms based on expected improvement, such as the Efficient Global Optimization (EGO) algorithm proposed by [4], and the earlier algorithms of Mockus (see, e.g., the monograph [2]), then recommend sampling at the point with the largest expected improvement. That is,

$$x_{n+1} \in \underset{x}{\operatorname{argmax}} EI(x). \quad (3.17)$$

Finding the point with largest expected improvement is itself a global optimization problem, like the original problem that we wished to solve (3.1). Unlike (3.1), however, $EI(x)$ can be computed quickly, and its first and second derivatives can also be computed quickly. Thus, we can expect to be able to solve (3.1) relatively well using an off-the-shelf optimization method for continuous global optimization. A common approach is to use a local solver for continuous optimization, such as gradient ascent, in a multistart framework, where we start the local solver from many starting points chosen at random, and then select the best local solution discovered. In Sect. 3.5 we describe several codes that implement expected improvement methods, and each makes its own choice about how to solve (3.17).

The algorithm given by (3.17) is *optimal* under three assumptions: (1) that we will take only a single sample; (2) there is no noise in our samples; and (3) that the x we will report as our final solution (i.e., the one that we will *implement*) must be among those previously sampled.

In practice, assumption (1) is violated, as Bayesian optimization methods like (3.17) are applied iteratively, and is made simply because it simplifies the analysis. Being able to handle violations of assumption (1) in a more principled way is of great interest to researchers working on Bayesian optimization methodology, and some partial progress in that direction is discussed in Sect. 3.4.3. Assumption

(2) is also often violated in a broad class of applications, especially those involving physical experiments or stochastic simulations. In the next section, we present an algorithm, the knowledge-gradient algorithm [24, 25], that relaxes this assumption (2), and also allows relaxing assumption (3) if this is desired.

3.4.2 Knowledge Gradient

When we have noise in our samples, the derivation of expected improvement meets with difficulty. In particular, if we have noise, then $f_n^* = \max_{i=1,\dots,n} f(x_i)$ is not precisely known, preventing us from using the expression (3.16).

One may simply take a quantity like $\max_{i=1,\dots,n} y_i$ that is similar in spirit to $f_n^* = \max_{i=1,\dots,n} f(x_i)$, and replace f_n^* in (3.16) with this quantity, but the resulting algorithm is no longer justified by an optimality analysis. Indeed, for problems with a great deal of noise, $\max_{i=1,\dots,n} y_i$ tends to be significantly larger than the true underlying value of the best point previously sampled, and so the resulting algorithm may be led to make a poor tradeoff between exploration and exploitation, and exhibit poor performance in such situations.

Instead, the knowledge-gradient algorithm [24, 25] takes a more principled approach, and starts where the derivation of expected improvement began, but fully accounts for the introduction of noise (assumption 2 in Sect. 3.4.1), and the possibility that we wish to search over a class of solutions broader than just those that have been previously evaluated when recommending the final solution (assumption 3 in Sect. 3.4.1).

We first introduce a set A_n , which is the set of points from which we would choose the final solution, if we were asked to recommend a final solution at time n , based on $x_{1:n}, y_{1:n}$. For tractability, we suppose A_n is finite. For example, if A is finite, as it often is in discrete optimization via simulation problems, we could take $A_n = A$, allowing the whole space of feasible solutions. This choice was considered in [24]. Alternatively, one could take $A_n = \{x_1, \dots, x_n\}$, stating that one is willing to consider only those points that have been previously evaluated. This choice is consistent with the expected improvement algorithm. Indeed, we will see that when one makes this choice, and measurements are free from noise, then the knowledge-gradient algorithm is identical to the expected improvement algorithm. Thus, the knowledge-gradient algorithm generalizes the expected improvement algorithm.

If we were to stop sampling at time n , then the expected value of a point $x \in A_n$ based on the information available would be $E_n[f(x)] = \mu_n(x)$. In the special case when evaluations are free from noise, this is equal to $f(x)$, but when there is noise, these two quantities may differ. If we needed to report a final solution, we would then choose the point in A_n for which this quantity is the largest, i.e., we would choose $\operatorname{argmax}_{x \in A_n} \mu_n(x)$. Moreover, the expected value of this solution would be

$$\mu_n^* = \max_{x \in A_n} \mu_n(x).$$

If evaluations are free from noise and $A_n = \{x_{1:n}\}$, then μ_n^* is equal to f_n^* , but in general these quantities may differ.

If we take one additional sample, then the expected value of the solution we would report based on this additional information is

$$\mu_{n+1}^* = \max_{x \in A_{n+1}} \mu_{n+1}(x),$$

where as before, A_{n+1} is some finite set of points we would be willing to consider when choosing a final solution. Observe in this expression that $\mu_{n+1}(x)$ is not necessarily the same as $\mu_n(x)$, even for points $x \in \{x_{1:n}\}$ that we had previously evaluated, but that $\mu_{n+1}(x)$ can be computed from the history of observations $x_{1:n+1}, y_{1:n+1}$.

The improvement in our expected solution value is then the difference between these two quantities, $\mu_{n+1}^* - \mu_n^*$. This improvement is random at time n , even fixing x_{n+1} , through its dependence on y_{n+1} , but we can take its expectation. The resulting quantity is called the *knowledge-gradient (KG) factor*, and is written,

$$\text{KG}_n(x) = E_n [\mu_{n+1}^* - \mu_n^* \mid x_{n+1} = x]. \quad (3.18)$$

Calculating this expectation is more involved than calculating the expected improvement, but nevertheless can also be done analytically in terms of the normal pdf and normal cdf. This is described in more detail in the Derivations and Proofs section.

The knowledge-gradient algorithm is then the one that chooses the point to sample next that maximizes the KG factor,

$$\underset{x}{\operatorname{argmax}} \text{KG}_n(x).$$

The KG factor for a one-dimensional optimization problem with noise is pictured in Fig. 3.6. We see a similar tradeoff between exploration and exploitation, where the KG factor favors measuring points with a large $\mu_n(x)$ and a large $\sigma_n(x)$. We also see local minima of the KG factor at points where we previously evaluated, just as with the expected improvement, but because there is noise in our samples, the value at these points is not 0—indeed, when there is noise, it may be useful to sample repeatedly at a point.

Choice of A_n and A_{n+1}

Recall that the KG factor depends on the choice of the sets A_n and A_{n+1} , through the dependence of μ_n^* and μ_{n+1}^* on these sets. Typically, if we choose these sets to contain more elements, then we allow μ_n^* and μ_{n+1}^* to range over a larger portion of the space, and we allow the KG factor calculation to more accurately approximate the value that would result if we allowed ourself to implement the best option. However, as we increase the size of these sets, computing the KG factor is slower, making implementation of the KG method more computationally intensive.




Fig. 3.6 Upper panel shows the posterior distribution in a problem with independent normal homoscedastic noise and a one-dimensional input space, where the circles are previously measured points, the solid line is the posterior mean $\mu_n(x)$, and the dashed lines are at $\mu_n(x) \pm 2\sigma_n(x)$. Lower panel shows the natural logarithm of the knowledge-gradient factor $\text{KG}(x)$ computed from this posterior distribution, where $A_n = A_{n+1}$ are the discrete grid $\{1, \dots, 300\}$. An “x” is marked at the point with the largest KG factor, which is where the KG algorithm would evaluate next

For applications with a finite A , [24] proposed setting $A_{n+1} = A_n = A$, which was seen to require fewer function evaluations to find points with large f , in comparison with expected improvement on noise-free problems, and in comparison with another Bayesian optimization method, sequential kriging optimization (SKO) [30] on noisy problems. However, the computation and memory required grows rapidly with the size of A , and is typically not feasible when A contains more than 10,000 points.

For large-scale applications, [25] proposed setting $A_{n+1} = A_n = \{x_{1:n+1}\}$ in (3.18), and called the resulting quantity the *approximate knowledge gradient (AKG)*, observing that this choice maintains computational tractability as A grows, but also offers good performance. This algorithm is implemented in the DiceKriging package [31].

Finally, in noise-free problems (but not in problems with noise), setting $A_{n+1} = \{x_{1:n+1}\}$ and $A_n = \{x_{1:n}\}$ recovers expected improvement.

3.4.3 Going Beyond One-Step Analyses, and Other Methods

Both expected improvement and the knowledge-gradient method are designed to be optimal, in the special case where we will take just one more function evaluation and then choose a final solution. They are not, however, known to be optimal for the more general case in which we will take multiple measurements, which is the way they are used in practice.

The optimal algorithm for this more general setting is understood to be the solution to a partially observable Markov decision process, but actually computing the optimal solution using this understanding is intractable using current methods [32]. Some work has been done toward the goal of developing such an optimal algorithm [33], but computing the optimal algorithm remains out of reach. Optimal strategies have been computed for other closely related problems in optimization of expensive noisy functions, including stochastic root-finding [34], multiple comparisons with a standard [35], and small instances of discrete noisy optimization with normally distributed noise (also called “ranking and selection”) [36].

Expected improvement and the knowledge gradient are both special cases of the more general concept of value of information, or expected value of sample information (EVSI) [37], as they calculate the expected reward of a final implementation decision as a function of the posterior distribution resulting from some information, subtract from this the expected reward that would result from not having the information, and then take the expectation of this difference with respect to the information itself.

Many other Bayesian optimization methods have been proposed. A few of these methods optimize the value of information, but are calculated using different assumptions than those used to derive expected improvement or value of information. A larger number of these methods optimize quantities that do not correspond to a value of information, but are derived using analyses that are similar in spirit. These include methods that optimize the probability of improvement [1, 38, 39], the entropy of the posterior distribution on the location of the maximum [40], and other composite measures involving the mean and the standard deviation of the posterior [30].

Other Bayesian optimization methods are designed for problem settings that do not match the assumptions made in this tutorial. These include [41–43], which consider multiple objectives; [6, 44–46], which consider multiple simultaneous function evaluations; [47–49], which consider objective functions that can be evaluated with multiple fidelities and costs; [50], which considers Bernoulli outcomes, rather than normally distributed ones; [51], which considers expensive-to-evaluate inequality constraints; and [52], which considers optimization over the space of small molecules.

3.5 Software

There are a number of excellent software packages, both freely available and commercial, that implement the methods described in this chapter, and other similar methods.

- Metrics Optimization Engine (MOE), an open-source code in C++ and Python, developed by the authors and engineers at Yelp. <http://yelp.github.io/MOE/>,
- Spearmint, an open-source code in Python, implementing algorithms described in [6]. <https://github.com/JasperSnoek/spearmint>
- DiceKriging and DiceOptim, an open-source R package that implements expected improvement, the approximate knowledge-gradient method, and a variety of algorithms for parallel evaluations. An overview is provided in [31]. <http://cran.r-project.org/web/packages/DiceOptim/index.html>,
- TOMLAB, a commercial package for MATLAB. <http://tomopt.com/tomlab/>
- matlabKG, an open-source research code that implements the discrete knowledge-gradient method for small-scale problems. <http://people.orie.cornell.edu/pfrazier/src.html>

A list of software packages focused on Gaussian process regression (but not Bayesian optimization) may be found at <http://www.gaussianprocess.org/>.

3.6 Conclusion

We have presented Bayesian optimization, including Gaussian process regression, the expected improvement method, and the knowledge-gradient method. In making this presentation, we wish to emphasize that this approach to materials design acknowledges the inherent uncertainty in statistical prediction and seeks to guide experimentation in a way that is robust to this uncertainty. It is inherently iterative, and does not seek to circumvent the fundamental trial-and-error process.

This is in contrast with another approach to informatics in materials design, which holds the hope that predictive methods can short-circuit the iterative loop entirely. In this alternative view of the world, one hopes to create extremely accurate prediction techniques, either through physically-motivated *ab initio* calculations, or using data-driven machine learning approaches, that are so accurate that one can rely on the predictions alone rather than on physical experiments. If this can be achieved, then we can search over materials designs *in silico*, find those designs that are predicted to perform best, and test those designs alone in physical experiments.

For this approach to be successful, one must have extremely accurate predictions, which limits its applicability to settings where this is possible. We argue that, in contrast, predictive techniques can be extremely powerful even if they are not perfectly accurate, as long as they are used in a way that acknowledges inaccuracy, builds in robustness, and reduces this inaccuracy through an iterative dialog with physical

reality mediated by physical experiments. Moreover, we argue that mathematical techniques like Bayesian optimization, Bayesian experimental design, and optimal learning provide us the mathematical framework for accomplishing this goal in a principled manner, and for using our power to predict as effectively as possible.

Acknowledgments Peter I. Frazier was supported by AFOSR FA9550-12-1-0200, AFOSR FA9550-15-1-0038, NSF CAREER CMMI-1254298, NSF IIS-1247696, and the ACSF's AVF. Jialei Wang was supported by AFOSR FA9550-12-1-0200.

Derivations and Proofs

This section contains derivations and proofs of equations and theoretical results found in the main text.

Proof of Proposition 1

Proof Using Bayes' rule, the conditional probability density of $\theta_{[2]}$ at a point $u_{[2]}$ given that $\theta_{[1]} = u_{[1]}$ is

$$p(\theta_{[2]} = u_{[2]} \mid \theta_{[1]} = u_{[1]}) = \frac{p(\theta_{[1]} = u_{[1]}, \theta_{[2]} = u_{[2]})}{p(\theta_{[1]} = u_{[1]})} \propto p(\theta_{[1]} = u_{[1]}, \theta_{[2]} = u_{[2]}) \propto \exp\left(-\frac{1}{2} \begin{bmatrix} u_{[1]} - \mu_{[1]} \\ u_{[2]} - \mu_{[2]} \end{bmatrix}^T \begin{bmatrix} \Sigma_{[1,1]} & \Sigma_{[1,2]} \\ \Sigma_{[2,1]} & \Sigma_{[2,2]} \end{bmatrix}^{-1} \begin{bmatrix} u_{[1]} - \mu_{[1]} \\ u_{[2]} - \mu_{[2]} \end{bmatrix}\right). \quad (3.19)$$

To deal with the inverse matrix in this expression, we use the following identity for inverting a block matrix: the inverse of the block matrix $\begin{bmatrix} A & B \\ C & D \end{bmatrix}$, where both A and D are invertible square matrices, is

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix}^{-1} = \begin{bmatrix} (A - BD^{-1}C)^{-1} & -(A - BD^{-1}C)^{-1}BD^{-1} \\ -(D - CA^{-1}B)^{-1}CA^{-1} & (D - CA^{-1}B)^{-1} \end{bmatrix}. \quad (3.20)$$

Applying (3.20) to (3.19), and using a bit of algebraic manipulation to get rid of constants, we have

$$p(\theta_{[2]} = u_{[2]} \mid \theta_{[1]} = u_{[1]}) \propto \exp\left(-\frac{1}{2}(u_{[2]} - \mu^{\text{new}})^T (\Sigma^{\text{new}})^{-1} (u_{[2]} - \mu^{\text{new}})\right), \quad (3.21)$$

where $\mu^{\text{new}} = \mu_{[2]} - \Sigma_{[2,1]} \Sigma_{[1,1]}^{-1} (u_{[1]} - \mu_{[1]})$ and $\Sigma^{\text{new}} = \Sigma_{[2,2]} - \Sigma_{[2,1]} \Sigma_{[1,1]}^{-1} \Sigma_{[1,2]}$.

We see that (3.21) is simply the unnormalized probability density function of a normal distribution. Thus the conditional distribution of $\theta_{[2]}$ given $\theta_{[1]} = u_{[1]}$ is multivariate normal, with mean μ^{new} and covariance matrix Σ^{new} .

Derivation of Equation (3.16)

Since $f(x) \sim \text{Normal}(\mu_n(x), \sigma_n^2(x))$, the probability density of $f(x)$ is $p(f(x) = z) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(z - \mu_n(x))^2}{2\sigma_n^2(x)}\right)$. We use this to calculate $\text{EI}(x)$:

$$\begin{aligned}
\text{EI}(x) &= E_n[(f(x) - f_n^*)^+] \\
&= \int_{f_n^*}^{\infty} (z - f_n^*) \frac{1}{\sqrt{2\pi}\sigma_n(x)} e^{-\frac{(z - \mu_n(x))^2}{2\sigma_n^2(x)}} dz \\
&= \int_{f_n^*}^{\infty} z \frac{1}{\sqrt{2\pi}\sigma_n(x)} e^{-\frac{(z - \mu_n(x))^2}{2\sigma_n^2(x)}} dz - f_n^* \left(1 - \Phi\left(\frac{f_n^* - \mu_n(x)}{\sigma_n(x)}\right)\right) \\
&= \int_{f_n^*}^{\infty} (\mu_n(x) + (z - \mu_n(x))) \frac{1}{\sqrt{2\pi}\sigma_n(x)} e^{-\frac{(z - \mu_n(x))^2}{2\sigma_n^2(x)}} dz - f_n^* \left(1 - \Phi\left(\frac{f_n^* - \mu_n(x)}{\sigma_n(x)}\right)\right) \\
&= \int_{f_n^*}^{\infty} (z - \mu_n(x)) \frac{1}{\sqrt{2\pi}\sigma_n(x)} e^{-\frac{(z - \mu_n(x))^2}{2\sigma_n^2(x)}} dz + (\mu_n(x) - f_n^*) \left(1 - \Phi\left(\frac{f_n^* - \mu_n(x)}{\sigma_n(x)}\right)\right) \\
&= \sigma_n(x) \frac{1}{\sqrt{2\pi}} e^{-\frac{(f_n^* - \mu_n(x))^2}{2\sigma_n^2(x)}} + (\mu_n(x) - f_n^*) \left(1 - \Phi\left(\frac{f_n^* - \mu_n(x)}{\sigma_n(x)}\right)\right) \\
&= (\mu_n(x) - f_n^*) \left(1 - \Phi\left(\frac{f_n^* - \mu_n(x)}{\sigma_n(x)}\right)\right) + \sigma_n(x) \varphi\left(\frac{f_n^* - \mu_n(x)}{\sigma_n(x)}\right) \\
&= (\mu_n(x) - f_n^*) \Phi\left(\frac{\mu_n(x) - f_n^*}{\sigma_n(x)}\right) + \sigma_n(x) \varphi\left(\frac{\mu_n(x) - f_n^*}{\sigma_n(x)}\right).
\end{aligned}$$

Calculation of the KG factor

The KG factor (3.18) is calculated by first considering how the quantity $\mu_{n+1}^* - \mu_n^*$ depends on the information that we have at time n , and the additional datapoint that we will obtain, y_{n+1} .

First observe that $\mu_{n+1}^* - \mu_n^*$ is a deterministic function of the vector $[\mu_{n+1}(x) : x \in A_{n+1}]$ and other quantities that are known at time n . Then, by applying the analysis in Sect. 3.3.5, but letting the posterior given $x_{1:n}, y_{1:n}$ play the role of the prior, we obtain the following version of (3.10), which applies to any given x ,

$$\mu_{n+1}(x) = \mu_n(x) + \frac{\Sigma_n(x, x_{n+1})}{\Sigma_n(x_{n+1}, x_{n+1}) + \lambda^2} (y_{n+1} - \mu_n(x_{n+1})). \quad (3.22)$$

In this expression, $\mu_n(\cdot)$ and $\Sigma_n(\cdot, \cdot)$ are given by (3.13) and (3.14).

We see from this expression that $\mu_{n+1}(x)$ is a linear function of y_{n+1} , with an intercept and a slope that can be computed based on what we know after the n th measurement.

We will calculate the distribution of y_{n+1} , given what we have observed at time n . First, $f(x_{n+1})|x_{1:n}, y_{1:n} \sim \text{Normal}(\mu_n(x_{n+1}), \Sigma_n(x_{n+1}, x_{n+1}))$. Since $y_{n+1} = f(x_{n+1}) + \varepsilon_{n+1}$, where ε_{n+1} is independent with distribution $\varepsilon_{n+1} \sim \text{Normal}(0, \lambda^2)$, we have

$$y_{n+1}|x_{1:n}, y_{1:n} \sim \text{Normal}(\mu_n(x_{n+1}), \Sigma_n(x_{n+1}, x_{n+1}) + \lambda^2).$$

Plugging the distribution of y_{n+1} into (3.22) and doing some algebra, we have

$$\mu_{n+1}(x)|x_{1:n}, y_{1:n} \sim \text{Normal}(\mu_n(x), \tilde{\sigma}^2(x, x_{n+1})),$$

where $\tilde{\sigma}(x, x_{n+1}) = \frac{\Sigma_n(x, x_{n+1})}{\sqrt{\Sigma_n(x_{n+1}, x_{n+1}) + \lambda^2}}$. Moreover, we can write $\mu_{n+1}(x)$ as

$$\mu_{n+1}(x) = \mu_n(x) + \tilde{\sigma}(x, x_{n+1})Z,$$

where $Z = (y_{n+1} - \mu_n(x_{n+1}))/\sqrt{\Sigma_n(x_{n+1}, x_{n+1}) + \lambda^2}$ is a standard normal random variable, given $x_{1:n}$ and $y_{1:n}$.

Now (3.18) becomes

$$\text{KG}_n(x) = E_n \left[\max_{x' \in A_{n+1}} \mu_n(x') + \tilde{\sigma}(x', x_{n+1})Z \mid x_{n+1} = x \right] - \mu_n^*.$$

Thus, computing the KG factor comes down to being able to compute the expectation of the maximum of a collection of linear functions of a scalar normal random variable. Algorithm 2 of [24], with software provided as part of the matlabKG library [53], computes the quantity

$$h(a, b) = \mathbb{E} \left[\max_{i=1, \dots, |a|} (a_i + b_i Z) \right] - \max_{i=1, \dots, |a|} a_i$$

for arbitrary equal-length vectors a and b . Using this ability, and letting $\mu_n(A_{n+1})$ be the vector $[\mu_n(x') : x' \in A_{n+1}]$ and $\tilde{\sigma}(A_{n+1}, x)$ be the vector $[\tilde{\sigma}(x', x) : x' \in A_{n+1}]$, we can write the KG factor as

$$\text{KG}_n(x) = h(\mu_n(A_{n+1}), \tilde{\sigma}(A_{n+1}, x)) + [\max(\mu_n(A_{n+1})) - \mu_n^*].$$

If $A_{n+1} = A_n$, as it is in the versions of the knowledge-gradient method proposed in [24, 25], then the last term $\max(\mu_n(A_{n+1})) - \mu_n^*$ is equal to 0 and vanishes.

References

1. H.J. Kushner, A new method of locating the maximum of an arbitrary multi-peak curve in the presence of noise. *J. Basic Eng.* **86**, 97–106 (1964)
2. J. Mockus, *Bayesian Approach to Global Optimization: Theory and Applications* (Kluwer Academic, Dordrecht, 1989)
3. J. Mockus, V. Tiesis, A. Zilinskas, The application of Bayesian methods for seeking the extremum, in *Towards Global Optimisation*, ed. by L.C.W. Dixon, G.P. Szego, vol. 2 (Elsevier Science Ltd., North Holland, Amsterdam, 1978), pp. 117–129
4. D.R. Jones, M. Schonlau, W.J. Welch, Efficient Global Optimization of Expensive Black-Box Functions. *J. Global Optim.* **13**(4), 455–492 (1998)
5. A. Booker, J. Dennis, P. Frank, D. Serafini, V. Torczon, M.W. Trosset, Optimization using surrogate objectives on a helicopter test example. *Prog. Syst. Control Theor.* **24**, 49–58 (1998)
6. J. Snoek, H. Larochelle, R.P. Adams, Practical bayesian optimization of machine learning algorithms. in *Advances in Neural Information Processing Systems*, pp. 2951–2959 (2012)
7. E. Brochu, M. Cora, N. de Freitas, *A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning*. Technical Report TR-2009-023, Department of Computer Science, University of British Columbia, November 2009
8. A. Forrester, A. Sobester, A. Keane, *Engineering Design Via Surrogate Modelling: A Practical Guide* (Wiley, West Sussex, UK, 2008)
9. T.J. Santner, B.W. Williams, W. Notz, *The Design and Analysis of Computer Experiments* (Springer, New York, 2003)
10. M.J. Sasena, *Flexibility and Efficiency Enhancements for Constrained Global Design Optimization with Kriging Approximations*. Ph.D. thesis, University of Michigan (2002)
11. D.G. Kibib, A statistical approach to some basic mine valuation problems on the witwatersrand. *J. Chem. Metall. Min. Soc. S. Afr.* (1951)
12. G. Matheron, *The theory of regionalized variables and its applications*, vol 5. École national supérieure des mines (1971)
13. N. Cressie, The origins of kriging. *Math. Geol.* **22**(3), 239–252 (1990)
14. C.E. Rasmussen, C.K.I. Williams, *Gaussian Processes for Machine Learning* (MIT Press, Cambridge, MA, 2006)
15. C.E. Rasmussen (2011), <http://www.gaussianprocess.org/code>, Accessed 15 July 2015
16. A.B. Gelman, J.B. Carlin, H.S. Stern, D.B. Rubin, *Bayesian Data Analysis* (CRC Press, Boca Raton, FL, second edition, 2004)
17. J.O. Berger, *Statistical decision theory and Bayesian analysis* (Springer-Verlag, New York, second edition) (1985)
18. B. Ankenman, B.L. Nelson, J. Staum, Stochastic kriging for simulation metamodeling. *Oper. Res.* **58**(2), 371–382 (2010)
19. P.W. Goldberg, C.K.I. Williams, C.M. Bishop, Regression with input-dependent noise: a gaussian process treatment. *Advances in neural information processing systems*, p. 493–499 (1998)
20. K. Kersting, C. Plagemann, P. Pfaff, W. Burgard, Most likely heteroscedastic Gaussian process regression. In *Proceedings of the 24th international conference on Machine learning*, ACM, pp. 393–400 (2007)
21. C. Wang, *Gaussian Process Regression with Heteroscedastic Residuals and Fast MCMC Methods*. Ph.D. thesis, University of Toronto (2014)
22. P.I. Frazier, J. Xie, S.E. Chick, Value of information methods for pairwise sampling with correlations, in *Proceedings of the 2011 Winter Simulation Conference*, ed. by S. Jain, R.R. Creasey, J. Himmelspach, K.P. White, M. Fu (Institute of Electrical and Electronics Engineers Inc, Piscataway, New Jersey, 2011), pp. 3979–3991
23. S. Sankaran, A.L. Marsden, The impact of uncertainty on shape optimization of idealized bypass graft models in unsteady flow. *Physics of Fluids (1994-present)*, **22**(12):121–902 (2010)

24. P.I. Frazier, W.B. Powell, S. Dayanik, The knowledge gradient policy for correlated normal beliefs. *INFORMS J. Comput.* **21**(4), 599–613 (2009)
25. W. Scott, P.I. Frazier, W.B. Powell, The correlated knowledge gradient for simulation optimization of continuous parameters using gaussian process regression. *SIAM J. Optim.* **21**(3), 996–1026 (2011)
26. L.P. Kaelbling, *Learning in Embedded Systems* (MIT Press, Cambridge, MA, 1993)
27. R.S. Sutton, A.G. Barto, *Reinforcement Learning* (The MIT Press, Cambridge, Massachusetts, 1998)
28. J. Gittins, K. Glazebrook, R. Weber. *Multi-armed Bandit Allocation Indices*. Wiley, 2nd edition (2011)
29. A. Mahajan, D. Teneketzis, Multi-armed bandit problems. In D. Cochran A. O. Hero III, D. A. Castanon, K. Kastella, (Ed.). *Foundations and Applications of Sensor Management*. Springer-Verlag (2007)
30. D. Huang, T.T. Allen, W.I. Notz, N. Zeng, Global Optimization of Stochastic Black-Box Systems via Sequential Kriging Meta-Models. *J. Global Optim.* **34**(3), 441–466 (2006)
31. O. Roustant, D. Ginsbourger, Y. Deville, Dicekriging, diceoptim: two R packages for the analysis of computer experiments by kriging-based metamodeling and optimization. *J. Stat. Softw.* **51**(1), p. 54 (2012)
32. P.I. Frazier, *Learning with Dynamic Programming*. John Wiley and Sons (2011)
33. D. Ginsbourger, R. Riche, Towards gaussian process-based optimization with finite time horizon. *mOda 9—Advances in Model-Oriented Design and Analysis*, p. 89–96 (2010)
34. R. Waeber, P.I. Frazier, S.G. Henderson, Bisection search with noisy responses. *SIAM J. Control Optim.* **51**(3), 2261–2279 (2013)
35. J. Xie, P.I. Frazier, Sequential bayes-optimal policies for multiple comparisons with a known standard. *Oper. Res.* **61**(5), 1174–1189 (2013)
36. P.I. Frazier, Tutorial: Optimization via simulation with bayesian statistics and dynamic programming, in *Proceedings of the 2012 Winter Simulation Conference Proceedings*, ed. by C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, A.M. Uhrmacher (Institute of Electrical and Electronics Engineers Inc., Piscataway, New Jersey, 2012), pp. 79–94
37. R.A. Howard, Information Value Theory. *Syst. Sci. Cybern. IEEE Trans.* **2**(1), 22–26 (1966)
38. C.D. Perttunen, A computational geometric approach to feasible region division unconstrained global optimization. In *Proceedings of 1991 IEEE International Conference on Systems, Man, and Cybernetics, 1991. 'Decision Aiding for Complex Systems*, pp. 585–590 (1991)
39. B.E. Stuckman, A global search method for optimizing nonlinear systems. *Syst. Man Cybern. IEEE Trans.* **18**(6), 965–977 (1988)
40. J. Villemonteix, E. Vazquez, E. Walter, An informational approach to the global optimization of expensive-to-evaluate functions. *J. Global Optim.* **44**(4), 509–534 (2009)
41. D.C.T. Bautista, *A Sequential Design for Approximating the Pareto Front using the Expected Pareto Improvement Function*. Ph.D. thesis, The Ohio State University (2009)
42. P.I. Frazier, A.M. Kazachkov, Guessing preferences: a new approach to multi-attribute ranking and selection, in *Proceedings of the 2011 Winter Simulation Conference*, ed. by S. Jain, R.R. Creasey, J. Himmelspach, K.P. White, M. Fu (Institute of Electrical and Electronics Engineers Inc, Piscataway, New Jersey, 2011), pp. 4324–4336
43. J. Knowles, ParEGO: A hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *Evol. Comput. IEEE Trans.* **10**(1), 50–66 (2006)
44. S.C. Clark, J. Wang, E. Liu, P.I. Frazier, Parallel global optimization using an improved multi-points expected improvement criterion (working paper, 2014)
45. D. Ginsbourger, R. Le Riche, L. Carraro, A multi-points criterion for deterministic parallel global optimization based on kriging. In *International Conference on Nonconvex Programming, NCP07*, Rouen, France, December 2007
46. D. Ginsbourger, R. Le Riche, and L. Carraro, Kriging is well-suited to parallelize optimization. In *Computational Intelligence in Expensive Optimization Problems*, Springer, vol. 2, p. 131–162 (2010)

47. A.I.J. Forrester, A. Sóbester, A.J. Keane, Multi-fidelity optimization via surrogate modelling. *Proc. R. Soc. A: Math. Phys. Eng. Sci.* **463**(2088), 3251–3269 (2007)
48. P.I. Frazier, W.B. Powell, H.P. Simão, Simulation model calibration with correlated knowledge-gradients, in *Proceedings of the 2009 Winter Simulation Conference Proceedings*, ed. by M.D. Rossetti, R.R. Hill, B. Johansson, A. Dunkin, R.G. Ingalls (Institute of Electrical and Electronics Engineers Inc, Piscataway, New Jersey, 2009), pp. 339–351
49. D. Huang, T.T. Allen, W.I. Notz, R.A. Miller, Sequential kriging optimization using multiple-fidelity evaluations. *Struct. Multi. Optim.* **32**(5), 369–382 (2006)
50. J. Bect, D. Ginsbourger, L. Li, V. Picheny, E. Vazquez, Sequential design of computer experiments for the estimation of a probability of failure. *Stat. Comput.* **22**(3), 773–793 (2012)
51. J.R. Gardner, M.J. Kusner, Z. Xu, K. Weinberger, J.P. Cunningham, Bayesian optimization with inequality constraints. In *Proceedings of The 31st International Conference on Machine Learning*, pp. 937–945 (2014)
52. D.M. Negoescu, P.I. Frazier, W.B. Powell, The knowledge gradient algorithm for sequencing experiments in drug discovery. *INFORMS J. Comput.* **23**(1) (2011)
53. P.I. Frazier (2009–2010), <http://people.orie.cornell.edu/pfrazier/src.html>

Chapter 4

Small-Sample Classification

Lori A. Dalton and Edward R. Dougherty

Abstract In a number of application areas, such as materials and genomics, where one wishes to classify objects, sample sizes are often small owing to the expense or unavailability of data points. Many classifier design procedures work well with large samples but are ineffectual or, at best, problematic with small samples. Worse yet, small-samples make it difficult to impossible to guarantee an accurate error estimate without modeling assumptions, and absent a good error estimate a classifier is useless. The present chapter discusses the problem of small-sample error estimation and how modeling assumptions can be used to obtain bounds on error estimation accuracy. Given the necessity of modeling assumptions, we go on to discuss minimum-mean-square-error (MMSE) error estimation and the design of optimal classifiers relative to prior knowledge and data in a Bayesian context.

4.1 Introduction

Given several classes of objects, one of the most basic problems of engineering and statistics is making a decision as to which class an object belongs to based on some set of features. The standard approach to the problem is to utilize labeled training data sampled from the class populations as inputs to a design algorithm that yields a decision function, known as a classifier. The designed classifier is then used to make decisions regarding future unlabeled observations. Classifier design alone is insufficient: one must also use sample data to estimate the error of the classifier on the class populations. A classifier whose misclassification rate is not known to some satisfactory degree of approximation is useless.

L.A. Dalton (✉)
The Ohio State University, Columbus, OH 43210, USA
e-mail: dalton@ece.osu.edu

E.R. Dougherty
Texas A&M University, College Station, TX 77843, USA
e-mail: edward@ece.tamu.edu

In application areas where data are plentiful and cheap, one can obtain a large training sample to design a classifier and a large independent test sample on which to estimate the error by the proportion of errors on the test sample. When data are limited, not only does this impact classifier design, but it forces one to use the same data for training and testing, else one would have little hope of obtaining a good classifier. With a small training set, one might still hope to design a well-performing classifier and let the estimated error decide if it is actually good. Unfortunately, error estimation is problematic with small samples; indeed, this is the most fundamental problem with small samples, which are ubiquitous in certain application areas, for instance genomics and materials where sample sizes less than 100 are commonplace. This chapter considers small-sample classification, demonstrating the issues with purely data-driven methods, and how these can be addressed using Bayesian approaches. For simplicity we restrict our attention to binary classification, where there are two classes.

4.2 Classification

Classification involves a *feature vector* $\mathbf{X} = (X_1, X_2, \dots, X_D)$ on D -dimensional Euclidean space \mathfrak{R}^D composed of random variables (*features*), a binary random variable $Y \in \{0, 1\}$ (0 and 1 are called *labels*), and a *classifier* $\psi : \mathfrak{R}^D \rightarrow \{0, 1\}$ to predict Y by $\psi(\mathbf{X})$. Classification is probabilistically characterized via the joint *feature-label distribution* F for the pair (\mathbf{X}, Y) . The space of all classifiers, which consists of the space of all binary functions on \mathfrak{R}^D , will be denoted by \mathcal{F} . The error $\varepsilon[\psi]$ of $\psi \in \mathcal{F}$ is the probability of misclassification,

$$\varepsilon[\psi] = P(\psi(\mathbf{X}) \neq Y) = E[|Y - \psi(\mathbf{X})|], \quad (4.1)$$

the probability and expectation being taken relative to F . An optimal classifier, ψ_{Bayes} , is one having minimal error, $\varepsilon_{\text{Bayes}}$, among all $\psi \in \mathcal{F}$. ψ_{Bayes} and $\varepsilon_{\text{Bayes}}$ are called a *Bayes classifier* and the *Bayes error*, respectively. A Bayes classifier, which need not be unique, and the Bayes error, depend on F .

Define $\eta_0(\mathbf{x}) = f_{\mathbf{X},Y}(\mathbf{x}, 0)/f_{\mathbf{X}}(\mathbf{x})$ and $\eta_1(\mathbf{x}) = f_{\mathbf{X},Y}(\mathbf{x}, 1)/f_{\mathbf{X}}(\mathbf{x})$, where $f_{\mathbf{X},Y}(\mathbf{x}, y)$ and $f_{\mathbf{X}}(\mathbf{x})$ are the densities for (\mathbf{X}, Y) and \mathbf{X} , respectively. The posteriors $\eta_0(\mathbf{x})$ and $\eta_1(\mathbf{x})$ give the probability that $Y = 0$ and $Y = 1$, respectively, given $\mathbf{X} = \mathbf{x}$. Classifier error can be expressed as

$$\varepsilon[\psi] = \int_{\{\mathbf{x}|\psi(\mathbf{x})=0\}} \eta_1(\mathbf{x}) f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x} + \int_{\{\mathbf{x}|\psi(\mathbf{x})=1\}} \eta_0(\mathbf{x}) f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x}. \quad (4.2)$$

The right-hand side of (4.2) is minimized by

$$\psi_{\text{Bayes}}(\mathbf{x}) = \begin{cases} 1, & \text{if } \eta_1(\mathbf{x}) \geq \eta_0(\mathbf{x}) \\ 0, & \text{otherwise} \end{cases}. \quad (4.3)$$

It follows from (4.2) and (4.3) that the Bayes error is given by

$$\begin{aligned} \varepsilon_{\text{Bayes}} &= \int_{\{\mathbf{x}|\eta_1(\mathbf{x}) < \eta_0(\mathbf{x})\}} \eta_1(\mathbf{x}) f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x} + \int_{\{\mathbf{x}|\eta_1(\mathbf{x}) \geq \eta_0(\mathbf{x})\}} \eta_0(\mathbf{x}) f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x} \\ &= E[\min\{\eta_0(\mathbf{X}), \eta_1(\mathbf{X})\}]. \end{aligned} \quad (4.4)$$

By Jensen's inequality,

$$\varepsilon_{\text{Bayes}} \leq \min\{E[\eta_0(\mathbf{X})], E[\eta_1(\mathbf{X})]\} = \min\{P(Y = 0), P(Y = 1)\},$$

where $P(Y = y)$ is the prior probability that a sample point is from class y . Thus, if either prior is small, then the Bayes error is necessarily small. This occurs if one class is much more likely than the other.

Each class, $y \in \{0, 1\}$, is described by its *class-conditional distribution* $f_{\mathbf{X}|Y}(\mathbf{x}|y)$. In the Gaussian model, each sample point in a given class is a column vector of D multivariate Gaussian features. In particular, the class-conditional distribution for class y is Gaussian with mean μ_y and covariance matrix Σ_y . Letting $c = P(Y = 0)$, the optimal classifier is quadratic and given by

$$\psi_{\text{Bayes}}(\mathbf{x}) = \begin{cases} 1, & \text{if } g_{\text{Bayes}}(\mathbf{x}) > 0 \\ 0, & \text{if } g_{\text{Bayes}}(\mathbf{x}) \leq 0 \end{cases}, \quad (4.5)$$

where

$$g_{\text{Bayes}}(\mathbf{x}) = \mathbf{x}^T A_{\text{Bayes}} \mathbf{x} + \mathbf{a}_{\text{Bayes}}^T \mathbf{x} + b_{\text{Bayes}}, \quad (4.6)$$

with constant matrix A_{Bayes} , column vector $\mathbf{a}_{\text{Bayes}}$ and scalar b_{Bayes} given by

$$\begin{aligned} A_{\text{Bayes}} &= -\frac{1}{2} (\Sigma_1^{-1} - \Sigma_0^{-1}), \\ \mathbf{a}_{\text{Bayes}} &= \Sigma_1^{-1} \mu_1 - \Sigma_0^{-1} \mu_0, \\ b_{\text{Bayes}} &= -\frac{1}{2} (\mu_1^T \Sigma_1^{-1} \mu_1 - \mu_0^T \Sigma_0^{-1} \mu_0) + \ln \left(\frac{1-c}{c} \left(\frac{|\Sigma_0|}{|\Sigma_1|} \right)^{1/2} \right). \end{aligned} \quad (4.7)$$

When $\Sigma = \Sigma_0 = \Sigma_1$, this classifier is linear and defined by

$$g_{\text{Bayes}}(\mathbf{x}) = \mathbf{a}_{\text{Bayes}}^T \mathbf{x} + b_{\text{Bayes}}, \quad (4.8)$$

where

$$\begin{aligned} \mathbf{a}_{\text{Bayes}} &= \Sigma^{-1} (\mu_1 - \mu_0), \\ b_{\text{Bayes}} &= -\frac{1}{2} \mathbf{a}_{\text{Bayes}}^T (\mu_1 + \mu_0) + \ln \frac{1-c}{c}. \end{aligned} \quad (4.9)$$

In practice, the feature-label distribution is unknown and a classifier is designed from sample data. A common assumption, and one we make here, is that a classifier ψ_n is designed using a random sample $S_n = \{(\mathbf{X}_1, Y_1), (\mathbf{X}_2, Y_2), \dots, (\mathbf{X}_n, Y_n)\}$ of vector-label pairs drawn from the feature-label distribution. While random sampling is usually assumed, in some applications sampling is often not random and this leads to misapplication of classification theory developed in the framework of random sampling. For instance, with *separate sampling* data are drawn randomly from each class but the number from each class is set outside the sampling procedure. Separate sampling is common in biomedical areas such as genomics and this can lead to serious problems for both classifier design [1, 2] and error estimation [3] if not taken into account.

Classifier design requires a procedure that operates on a sample to yield a classifier. A *classification rule* is a mapping $\Psi_n : [\mathfrak{R}^D \times \{0, 1\}]^n \rightarrow \mathcal{F}$. Given a sample S_n , Ψ_n yields a designed classifier $\psi_n = \Psi_n(S_n) \in \mathcal{F}$. To be fully formal, one might write $\psi_n(S_n; \mathbf{X})$ rather than $\psi_n(\mathbf{X})$; however, we will use the simpler notation, keeping in mind that ψ_n derives from a classification rule applied to a feature-label sample. Note that a classification rule is really a sequence of classification rules, each depending on the sample size, n .

Under a Gaussian assumption in which the means and covariances are unknown, a natural classification rule when the covariance matrices are unequal is to replace the means and covariances in (4.7) with the sample means and covariances computed from sample data and to replace c by the estimate $\hat{c} = n_0/n$, where n_y is the number of sample points in class y . This yields the *quadratic discriminant analysis* (QDA) classification rule. Assuming equal covariance matrices, the *linear discriminant analysis* (LDA) classification rule results from replacing the means and covariance in (4.9) by the sample means and a pooled sample covariance matrix, and replacing c by $\hat{c} = n_0/n$. Although constructed under the Gaussian assumption, QDA and LDA can be used without a Gaussian assumption and may perform fairly well so long as the true class-conditional distributions are not too far from Gaussian, and the sample is sufficiently large that the sample estimates are accurate.

Since the optimal error is the Bayes error, sample-based design suffers a *design cost*, $\Delta_n = \varepsilon_n - \varepsilon_{\text{Bayes}}$, where $\varepsilon_n = \varepsilon[\psi_n]$ and ε_n and Δ_n are sample-dependent random variables. The expected design cost is $E[\Delta_n]$, the expectation here being relative to the random sample drawn from F . The expected error of ψ_n is decomposed according to $E[\varepsilon_n] = \varepsilon_{\text{Bayes}} + E[\Delta_n]$. A classification rule is said to be *consistent* for a feature-label distribution F if $\Delta_n \rightarrow 0$ in the mean, meaning $E[\Delta_n] \rightarrow 0$ as $n \rightarrow \infty$. For a consistent rule, the expected design cost can be made arbitrarily small

for a sufficiently large amount of data. A classification rule is *universally consistent* if $\Delta_n \rightarrow 0$ in the mean for any feature-label distribution of (\mathbf{X}, Y) . Consistency is useful for large samples, but has negligible value for small samples.

A classification rule can yield a classifier that makes few errors, or even no errors, on the training data but performs poorly on the distribution as a whole, a situation called *overfitting*. This situation is exacerbated by the use of complex classifiers with small samples. The essential idea is that a classifier should not cut up the space too finely for the amount of training data. Overfitting can be mitigated by constraining classifier design, which means restricting classifiers to a subfamily $\mathcal{C} \subseteq \mathcal{F}$. The aim is to find an optimal *constrained classifier* $\psi_{\mathcal{C}} \in \mathcal{C}$ having error $\varepsilon_{\mathcal{C}}$. Since optimization in \mathcal{C} is over a subfamily of classifiers, $\varepsilon_{\mathcal{C}} \geq \varepsilon_{\text{Bayes}}$. The cost of constraint is $\Delta_{\mathcal{C}} = \varepsilon_{\mathcal{C}} - \varepsilon_{\text{Bayes}} \geq 0$. When only data is available, a classification rule yields a classifier $\psi_{n,\mathcal{C}} \in \mathcal{C}$, with error $\varepsilon_{n,\mathcal{C}}$ such that $\varepsilon_{n,\mathcal{C}} \geq \varepsilon_{\mathcal{C}} \geq \varepsilon_{\text{Bayes}}$. The design cost for constrained classification is $\Delta_{n,\mathcal{C}} = \varepsilon_{n,\mathcal{C}} - \varepsilon_{\mathcal{C}}$. For small samples, this can be substantially less than Δ_n , depending on \mathcal{C} and the classification rule. For instance, although LDA is constructed under the assumption of equal covariance matrices, with small samples it can outperform QDA when the covariance matrices are unequal because it only requires estimation of a single covariance matrix rather than two. The error of a designed constrained classifier is decomposed as $\varepsilon_{n,\mathcal{C}} = \varepsilon_{\text{Bayes}} + \Delta_{\mathcal{C}} + \Delta_{n,\mathcal{C}}$. Hence, the expected error of a constrained designed classifier can be decomposed as

$$E[\varepsilon_{n,\mathcal{C}}] = \varepsilon_{\text{Bayes}} + \Delta_{\mathcal{C}} + E[\Delta_{n,\mathcal{C}}]. \quad (4.10)$$

The constraint is beneficial if and only if $E[\varepsilon_{n,\mathcal{C}}] < E[\varepsilon_n]$, that is, if $\Delta_{\mathcal{C}} < E[\Delta_n] - E[\Delta_{n,\mathcal{C}}]$. If the cost of constraint is less than the decrease in expected design error, then $E[\varepsilon_{n,\mathcal{C}}] < E[\varepsilon_n]$. The dilemma is that strong constraint reduces $E[\Delta_{n,\mathcal{C}}]$ at the cost of increasing $\varepsilon_{\mathcal{C}}$.

A fundamental theorem provides bounds for $E[\Delta_{n,\mathcal{C}}]$ [4]. The idea of choosing a classifier in \mathcal{C} that minimizes the number of errors on the sample data is known as *empirical risk minimization*. A distribution-free bound on the design error for any classification rule that employs empirical risk minimization is given by

$$E[\Delta_{n,\mathcal{C}}] \leq 8\sqrt{\frac{V_{\mathcal{C}} \log n + 4}{2n}}, \quad (4.11)$$

where $V_{\mathcal{C}}$ is a constant known as the *VC (Vapnik-Chervonenkis) dimension* of \mathcal{C} (see [5] for a detailed discussion of the VC dimension). It is obvious that n must greatly exceed $V_{\mathcal{C}}$ for the bound to be small.

4.3 Error Estimation

With the feature-label distribution unknown, the classifier error must be estimated by an *estimation rule*, \mathcal{E}_n , which given the random sample S_n yields an error estimate $\hat{\varepsilon}[\psi_n] = \mathcal{E}_n(S_n)$. The key issue is accuracy. Given a feature-label distribution, error estimation accuracy is commonly measured by the *mean-square error (MSE)*, $\text{MSE}(\hat{\varepsilon}) = E[(\hat{\varepsilon} - \varepsilon)^2]$, where for notational ease we denote $\varepsilon[\psi_n]$ and $\hat{\varepsilon}[\psi_n]$ by ε and $\hat{\varepsilon}$, respectively. The square root of the MSE is known as the *root-mean-square (RMS)*. The expectation is relative to the sampling distribution. The MSE is decomposed into the bias, $\text{Bias}(\hat{\varepsilon}) = E[\hat{\varepsilon} - \varepsilon]$, of the error estimator relative to the true error, and the deviation variance, $\text{Var}_{\text{dev}}(\hat{\varepsilon}) = \text{Var}(\hat{\varepsilon} - \varepsilon)$, according to

$$\text{MSE}(\hat{\varepsilon}) = \text{Var}_{\text{dev}}(\hat{\varepsilon}) + \text{Bias}(\hat{\varepsilon})^2. \quad (4.12)$$

When a large amount of data is available, the sample can be split into independent training and test sets, the error being estimated by the proportion of errors on the test data. For this *holdout estimate*, we have the distribution-free bound $\text{RMS}(\hat{\varepsilon}_{\text{holdout}}) \leq 1/\sqrt{4m}$, where m is the size of the test sample [6]. For $m = 100$, and any feature-label distribution, F , we have that $\text{RMS}(\hat{\varepsilon}_{\text{holdout}}) \leq 0.05$.

With small samples, training and error estimation must take place on the same data set. The consequences of training-set error estimation are seen in the following formula for the deviation variance:

$$\text{Var}_{\text{dev}}(\hat{\varepsilon}) = \sigma_{\hat{\varepsilon}}^2 + \sigma_{\varepsilon}^2 - 2\rho\sigma_{\hat{\varepsilon}}\sigma_{\varepsilon}, \quad (4.13)$$

where $\sigma_{\hat{\varepsilon}}^2$, σ_{ε}^2 , and ρ are the variance of the error estimate, the variance of the error, and the correlation between the estimated and true errors, respectively. The deviation variance is driven down by small variances or a correlation coefficient near 1. Unfortunately, for small samples, precisely the situation when one wishes to use training-set error estimation, neither condition typically holds.


Consider the popular *cross-validation error estimator*. For it, the error is estimated on the training data by randomly splitting the training data into k folds (subsets), S_n^i , for $i = 1, 2, \dots, k$, training k classifiers on $S_n - S_n^i$, for $i = 1, 2, \dots, k$, calculating the proportion of errors of each designed classifier on the appropriate left-out fold, and then averaging these proportions to obtain the cross-validation estimate of the originally designed classifier. Various enhancements are made, such as by repeating the process some number of times and averaging. Letting $k = n$ yields the *leave-one-out estimator*. The problem with cross-validation is that, for small samples, it typically has large variance and little correlation with the true error. Hence, although with large number of folds cross-validation does not suffer too much from bias, it typically has large deviation variance.

To illustrate with a materials dataset, consider predicting the formability of ABO_3 cubic perovskites. A dataset of 223 binary oxide systems, 34 of which can form cubic perovskites, is available in [7]. From this dataset we use two features that have been

shown to be predictive of formability: the octahedral factor and tolerance factor. We emulate the classification and error estimation procedure by drawing a small subset of examples from the full dataset for training, while using the left out points to estimate the ground truth true error. In particular, suppose that only 50 of the 223 compounds in the full dataset are available for classifier training, 8 of which can form a cubic structure and 42 cannot (the proportion is kept close to that of the full dataset). We train a radial-basis-function support vector machine (RBF-SVM) classifier on the 50 training points, use the same 50 points to estimate the error of this classifier using cross-validation with 5 folds and 10 repetitions, and approximate the true error rate of this classifier by evaluating the proportion of misclassified points among the 173 points left out of training (note the distribution free bound on the RMS of holdout applies here with $\text{RMS} \leq 1/\sqrt{4 \times 173} \approx 0.038$). We repeat this process 10,000 times to emulate the sampling procedure, each time drawing a different training set of 50 points. A scatter plot of the cross-validation error estimates and true errors is shown in Fig. 4.1, along with the least-squares regression line. The mean of the true errors and cross-validation estimates is indicated by a solid triangle, which shows that the cross-validation estimate is approximately unbiased (in fact, slightly high-biased). Because the class sizes are so unbalanced, the classifier error should be small, in particular, if we assume that $P(Y = 0) \approx 34/223 \approx 0.15$ then $\varepsilon_{\text{Bayes}}$ is upper bounded by $\min\{0.15, 0.85\} = 0.15$. Relative to the small true error, the dispersion of the scatter plot is very large. Moreover, the regression line has a slightly negative slope, certainly not a desirable property if one is going to estimate the true error by the cross-validation estimate.

What we observe in Fig. 4.1 is typical for small samples: large variance [8] and negligible regression between the true and estimated errors [9]. As seen, negatively sloping regression lines are possible; indeed, for cross-validation, negative correlation between the true and cross-validation estimated errors has been mathematically demonstrated in some basic models [10]. Such error estimates are worthless and

Fig. 4.1 Scatter plot and linear regression between cross-validation (horizontal axis) and the true error (vertical axis) with sample size 50 for RBF-SVM classification of the formability of ABO_3 cubic perovskites



can result in various problems that may not be immediately recognized: lack of reproducibility [11], optimistic bias when evaluating performance over several data sets [12], optimistic bias when considering several classification rules for a single classification problem [13], and inaccurate ROC curves [14]. Optimistic bias occurs because the high variance of the estimator gives a wide array of optimistic and pessimistic estimates when using different data sets or different classification rules, so that when one chooses the apparent best, he merely selects the one most optimistically biased—and bias can be severe.

4.4 Validity

A *pattern recognition model* (ψ, ε_ψ) consists of a classifier ψ and an error rate ε_ψ , where ε_ψ is simply a real number between 0 and 1. Intuitively, one might wish to say that (ψ, ε_ψ) is valid for the feature-label distribution F to the extent that ε_ψ approximates the classifier error, $\varepsilon[\psi]$, on F , where the degree of approximation is measured by some distance between ε_ψ and $\varepsilon[\psi]$. For a classifier ψ_n designed from a specific sample, this would mean that we want to measure some distance between the true error $\varepsilon = \varepsilon[\psi_n]$ and the estimated error $\hat{\varepsilon} = \hat{\varepsilon}[\psi_n]$, say $|\varepsilon - \hat{\varepsilon}|$. To know the true error we would need to know F , but if we knew F then we would use the Bayes classifier and not design a classifier from sample data. Since it is the precision of the error estimate that is of consequence, a natural way to proceed would be to characterize validity in terms of the precision of the error estimator $\hat{\varepsilon}[\psi_n] = \bar{\varepsilon}_n(S_n)$ as an estimator of $\varepsilon[\psi_n]$, say by $\text{RMS}(\hat{\varepsilon})$. This makes sense because the RMS measures the closeness of $\hat{\varepsilon}$ and ε across the sampling distribution. However, to compute the RMS again we need to know F , which we do not know. One way to proceed is to find a distribution-free bound on the RMS. For instance, for the leave-one-out error estimator with the discrete histogram rule and tie-breaking in the direction of class 0 [6],

$$\text{RMS}(\hat{\varepsilon}_{\text{loo}}) \leq \sqrt{\frac{1 + 6/e}{n} + \frac{6}{\sqrt{\pi}(n-1)}}. \quad (4.14)$$

The discrete histogram rule applies to a finite sample space $\{1, 2, \dots, b\}$, and defines $\psi_n(i) = 0$ if training samples with value i are labeled as class 0 at least as often as they are labeled class 1, and $\psi_n(i) = 1$ otherwise. Although this bound is distribution-free, it is useless for small samples: for $n = 200$ this bound is 0.506. In general, there are very few cases in which distribution-free bounds are known and, when they are known, they are useless for small samples.

Distribution-based bounds on the RMS are needed, which requires knowledge concerning the second-order moments of the joint distribution between the true and estimated errors. More generally, to fully understand an error estimator we need to know its joint distribution with the true error. Given that a classifier is epistemologically vacuous absent an accurate estimate of its error, one might think that over the

years much effort would have gone into studying the moments of the joint distribution between the true and estimated errors, especially the mixed second moment; however, this has not been the case. Going back half a century, there were some results on the mean and variance of some error estimators for the Gaussian model using LDA. In 1966, Hills obtained the expected value of the resubstitution and plug-in estimators in the univariate model with known common variance [15]. The *resubstitution error estimate* is simply a count of the classification errors on the training data. The *plug-in estimate* is found by using the data to estimate the feature-label distribution and then finding the error of the designed classifier on the estimated distribution. In 1972, Foley obtained the expected value of resubstitution in the multivariate model with known common covariance matrix [16]. In 1973, Sorum derived results for the expected value and variance for both resubstitution and leave-one-out in the univariate model with known common variance [17]. In 1973, McLachlan derived an asymptotic representation for the expected value of resubstitution in the multivariate model with unknown common covariance matrix [18]. In 1975, Moran obtained new results for the expected value of resubstitution and plug-in in the multivariate model with known covariance matrix [19]. In 1977, Goldstein and Wolf obtained the expected value of resubstitution for multinomial discrimination [20]. In 1992, Davison and Hall derived asymptotic representations for the expected value and variance of bootstrap and leave-one-out in the univariate Gaussian model with unknown and possibly different covariances [21]. Prior to 2005, we know of no other paper providing analytic results for moments of common error estimators. In total, none of these papers provide representation of the joint distribution or representation of second-order mixed moments, which are needed for the RMS.

Motivated by small samples ubiquitous in genomics, efforts commenced to obtain these types of representations, in particular, for the resubstitution and leave-one-out estimators. For the multinomial model, complete enumeration was used to obtain marginal distributions for both error estimators [10], followed by the full joint distributions [22]. Subsequently, exact closed-form representations for second-order moments, including the mixed moments, were obtained, thereby providing exact RMS representations for both estimators [10]. For the Gaussian model using LDA, in 2009 exact marginal distributions for both estimators in the univariate model (with known but not necessarily equal class variances) and approximations in the multivariate model (with known and equal class covariance matrices) were obtained [23]. Subsequently, these were extended to joint distributions of the true and estimated errors in a Gaussian model [24]. Recently exact closed-form representations for the second-order moments in the univariate model without assuming equal covariances were discovered, thereby providing exact expressions of the RMS for both estimators [25]. Moreover, double asymptotic representations for the second-order moments in the multivariate model, sample size and dimension approaching infinity at a fixed rate between the two, were found, thereby providing double asymptotic expressions for the RMS [26]. Finite-sample approximations from the double asymptotic method have been shown to possess better accuracy than various simple asymptotic representations, although much more work is needed on this issue [27, 28].




Fig. 4.2 **a** RMS (y-axis) as a function of the Bayes error (x-axis) for leave-one-out with dimension $D = 10$ and sample sizes $n = 20, 40, 60$; **b** $\max\text{Bayes}(\lambda)$ (y-axis) as a function of RMS (x-axis) corresponding to the RMS curves in part **(a)**

To utilize mixed-moment theory, prior knowledge is required, in the sense that the actual (unknown) feature-label distribution belongs to some *uncertainty class*, \mathcal{U} , of feature-label distributions. Once RMS representations have been obtained for feature-label distributions in \mathcal{U} , distribution-based RMS bounds follow: $\text{RMS}(\hat{\epsilon}) \leq \max_{G \in \mathcal{U}} \text{RMS}(\hat{\epsilon}|G)$, where $\text{RMS}(\hat{\epsilon}|G)$ is the RMS of the error estimator under the assumption that the feature-label distribution is G . We do not know the actual feature-label distribution precisely, but prior knowledge allows us to bound the RMS. For instance, consider using LDA with a feature-label distribution having two equally probable Gaussian class-conditional densities sharing a known covariance matrix. For this model the Bayes error is a one-to-one decreasing function of the distance, m , between the means. Figure 4.2a shows the RMS to be a one-to-one increasing function of the Bayes error for leave-one-out in dimension $D = 10$ and sample sizes $n = 20, 40, 60$, the RMS and Bayes errors being on the y and x axes, respectively.

Assuming a parameterized model in which the RMS is an increasing function of the Bayes error, ϵ_{Bayes} , we can pose the following question: Given sample size n and $\lambda > 0$, what is the maximum value, $\max\text{Bayes}(\lambda)$, of the Bayes error such that $\text{RMS}(\hat{\epsilon}) \leq \lambda$? If RMS is the measure of validity and λ represents the largest acceptable RMS for the classifier model to be considered meaningful, then the epistemological requirement is characterized by $\max\text{Bayes}(\lambda)$. Given the relationship between model parameters and the Bayes error, the inequality $\epsilon_{\text{Bayes}} \leq \max\text{Bayes}(\lambda)$ can be solved in terms of the parameters to arrive at a necessary modeling assumption. In the preceding Gaussian example, since ϵ_{Bayes} is a decreasing function of m , we obtain an inequality $m \geq m(\lambda)$. Figure 4.2b shows the $\max\text{Bayes}(\lambda)$ curves corresponding to the RMS curves in Fig. 4.2a [29]. These curves show that, assuming Gaussian class-conditional densities and a known common covariance matrix, further assumptions must be made to insure that the RMS is sufficiently small to make the classifier model meaningful.

To have scientific content, small-sample classification requires prior knowledge. Regarding the feature-label distribution there are two extremes: (1) the feature-label distribution is known, in which case the entire classification problem collapses to

finding a Bayes classifier and Bayes error, so there is no classifier design or error estimation issue; and (2) the uncertainty class consists of all feature-label distributions, the distribution-free case, and we typically have no bound on performance, or one that is too loose for practice. In the middle ground, there is a trade-off between the size of the uncertainty class and the size of the sample. The uncertainty class must be sufficiently constrained (equivalently, the prior knowledge must be sufficiently great) that an acceptable bound can be achieved with an acceptable sample size.

We have focused on cross-validation for two reasons: (1) it is probably the most commonly used training-data-based error estimator and (2) its moments, along with resubstitution, are the most studied. Another often employed re-sampling-based error estimator is bootstrap [30]. It generally has smaller variance than cross-validation; however, it can suffer from significant bias, depending on the feature-label distribution and classification rule. Analytic representation of bootstrap expectation in the Gaussian model with LDA classification has recently been found and, since the bootstrap has a weighting parameter, under these conditions it can be weighted to be unbiased [31]. In general, using its free weight together with the fact that the bootstrap is formed by a convex combination, given the model and the classification rule, the Lagrangian multiplier technique can be used to determine a weight that minimizes the RMS between this optimized bootstrap and the true error [32].

Given that one needs a distributional model to assure satisfactory performance for classifier error estimation, a natural way to proceed is to define a prior distribution over the uncertainty class of feature-label distributions and then find an optimal minimum-mean-square-error (MMSE) error estimator relative to the prior [33]. This results in a Bayesian approach with the uncertainty class governed by the prior distribution and the data being used to construct a posterior distribution that quantifies everything we know about the feature-label distribution. In this way we can incorporate prior knowledge in the whole classification procedure, both classifier design and error estimation.

4.5 MMSE Error Estimation

If the class-conditional distribution for class y , denoted $f_{\theta_y}(\mathbf{x}|y)$, is parameterized by θ_y , then the feature-label distribution is completely specified by the modeling parameters $\theta = [c, \theta_0, \theta_1]$, where $c = P(Y = 0)$. Writing the parameter space of θ_y as Θ_y , the parameter space of θ is $\Theta = [0, 1] \times \Theta_0 \times \Theta_1$. We denote the prior distribution on θ by $\pi(\theta)$ and the posterior, derived from a random sample of size n with n_y points from class y , by $\pi^*(\theta)$. Here we assume that c is independent from θ_0 and θ_1 prior to observing the data and denote its prior by $\pi(c)$. Under a given sampling method, the posterior $\pi^*(c)$ for c may be obtained from the number of sample points in each class using Bayes' rule. For instance, under random sampling and assuming a beta(α^0, α^1) prior for c , the posterior of c is also beta with hyperparameters $\alpha^0 + n_0$ and $\alpha^1 + n_1$. In particular, letting B be the beta function,

$$\pi^*(c) = \frac{c^{\alpha^0+n_0-1}(1-c)^{\alpha^1+n_1-1}}{B(\alpha^0+n_0, \alpha^1+n_1)}, \quad (4.15)$$

$$E_{\pi^*}[c] = \frac{n_0 + \alpha^0}{n + \alpha^0 + \alpha^1}, \quad (4.16)$$

where E_{π^*} represents expectation relative to the posterior (conditioned on the sample). A uniform prior on c is achieved with $\alpha^0 = \alpha^1 = 1$.

The Bayesian framework in [33, 34] not only assumes that c is independent, but that c , θ_0 and θ_1 are all independent prior to observing the data. Writing the prior for θ_y as $\pi(\theta_y)$, this means that $\pi(\theta) = \pi(c)\pi(\theta_0)\pi(\theta_1)$. We also write the posterior as $\pi^*(\theta_y)$, where it has been shown that independence is preserved after observing the data, that is, $\pi^*(\theta) = \pi^*(c)\pi^*(\theta_0)\pi^*(\theta_1)$. $\pi^*(\theta_y)$ is proportional to the product of the prior and a likelihood function for sample points observed from the corresponding class:

$$\pi^*(\theta_y) \propto \pi(\theta_y) \prod_{i=1}^{n_y} f_{\theta_y}(\mathbf{x}_i^y | y), \quad (4.17)$$

where \mathbf{x}_i^y is the i th sample point in class y and the constant of proportionality is found by normalizing the integral of $\pi^*(\theta_y)$ to 1. When the prior is a proper density, this follows from Bayes' rule; if $\pi(\theta_y)$ is improper (i.e., if the integral of $\pi(\theta_y)$ cannot be normalized to 1), then this is taken as a definition, but in all cases it is mandatory that $\pi^*(\theta_y)$ be a proper density. Priors quantify the information we have about the distribution before observing the data. We have the option of using flat, or non-informative, priors, as long as the posterior is a valid density function. Alternatively, informative priors can supplement the classification problem with additional information.

The Bayesian model characterizes our initial uncertainty in the actual distribution through the prior. As we observe sample points, this uncertainty should converge to a certainty on the true distribution. More precisely, it has been proven in [35] that under mild regularity conditions, the posteriors converge to a point mass at the true parameters for an independent covariance Gaussian model, which we will discuss shortly. More informative priors may help the posteriors converge faster, but, essentially, as long as the prior does not assign zero probability to any a neighborhood around the true distribution, convergence is assured.

The Bayesian model defines priors on the feature-label distribution itself; nevertheless, the posteriors of the distribution parameters imply a (sample-conditioned) distribution on the true classifier error. This randomness in the true error comes from our uncertainty in the underlying feature-label distribution (given the sample), which is in contrast to the classical analysis discussed in previous sections, where randomness in the true error for a fixed distribution comes only from randomness in the trained classifier through the sampling distribution. In addition, we may speak of moments of the true error for a fixed sample and classifier.

The true error of a designed classifier ψ_n may be decomposed as

$$\varepsilon(\theta, \psi_n) = c\varepsilon^0(\theta_0, \psi_n) + (1-c)\varepsilon^1(\theta_1, \psi_n), \quad (4.18)$$

where $\varepsilon^y(\theta_y, \psi_n)$ is the probability that ψ_n mislabels a class y point under true parameter θ_y . Since the Bayesian framework quantifies uncertainty in the feature-label distribution parameters, we may find the MMSE estimate of the true error, $\hat{\varepsilon}(\psi_n, S_n)$, which is equal to the first moment of the true error conditioned on the observed sample [33]. We call this the *Bayesian error estimate*. As long as c is independent from θ_0 and θ_1 *a posteriori*,

$$\begin{aligned} \hat{\varepsilon}(\psi_n, S_n) &= E_{\pi^*}[\varepsilon(\theta, \psi_n)] \\ &= E_{\pi^*}[c]\hat{\varepsilon}^0(\psi_n, S_n) + (1 - E_{\pi^*}[c])\hat{\varepsilon}^1(\psi_n, S_n), \end{aligned} \quad (4.19)$$

where $\hat{\varepsilon}^y(\psi_n, S_n) = E_{\pi^*}[\varepsilon^y(\theta_y, \psi_n)]$ is the posterior expected error contributed by class y . Both $\hat{\varepsilon}$ and $\hat{\varepsilon}^y$ are functions of the classifier ψ_n , and the sample via π^* .

The expectation of c depends on our prior model for c , but is straightforward to find analytically. For example, if c is fixed, then the expectation can be replaced with the fixed value of c , and if c has a beta(α^0, α^1) prior, then $E_{\pi^*}[c]$ is available in (4.16). Representation for $\hat{\varepsilon}^y(\psi_n, S_n)$ is known for the discrete and independent covariance Gaussian models [33, 34]. Owing to convergence of the posteriors, classical frequentist consistency holds for Bayesian error estimators in both models for any fixed distribution in the parameterized family [35].

We next present an example illustrating the optimal performance of MMSE error estimation. Consider a $D = 5$ dimensional Gaussian model with a uniform prior on c and independent arbitrary covariance matrices. In particular, we assume normal-inverse-Wishart priors with hyperparameters $v_y = \kappa_y = 25$, $\mathbf{m}_0 = [0, 0, 0, 0, 0]$, $\mathbf{m}_1 = [1, 0, 0, 0, 0]$, and $S_y = 13.19I_5$, where I_D is a $D \times D$ identity matrix. This is a moderately informative prior where the expected mean of class y is \mathbf{m}_y and the expected covariance for both classes is $0.7413^2 I_5$. We generate 100,000 feature-label distributions from the prior, each including a random realization for c and random μ_y and Σ_y pairs for each class $y \in \{0, 1\}$. For each fixed feature-label distribution, we generate 10 samples of a given size n ranging from 30 to 200, first determining the number of points in each class by drawing n_0 from a binomial(n, c) distribution, and then, for each class, drawing the appropriate number of i.i.d. points from a Gaussian distribution with the corresponding mean and covariance pair. From each sample we train an LDA classifier, we evaluate the true error of the trained classifier under the corresponding true feature-label distribution, and we estimate the error of this classifier using 4 training-data-based methods: the MMSE error estimator (Bayes), resubstitution (resub), cross-validation (cv), and bolstered resubstitution (bol). Bolstered resubstitution is similar to resubstitution except that each point of the training set is replaced with a density kernel and the error is estimated by integrating each kernel over the classifier decision region disagreeing with the label at the point, thereby “spreading” the incorrect mass and giving more error weight to incorrectly

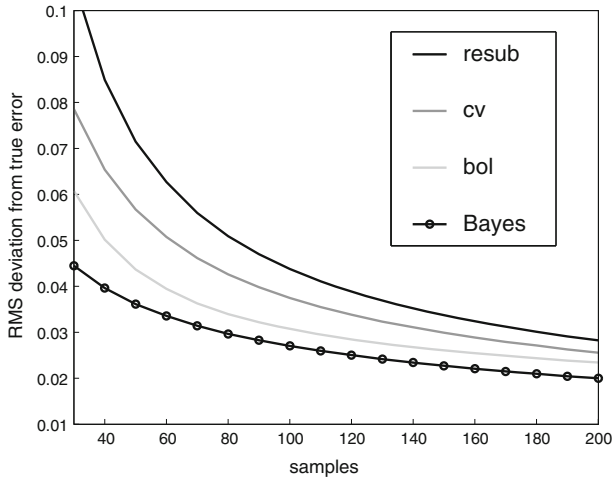


Fig. 4.3 RMS deviation from true error for linear classification of Gaussian distributions, averaged over all distributions and samples using a proper prior with $D = 5$

labeled points near the decision boundary (see [36] for details). We then approximate an RMS for each error estimator, that is, we evaluate the square root of the average square difference between each error estimator and the true error, where the average is taken over all 100,000 feature-label distributions and 10 samples. A graph of the RMS with respect to sample size is provided in Fig. 4.3. The performance of the MMSE error estimator here, averaged over all distributions and samples under the assumed prior, is optimal, outperforming all other error estimators, as it must. This does not mean that performance is optimal for any fixed feature-label distribution, only that it is optimal on average.

4.6 Optimal Bayesian Classification

An *optimal Bayesian classifier* (OBC) ψ_{OBC} is any classifier satisfying

$$E_{\pi^*} [\varepsilon(\theta, \psi_{\text{OBC}})] \leq E_{\pi^*} [\varepsilon(\theta, \psi)] \quad (4.20)$$

for all $\psi \in \mathcal{C}$, where \mathcal{C} is a family of classifiers. Under the Bayesian framework,

$$P(\psi(\mathbf{X}) \neq Y | S_n) = E_{\pi^*} [P(\psi(\mathbf{X}) \neq Y | \theta, S_n)] = E_{\pi^*} [\varepsilon(\theta, \psi)] = \hat{\varepsilon}(\psi, S_n).$$

Thus, optimal Bayesian classifiers minimize the misclassification probability relative to the assumed model or, equivalently, minimize the Bayesian error estimate.

The following representation of the Bayesian error estimator facilitates a straightforward approach for finding an OBC [33, 34]: If ψ is a fixed classifier defined by

$\psi(\mathbf{x}) = 0$ if $\mathbf{x} \in R_0$ and $\psi(\mathbf{x}) = 1$ if $\mathbf{x} \in R_1$, where R_0 and R_1 are measurable sets partitioning the sample space, then the Bayesian error estimator is given by

$$\hat{\varepsilon}(\psi, S_n) = E_{\pi^*[c]} \int_{R_1} f(\mathbf{x}|0) d\mathbf{x} + (1 - E_{\pi^*[c]}) \int_{R_0} f(\mathbf{x}|1) d\mathbf{x} \quad (4.21)$$

$$= \int_{\mathfrak{R}^D} (E_{\pi^*[c]} f(\mathbf{x}|0) \mathbf{I}_{\mathbf{x} \in R_1} + (1 - E_{\pi^*[c]}) f(\mathbf{x}|1) \mathbf{I}_{\mathbf{x} \in R_0}) d\mathbf{x}, \quad (4.22)$$

where \mathbf{I}_E is an indicator function equal to one if E is true and zero otherwise, and

$$f(\mathbf{x}|y) = \int_{\Theta_y} f_{\theta_y}(\mathbf{x}|y) \pi^*(\theta_y) d\theta_y, \quad (4.23)$$

which is called the *effective class-conditional density* with respect to the posterior.

An OBC can be found by brute force using the closed-form solutions for the expected true error (the Bayesian error estimator), when available; however, if \mathcal{C} is the set of all classifiers (with measurable decision regions), then an OBC, in the presence of model uncertainty, can be found analogously to a Bayes classifier under a known feature-label distribution. To wit, an OBC relative to the set of all classifiers with measurable decision regions exists and is given pointwise by [37]

$$\psi_{\text{OBC}}(\mathbf{x}) = \begin{cases} 0, & \text{if } E_{\pi^*[c]} f(\mathbf{x}|0) \geq (1 - E_{\pi^*[c]}) f(\mathbf{x}|1) \\ 1, & \text{otherwise} \end{cases}. \quad (4.24)$$

To find an OBC we can average the class-conditional densities $f_{\theta_y}(\mathbf{x}|y)$ relative to the posterior distribution to obtain the effective class-conditional density, $f(\mathbf{x}|y)$, whereby an OBC is found via (4.24). Essentially, the OBC is the Bayes classifier using $f(\mathbf{x}|0)$ and $f(\mathbf{x}|1)$ as the true class-conditional distributions.

In regard to both optimal Bayesian classification and MMSE error estimation, $f(\mathbf{x}|y)$ contains all of the necessary information in the model about the class-conditional distributions and we do not have to deal with the priors directly. Upon defining a model, we find $f(\mathbf{x}|y)$, which depends on the sample because it depends on π^* , and then several problems are solved by treating $f(\mathbf{x}|y)$ as the true distribution: optimal (unconstrained) classification, the optimal error estimate for the optimal classifier, and the optimal error estimate for arbitrary classifiers.

Henceforth, we will only consider optimal Bayesian classifiers over the space of all classifiers. Moreover, note that if $E_{\pi^*[c]} = 0$ then the OBC is a constant given by $\psi_{\text{OBC}} = 1$, and if $E_{\pi^*[c]} = 1$ then $\psi_{\text{OBC}} = 0$.

4.7 The Gaussian Model

In the Gaussian model, the uncertainty class is determined by the parameters $\theta_y = [\mu_y, \Lambda_y]$, where μ_y is the mean of the class-conditional distribution and Λ_y is a collection of parameters that determine the covariance matrix, Σ_y , of the class.

By defining Σ_y as a function of Λ_y , we may impose a structure on the covariance. Three types of models are considered in [37]: a fixed covariance model ($\Sigma_y = \Lambda_y$ is known perfectly), a scaled identity covariance model having uncorrelated features with equal variances ($\Lambda_y = \sigma_y^2$ is a scalar and $\Sigma_y = \sigma_y^2 I_D$), and an arbitrary (valid) covariance model ($\Sigma_y = \Lambda_y$ may be any invertible covariance matrix). Here we consider the known and arbitrary-covariance models in detail. If the arbitrary covariance model is used in both classes, then we assume that the covariance matrices in each class are independent. The parameter space of μ_y is \Re^D , and the parameter space of Λ_y must be carefully defined to permit only valid covariance matrices. As Σ_y and Λ_y are equivalent in the cases we will consider, we will write Σ_y in place of Λ_y without explicitly showing its dependence on Λ_y , i.e., we write Σ_y rather than $\Sigma_y(\Lambda_y)$. We also denote a multivariate Gaussian distribution with mean μ and covariance Σ by $f_{\mu, \Sigma}(\mathbf{x})$, so that the parameterized class-conditional distributions can be written as $f_{\theta_y}(\mathbf{x}|y) = f_{\mu_y, \Sigma_y}(\mathbf{x})$.

Under the independence assumption, c , $\theta_0 = [\mu_0, \Sigma_0]$ and $\theta_1 = [\mu_1, \Sigma_1]$ are all independent prior to observing the data, so that $\pi(\theta) = \pi(c)\pi(\theta_0)\pi(\theta_1)$. Assuming $\pi(c)$ and $\pi^*(c)$ have been established, we must define priors $\pi(\theta_y)$ and find posteriors $\pi^*(\theta_y)$ for both classes. We begin by specifying conjugate priors for θ_0 and θ_1 . Define

$$f_m(\mu; \nu, \mathbf{m}, \Sigma) = |\Sigma|^{-1/2} \exp\left(-\frac{\nu}{2}(\mu - \mathbf{m})^T \Sigma^{-1}(\mu - \mathbf{m})\right), \quad (4.25)$$

$$f_c(\Sigma; \kappa, S) = |\Sigma|^{-(\kappa+D+1)/2} \exp\left(-\frac{1}{2}\text{trace}(S\Sigma^{-1})\right), \quad (4.26)$$

which involve several constants: ν , \mathbf{m} , κ and S . If $\nu > 0$, then f_m is a (scaled) Gaussian distribution with mean \mathbf{m} and covariance Σ/ν . If $\kappa > D - 1$ and S is symmetric and positive definite, then f_c is a (scaled) inverse-Wishart (κ, S) distribution. However, to allow for improper priors we do not necessarily require f_m and f_c to be normalizable.

Consider class $y \in \{0, 1\}$. In the arbitrary covariance model, we assume Σ_y is invertible with probability 1 and that for invertible Σ_y the prior for θ_y is of the form

$$\pi(\theta_y) = \pi(\mu_y|\Sigma_y)\pi(\Sigma_y), \quad (4.27)$$

where

$$\pi(\mu_y|\Sigma_y) \propto f_m(\mu_y; \nu_y, \mathbf{m}_y, \Sigma_y), \quad (4.28)$$

$$\pi(\Sigma_y) \propto f_c(\Sigma_y; \kappa_y, S_y), \quad (4.29)$$

ν_y is a real number, \mathbf{m}_y is a length D real vector, κ_y is a real number, and S_y is a symmetric non-negative definite $D \times D$ matrix. If $\nu_y > 0$, then the prior for the mean conditioned on the covariance, $\pi(\mu_y|\Sigma_y)$, is proper and Gaussian with mean \mathbf{m}_y and covariance Σ_y/ν_y . The hyperparameter \mathbf{m}_y is the prior expected mean of class y , where the larger ν_y is the more confident we are that μ_y is close to \mathbf{m}_y .

In the arbitrary covariance model, $\pi(\Sigma_y)$ is a proper inverse-Wishart distribution if $\kappa_y > D - 1$ and S_y is symmetric and positive definite. If in addition $\nu_y > 0$, then $\pi(\theta_y)$ is a normal-inverse-Wishart distribution, which is the conjugate prior for the mean and covariance when sampling from normal distributions [38, 39]. As long as $\kappa_y > D + 1$, the prior mean of Σ_y exists and is given by $E_\pi[\Sigma_y] = S_y/(\kappa_y - D - 1)$. Thus, S_y determines the expected shape of the covariance, where the actual expected covariance is scaled. If S_y is scaled appropriately, then the larger κ_y is the more certainty we have about the covariance Σ_y .

In this model, the posterior has the same form as the prior [34],

$$\pi^*(\theta_y) \propto f_m(\mu_y; \nu_y^*, \mathbf{m}_y^*, \Sigma_y) f_c(\Sigma_y; \kappa_y^*, S_y^*), \quad (4.30)$$

with updated hyperparameters

$$\nu_y^* = \nu_y + n_y, \quad (4.31)$$

$$\mathbf{m}_y^* = \frac{\nu_y \mathbf{m}_y + n_y \hat{\mu}_y}{\nu_y + n_y}, \quad (4.32)$$

$$\kappa_y^* = \kappa_y + n_y, \quad (4.33)$$

$$S_y^* = S_y + (n_y - 1) \hat{\Sigma}_y + \frac{\nu_y n_y}{\nu_y + n_y} (\hat{\mu}_y - \mathbf{m}_y)(\hat{\mu}_y - \mathbf{m}_y)^T, \quad (4.34)$$

where $\hat{\mu}_y$ and $\hat{\Sigma}_y$ are the sample mean and sample covariance of the n_y training points in class y . Improper priors can still be used so long as the posterior is proper: for a proper posterior in the arbitrary covariance model we require $\nu_y^* > 0$, $\kappa_y^* > D - 1$, and that S_y^* is symmetric and positive definite. The previous discussion on properties of a proper prior again apply to the posterior, namely that $\pi^*(\mu_y | \Sigma_y)$ must be a valid Gaussian distribution and $\pi^*(\Sigma_y)$ must be a valid inverse-Wishart distribution.

Continuing with the arbitrary covariance model, the parameter space of θ_y is the product of the space of all valid mean vectors, \mathfrak{R}^D , and the space of all positive-definite matrices, which we denote by $\Sigma_y > 0$. By definition,

$$f(\mathbf{x} | y) = \int_{\Sigma_y > 0} \int_{\mathfrak{R}^D} f_{\mu_y, \Sigma_y}(\mathbf{x}) \pi^*(\mu_y | \Sigma_y) \pi^*(\Sigma_y) d\mu_y d\Sigma_y. \quad (4.35)$$

Given that $\pi^*(\mu_y | \Sigma_y)$ is Gaussian and $\pi^*(\Sigma_y)$ is inverse-Wishart, one can show that evaluation of the double integral yields a multivariate student's t -distribution [34]:

$$f(\mathbf{x} | y) = \frac{1}{k_y^{D/2} \pi^{D/2} |\Psi_y|^{1/2}} \times \frac{\Gamma\left(\frac{k_y + D}{2}\right)}{\Gamma\left(\frac{k_y}{2}\right)} \left(1 + \frac{1}{k_y} (\mathbf{x} - \mathbf{m}_y^*)^T \Psi_y^{-1} (\mathbf{x} - \mathbf{m}_y^*)\right)^{-\frac{k_y + D}{2}}, \quad (4.36)$$

with location vector \mathbf{m}_y^* , scale matrix $\Psi_y = \frac{v_y^*+1}{(\kappa_y^*-D+1)v_y^*} S_y^*$ and $k_y = \kappa_y^* - D + 1$ degrees of freedom. This distribution is proper because $\frac{v_y^*+1}{(\kappa_y^*-D+1)v_y^*} > 0$ and S_y^* is symmetric and positive definite (so the scale matrix is symmetric and positive definite) and $\kappa_y^* - D + 1 > 0$. As long as $\kappa_y^* > D$ the mean of this distribution is \mathbf{m}_y^* , and as long as $\kappa_y^* > D + 1$ the variance is $\frac{v_y^*+1}{(\kappa_y^*-D-1)v_y^*} S_y^*$.

Switching gears to the known covariance model, now assume that the prior for θ_y is of the form

$$\pi(\theta_y) = \pi(\mu_y | \Sigma_y) \pi(\Sigma_y), \quad (4.37)$$

where $\pi(\mu_y | \Sigma_y)$ is given in (4.28) and $\pi(\Sigma_y)$ is simply a point mass at the known value of Σ_y . Again, we require that v_y be a real number and \mathbf{m}_y be a length D real vector, where if $v_y > 0$ then the prior for the mean is proper and Gaussian with mean \mathbf{m}_y and covariance Σ_y/v_y . Also as before, the posterior has the same form as the prior with the same hyperparameter update equations, (4.31) and (4.32). For a proper posterior, we require $v_y^* > 0$. In the known covariance model, we may simplify the effective density in (4.35) as,

$$f(\mathbf{x}|y) = \int_{\mathbb{R}^D} f_{\mu_y, \Sigma_y}(\mathbf{x}) \pi^*(\mu_y | \Sigma_y) d\mu_y, \quad (4.38)$$

where Σ_y is the known covariance matrix. One can show that this integral yields a proper Gaussian distribution with mean \mathbf{m}_y^* and covariance $\frac{v_y^*+1}{v_y^*} \Sigma_y$ [34]:

$$f(\mathbf{x}|y) = \frac{(v_y^*)^{D/2}}{(v_y^*+1)^{D/2} (2\pi)^{D/2} |\Sigma_y|^{1/2}} \exp\left(-\frac{v_y^*}{2(v_y^*+1)} (\mathbf{x} - \mathbf{m}_y^*)^T \Sigma_y^{-1} (\mathbf{x} - \mathbf{m}_y^*)\right). \quad (4.39)$$

4.8 Optimal Bayesian Classifier in the Gaussian Model

There are three cases to consider when finding the OBC: the covariances are known in both classes, a covariance is known in only one class, and the covariances are unknown in both classes (for derivation details see [37]). It is interesting to consider the shape of the decision boundary for the OBC as compared to the shapes of the decision boundaries for each feature-label distribution in the uncertainty class; in particular, note how the effective class-conditional distributions become multivariate student's t -distributions.

When both covariances are known, in the previous section we showed that the effective class-conditional distributions are Gaussian with mean \mathbf{m}_y^* and covariance $\frac{v_y^*+1}{v_y^*} \Sigma_y$ for $y \in \{0, 1\}$. The OBC, $\psi_{\text{OBC}}(\mathbf{x})$, is the optimal classifier between the effective Gaussians with class 0 probability $E_{\pi^*}[c]$, and is of the same form as the

Bayes classifier in (4.5) and (4.6) with discriminant $g_{\text{OBC}}(\mathbf{x})$ given by

$$\begin{aligned}
 A_{\text{OBC}} &= -\frac{1}{2} \left(\frac{v_1^*}{v_1^* + 1} \Sigma_1^{-1} - \frac{v_0^*}{v_0^* + 1} \Sigma_0^{-1} \right), \\
 \mathbf{a}_{\text{OBC}} &= \frac{v_1^*}{v_1^* + 1} \Sigma_1^{-1} \mathbf{m}_1^* - \frac{v_0^*}{v_0^* + 1} \Sigma_0^{-1} \mathbf{m}_0^*, \\
 b_{\text{OBC}} &= -\frac{1}{2} \left(\frac{v_1^*}{v_1^* + 1} \mathbf{m}_1^{*T} \Sigma_1^{-1} \mathbf{m}_1^* - \frac{v_0^*}{v_0^* + 1} \mathbf{m}_0^{*T} \Sigma_0^{-1} \mathbf{m}_0^* \right) \\
 &\quad + \ln \left(\frac{1 - E_{\pi^*}[c]}{E_{\pi^*}[c]} \left(\frac{v_1^*(v_0^* + 1)}{v_0^*(v_1^* + 1)} \right)^{D/2} \left(\frac{|\Sigma_0|}{|\Sigma_1|} \right)^{1/2} \right). \quad (4.40)
 \end{aligned}$$

The expected true error for the OBC is simply the true error for this quadratic classifier under the effective Gaussian distributions.

If the covariance is known in only one class and modeled as arbitrary in the other, then the effective class-conditional distribution for the known class, say class 0, is Gaussian and the other class is a multivariate student's t -distribution, hence,

$$f(\mathbf{x}|0) = \frac{(v_0^*)^{D/2}}{(v_0^* + 1)^{D/2} (2\pi)^{D/2} |\Sigma_0|^{1/2}} \exp \left(-\frac{v_0^*}{2(v_0^* + 1)} (\mathbf{x} - \mathbf{m}_0^*)^T \Sigma_0^{-1} (\mathbf{x} - \mathbf{m}_0^*) \right), \quad (4.41)$$

$$f(\mathbf{x}|1) = \frac{1}{k_1^{D/2} \pi^{D/2} |\Psi_1|^{1/2}} \times \frac{\Gamma\left(\frac{k_1 + D}{2}\right)}{\Gamma\left(\frac{k_1}{2}\right)} \left(1 + \frac{1}{k_1} (\mathbf{x} - \mathbf{m}_1^*)^T \Psi_1^{-1} (\mathbf{x} - \mathbf{m}_1^*) \right)^{-\frac{k_1 + D}{2}}, \quad (4.42)$$

and, from (4.36), $\Psi_1 = \frac{v_1^* + 1}{(\kappa_1^* - D + 1)v_1^*} S_1^*$ and $k_1 = \kappa_1^* - D + 1$. The discriminant of the OBC can be simplified to

$$\begin{aligned}
 g_{\text{OBC}}(\mathbf{x}) &= \frac{v_0^*}{v_0^* + 1} (\mathbf{x} - \mathbf{m}_0^*)^T \Sigma_0^{-1} (\mathbf{x} - \mathbf{m}_0^*) \\
 &\quad - (k_1 + D) \ln \left(1 + \frac{1}{k_1} (\mathbf{x} - \mathbf{m}_1^*)^T \Psi_1^{-1} (\mathbf{x} - \mathbf{m}_1^*) \right) + K,
 \end{aligned}$$

where

$$K = 2 \ln \left(\frac{1 - E_{\pi^*}[c]}{E_{\pi^*}[c]} \left(\frac{2(v_0^* + 1)}{v_0^* k_1} \right)^{D/2} \left(\frac{|\Sigma_0|}{|\Psi_1|} \right)^{1/2} \frac{\Gamma((k_1 + D)/2)}{\Gamma(k_1/2)} \right). \quad (4.43)$$

The form of the OBC is not necessarily linear or quadratic.

When the covariances of both classes are unknown and arbitrary, the effective class-conditional distribution for each class is multivariate student's t with location vector \mathbf{m}_y^* , scale matrix Ψ_y and k_y degrees of freedom, as given in (4.36). The discriminant of the OBC can be simplified to

$$g_{\text{OBC}}(\mathbf{x}) = K \left(1 + \frac{1}{k_0} (\mathbf{x} - \mathbf{m}_0^*)^T \Psi_0^{-1} (\mathbf{x} - \mathbf{m}_0^*) \right)^{k_0+D} - \left(1 + \frac{1}{k_1} (\mathbf{x} - \mathbf{m}_1^*)^T \Psi_1^{-1} (\mathbf{x} - \mathbf{m}_1^*) \right)^{k_1+D}, \quad (4.44)$$

where

$$K = \left(\frac{1 - E_{\pi^*}[c]}{E_{\pi^*}[c]} \right)^2 \left(\frac{k_0}{k_1} \right)^D \frac{|\Psi_0|}{|\Psi_1|} \left(\frac{\Gamma(k_0/2)\Gamma((k_1+D)/2)}{\Gamma((k_0+D)/2)\Gamma(k_1/2)} \right)^2. \quad (4.45)$$

This classifier has a polynomial decision boundary that is not necessarily linear or quadratic as long as k_0 and k_1 are integers, which is satisfied for arbitrary covariance models with independent covariances if κ_0 and κ_1 are integers.

Consider an example with $D = 2$ features, where each class is equally likely ($c = 0.5$) and the class-conditional distributions are known to be Gaussian with unequal and arbitrary invertible covariances. We assume that the mean and covariance pairs associated with each class are independent and given by a normal-inverse-Wishart prior with hyperparameters $\nu_0 = \nu_1 = 0$, $\kappa_0 = \kappa_1 = 0$, $\mathbf{m}_0 = \mathbf{m}_1 = [0, 0]$ and $S_0 = S_1 = -2I_2$. Further suppose that we observe 40 sample points from class 0 and 4 sample points from class 1, where the sample mean of class 0 is $[0, 0]$, the sample mean of class 1 is $[1, 1]$, and the sample covariance of both classes is I_2 . Then the posteriors are proper normal-inverse-Wishart distributions given by hyperparameters $\nu_0^* = \kappa_0^* = 40$, $\mathbf{m}_0^* = [0, 0]$, $S_0^* = 37I_2$, $\nu_1^* = \kappa_1^* = 4$, $\mathbf{m}_1^* = [1, 1]$, and $S_1^* = I_2$. We will consider three classifiers. The first is a plug-in classifier, which substitutes the posterior expected means and covariances into the Bayes classifier, that is, we assume that μ_0 is $E_{\pi^*}[\mu_0] = \mathbf{m}_0^* = [0, 0]$, μ_1 is $\mathbf{m}_1^* = [1, 1]$, and Σ_y is $E_{\pi^*}[\Sigma_y] = \frac{1}{\kappa_y^* - D - 1} S_y^* = I_2$. Since the expected covariances are equal, this classifier is linear. Note for this prior the posterior expected parameters coincide with the sample means and covariances, so that the plug-in classifier is also equivalent to an LDA classifier. The second classifier that we consider is a state-constrained optimal Bayesian classifier (SCOBC), which is found by searching across mean and covariance pairs in the uncertainty class for a Bayes classifier having minimal expected error [40]. Since the Bayes classifier for any Gaussian distribution is quadratic, the SCOBC is quadratic. Finally, we have the optimal Bayesian classifier, which is available in closed form in (4.44). Since the effective densities are not Gaussian but multivariate student's t -distributions, the OBC has a polynomial decision boundary of greater than quadratic order. Figure 4.4 shows the plug-in classifier (light gray), SCOBC (dark gray) and OBC (black). Level curves for the class-conditional distributions corresponding to the expected parameters in the posteriors used in the plug-in rule are

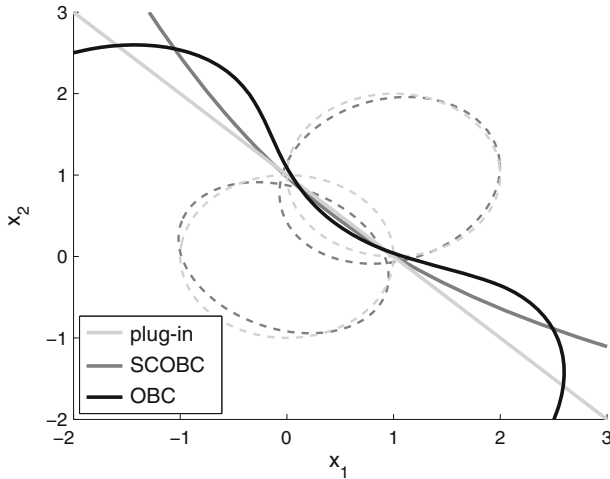


Fig. 4.4 Classifiers for an independent arbitrary covariance Gaussian model with $D = 2$ features and proper posteriors. The optimal Bayesian classifier is polynomial with expected true error 0.2007 (averaged over the posterior on the uncertainty class of states), the state-constrained optimal Bayesian classifier is quadratic with expected true error 0.2061 and the plug-in classifier is linear with expected true error 0.2078

shown in light gray dashed lines, and level curves for the distributions corresponding to the optimal parameters found in the SCOBC are shown in dark gray dashed lines. Each classifier is quite distinct, and in particular, the optimal Bayesian classifier is non-quadratic even though all class-conditional distributions in the uncertainty class are Gaussian.

4.9 Concluding Remarks

This chapter follows a natural progression: with small samples distributional knowledge has to be applied to obtain performance bounds, without which a classifier is epistemologically meaningless, and once distributional knowledge is assumed the obvious path to take is to engage in optimal error estimation and optimal classifier design. There is nothing surprising about these developments. As far back as 1925, R.A. Fisher wrote, “Little experience is sufficient to show that the traditional machinery of statistical processes is wholly unsuited to the needs of practical research. Not only does it take a cannon to shoot a sparrow, but it misses the sparrow! The elaborate mechanism built on the theory of infinitely large samples is not accurate enough for simple laboratory data. Only by systematically tackling small sample problems on their merits does it seem possible to apply accurate tests to practical data.” [41]

Before closing, let us mention some technical issues relating to the Bayesian theory. For the Gaussian model, the effective class-conditional distributions, the MMSE error estimate for linear classifiers, and the OBC can be found analytically; in particular, the posterior distribution has the same form as the prior. Although not covered herein, similar comments apply to the discrete multinomial model with Dirichlet priors. However, closed-form analytic solutions are not generally possible. For instance, in the Gaussian model with nonlinear classifiers there is no analytic expression for the MMSE error estimator and Monte Carlo methods must be employed [42]. Leaving the Gaussian model, one typically needs to employ numerical methods—for instance, Markov-chain-Monte-Carlo (MCMC) methods have been used to find the OBC with a hierarchical Poisson model [43].

A fundamental problem for any Bayesian approach is prior construction. Historically, various methods have been proposed to construct prior probabilities in different contexts [44–48]; however, these are general methodologies in that they do not target any specific type of prior information. If one tailors a prior to a specific problem in hand, then one can do better. For instance, in genomics biological knowledge in the form of regulatory pathways can be translated into feature-label knowledge for classification. This has been achieved for Gaussian network models [49], thereby significantly improving classification accuracy. The basic idea is that regulatory control constrains the feature-label distribution, in particular, the correlation between certain features in the Gaussian model. Priors are built according to the heuristic that there should be maximum uncertainty in the prior, given the regulatory constraints.

Under very general conditions, the posterior $\pi^*(\theta)$ converges to the true value of θ as the sample size goes to infinity, but this is of little interest when samples are small. What is of interest is the degree of uncertainty as it relates to classification accuracy. An obvious measure of uncertainty is the entropy of the posterior; however, what really matters is the uncertainty relating to our objective, not simply uncertainty in general. To this end, one can define the *objective cost of uncertainty*, which relates to the loss of classification performance obtained by the OBC relative to the performance should one know the true feature-label distribution [50].

In closing, we point out a critical advantage of Bayesian MMSE error estimation over classical non-Bayesian estimators. For standard data-driven error estimators, nothing can be said about the MSE of an error estimator given the sample. One can only compute the MSE as an expectation over all samples. However, in a Bayesian framework, one can compute the sample-conditioned MSE for a Bayesian error estimate, $\hat{\varepsilon}$, on a fixed classifier, ψ_n . This is equivalent to the variance of the true error conditioned on the observed sample [51],

$$\text{MSE}(\hat{\varepsilon}|S_n) = \text{Var}_{\pi^*}(\varepsilon(\theta, \psi_n)), \quad (4.46)$$

where the variance is taken with respect to $\pi^*(\theta)$. The sample-conditioned MSE converges to zero with probability 1 in both the discrete multinomial and independent covariance Gaussian models and closed-form expressions for the MSE are available [35].

References

1. T.W. Anderson, Classification by multivariate analysis. *Psychometrika* **16**(1), 31–50 (1951)
2. M.S. Esfahani, E.R. Dougherty, Effect of separate sampling on classification accuracy. *Bioinformatics* **30**(2), 242–250 (2014)
3. U.M. Braga-Neto, A. Zollanvari, E.R. Dougherty, Cross-validation under separate sampling: optimistic bias and how to correct it. *Bioinformatics* **30**(23), 3349–3355 (2014)
4. V.N. Vapnik, A. Chervonenkis, *Theory of Pattern Recognition* (Nauka, Moscow, 1974)
5. I. Shmulevich, E.R. Dougherty, *Genomic Signal Processing* (Princeton University Press, Princeton, 2007)
6. L. Devroye, L. Györfi, G. Lugosi, *A Probabilistic Theory of Pattern Recognition*, Stochastic Modelling and Applied Probability (Springer, New York, 1996)
7. C. Li, K.C.K. Soh, P. Wu, Formability of ABO₃ Perovskites. *J. Alloys Compd.* **372**(1), 40–48 (2004)
8. U.M. Braga-Neto, E.R. Dougherty, Is cross-validation valid for small-sample microarray classification? *Bioinformatics* **20**(3), 374–380 (2004)
9. B. Hanczar, J. Hua, E.R. Dougherty, Decorrelation of the true and estimated classifier errors in high-dimensional settings. *EURASIP J. Bioinform. Syst. Biol.* Article ID 38473, 12 pp (2007)
10. U. Braga-Neto, E.R. Dougherty, Exact performance of error estimators for discrete classifiers. *Pattern Recognit.* **38**(11), 1799–1814 (2005)
11. M.R. Yousefi, E.R. Dougherty, Performance reproducibility index for classification. *Bioinformatics* **28**(21), 2824–2833 (2012)
12. M.R. Yousefi, J. Hua, C. Sima, E.R. Dougherty, Reporting bias when using real data sets to analyze classification performance. *Bioinformatics* **26**(1), 68–76 (2010)
13. M.R. Yousefi, J. Hua, E.R. Dougherty, Multiple-rule bias in the comparison of classification rules. *Bioinformatics* **27**(12), 1675–1683 (2011)
14. B. Hanczar, J. Hua, C. Sima, J. Weinstein, M. Bittner, E.R. Dougherty, Small-sample precision of ROC-related estimates. *Bioinformatics* **26**, 822–830 (2010)
15. M. Hills, Allocation rules and their error rates. *J. R. Stat. Soc.: Ser. B (Stat. Methodol.)* **28**(1), 1–31 (1966)
16. D. Foley, Considerations of sample and feature size. *IEEE Trans. Inf. Theory* **18**(5), 618–626 (1972)
17. M.J. Sorum, Estimating the conditional probability of misclassification. *Technometrics* **13**, 333–343 (1971)
18. G.J. McLachlan, An asymptotic expansion of the expectation of the estimated error rate in discriminant analysis. *Aust. J. Stat.* **15**(3), 210–214 (1973)
19. M. Moran, On the expectation of errors of allocation associated with a linear discriminant function. *Biometrika* **62**(1), 141–148 (1975)
20. M. Goldstein, E. Wolf, On the problem of bias in multinomial classification. *Biometrics* **33**, 325–331 (1977)
21. A. Davison, P. Hall, On the bias and variability of bootstrap and cross-validation estimates of error rates in discrimination problems. *Biometrika* **79**, 274–284 (1992)
22. Q. Xu, J. Hua, U.M. Braga-Neto, Z. Xiong, E. Suh, E.R. Dougherty, Confidence intervals for the true classification error conditioned on the estimated error. *Technol. Cancer Res. Treat.* **5**, 579–590 (2006)
23. A. Zollanvari, U.M. Braga-Neto, E.R. Dougherty, On the sampling distribution of resubstitution and leave-one-out error estimators for linear classifiers. *Pattern Recognit.* **42**(11), 2705–2723 (2009)
24. A. Zollanvari, U.M. Braga-Neto, E.R. Dougherty, On the joint sampling distribution between the actual classification error and the resubstitution and leave-one-out error estimators for linear classifiers. *IEEE Trans. Inf. Theory* **56**(2), 784–804 (2010)
25. A. Zollanvari, U.M. Braga-Neto, E.R. Dougherty, Exact representation of the second-order moments for resubstitution and leave-one-out error estimation for linear discriminant analysis in the univariate heteroskedastic Gaussian model. *Pattern Recognit.* **45**(2), 908–917 (2012)

26. A. Zollanvari, U.M. Braga-Neto, E.R. Dougherty, Analytic study of performance of error estimators for linear discriminant analysis. *IEEE Trans. Signal Process.* **59**(9), 4238–4255 (2011)
27. F. Wyman, D. Young, D. Turner, A comparison of asymptotic error rate expansions for the sample linear discriminant function. *Pattern Recognit.* **23**, 775–783 (1990)
28. V. Pikelis, Comparison of methods of computing the expected classification errors. *Autom. Remote Control* **5**, 59–63 (1976)
29. E.R. Dougherty, A. Zollanvari, U.M. Braga-Neto, The illusion of distribution-free small-sample classification in genomics. *Curr. Genomics* **12**(5), 333–341 (2011)
30. B. Efron, Estimating the error rate of a prediction rule: improvement on cross-validation. *J. Am. Stat. Assoc.* **78**(382), 316–331 (1983)
31. T. Vu, C. Sima, U.M. Braga-Neto, E.R. Dougherty, Unbiased bootstrap error estimation for linear discriminant analysis. *EURASIP J. Bioinform. Syst. Biol.* **2014**(1), 15 (2014)
32. C. Sima, E.R. Dougherty, Optimal convex error estimators for classification. *Pattern Recognit.* **39**, 1763–1780 (2006)
33. L.A. Dalton, E.R. Dougherty, Bayesian minimum mean-square error estimation for classification error-Part I: Definition and the Bayesian MMSE error estimator for discrete classification. *IEEE Trans. Signal Process.* **59**(1), 115–129 (2011)
34. L.A. Dalton, E.R. Dougherty, Bayesian minimum mean-square error estimation for classification error-Part II: The Bayesian MMSE error estimator for linear classification of Gaussian distributions. *IEEE Trans. Signal Process.* **59**(1), 130–144 (2011)
35. L.A. Dalton, E.R. Dougherty, Exact sample conditioned MSE performance of the Bayesian MMSE estimator for classification error-Part II: Consistency and performance analysis. *IEEE Trans. Signal Process.* **60**(5), 2588–2603 (2012)
36. U. Braga-Neto, E. Dougherty, Bolstered error estimation. *Pattern Recognit.* **37**(6), 1267–1281 (2004)
37. L.A. Dalton, E.R. Dougherty, Optimal classifiers with minimum expected error within a Bayesian framework-Part I: Discrete and Gaussian models. *Pattern Recognit.* **46**(5), 1301–1314 (2013)
38. M.H. DeGroot, *Optimal Statistical Decisions* (McGraw-Hill, New York, 1970)
39. H. Raiffa, R. Schlaifer, *Appl. Stat. Decis. Theory* (MIT Press, Cambridge, 1961)
40. E.R. Dougherty, J. Hua, Z. Xiong, Y. Chen, Optimal robust classifiers. *Pattern Recognit.* **38**(10), 1520–1532 (2005)
41. R.A. Fisher, *Statistical Methods for Research Workers* (Oliver and Boyd, Edinburgh, 1925)
42. L.A. Dalton, E.R. Dougherty, Application of the Bayesian MMSE estimator for classification error to gene expression microarray data. *Bioinformatics* **27**(13), 1822–1831 (2011)
43. J.M. Knight, I. Ivanov, E.R. Dougherty, MCMC implementation of the optimal Bayesian classifier for non-Gaussian models: Model-based RNA-Seq classification. *BMC Bioinform.* **15**(1), 401 (2014)
44. J.M. Bernardo, Reference posterior distributions for Bayesian inference. *J. R. Stat. Soc. Ser. B (Methodol.)*, 113–147 (1979)
45. J. Rissanen, A universal prior for integers and estimation by minimum description length. *Ann. Stat.* 416–431 (1983)
46. J.C. Spall, S.D. Hill, Least-informative Bayesian prior distributions for finite samples based on information theory. *IEEE Trans. Autom. Control* **35**(5), 580–583 (1990)
47. J.O. Berger, J.M. Bernardo, On the development of reference priors. *Bayesian Stat.* **4**(4), 35–60 (1992)
48. R.E. Kass, L. Wasserman, The selection of prior distributions by formal rules. *J. Am. Stat. Assoc.* **91**(435), 1343–1370 (1996)
49. M.S. Esfahani, E. Dougherty, Incorporation of biological pathway knowledge in the construction of priors for optimal Bayesian classification. *IEEE/ACM Trans. Comput. Biol. Bioinform.* **11**(1), 202–218 (2014)
50. B.-J. Yoon, X. Qian, E.R. Dougherty, Quantifying the objective cost of uncertainty in complex dynamical systems. *Signal Process., IEEE Trans.* **61**(9), 2256–2266 (2013)

51. L.A. Dalton, E.R. Dougherty, Exact sample conditioned MSE performance of the Bayesian MMSE estimator for classification error-Part I: Representation. *IEEE Trans. Signal Process.* **60**(5), 2575–2587 (2012)

Chapter 5

Data Visualization and Structure Identification

J.E. Gubernatis

Abstract For three datasets, all dealing with materials with ABO_3 chemistries, the two data visualizations algorithms of Tsafrir et al. [Bioinformatics 21, 2301 (2005)] were studied and applied. These algorithms permute the distance matrix associated with the data in a way to unveil structure in one case by keeping large-distanced information afar or in the other case by keeping small-distanced information near. Modifications to their proposed numerical implementations were made to enhance effectiveness. The two algorithms were used both in space of the materials and the features, looking for groupings of features and materials. In general, for the datasets considered, when visualized, the features tended to show more distinctive structure (clustering) than the materials. For enhanced grouping of materials, the initial studies point to the importance of feature selection.

5.1 Introduction

The pre-emptive focus of Materials Informatics is gathering materials data and extracting from them sign-posts for candidate materials with enhanced properties. We studied three datasets, previously used in materials informatics studies [1] that had similar objectives, literally asking, What does the data look like? To assist in visualizing the data, we used the recent work by Tsafrir et al. [2] in bioinformatics that presented two seemingly simple algorithms to visualize the data in a way that also revealed structure in them, that is, correlations among the materials and features being visualized. Their algorithms reorder the data by permuting the rows and columns of a distance matrix constructed from the data matrix. The permutations minimize a cost function that favors placing data close together when the distances between them are small. With addenda, the algorithms become clustering methods which do not *a priori* assume the number of clusters [3, 4].

J.E. Gubernatis (✉)
Theoretical Division, Los Alamos National Laboratory,
Los Alamos, NM 87545, USA
e-mail: jg@lanl.gov

© Springer International Publishing Switzerland 2016
T. Lookman et al. (eds.), *Information Science for Materials
Discovery and Design*, Springer Series in Materials Science 225,
DOI 10.1007/978-3-319-23871-5_5

The distance matrix D is usually formed after normalizing the data. Normalization is necessary because the different features have different units and the values of different features vary by orders of magnitude. The rows and columns of the data set tables are viewed as an $M \times F$ matrix of materials and features. Each column of features is regarded as a M -vector which is normalized by first computing the mean value of its components and subtracting the mean from each component. Next, the standard deviation for each mean-centered column is computed and each component is divided by it. After normalization, the units have disappeared and the numerical values in each column of data have the same center (mean of zero) and the same range (variance of unity). The result is a new $M \times F$ -dimensional Data matrix, $\text{Data} = (\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_F)$.

Various definitions of a distance matrix exist. A Euclidean distance is the only type considered in this report. We computed a Euclidean distance matrix from the normalized data matrix in two ways. One way is what we call computing distances in Materials Space. In this space, the components of the $F \times F$ distance matrix are

$$D_{ij} = \sqrt{(\mathbf{f}_i - \mathbf{f}_j) \cdot (\mathbf{f}_i - \mathbf{f}_j)}. \quad (5.1)$$

Here, each feature i is regarded as a M -dimensional vector \mathbf{f}_i of materials. The second way is what we call computing distances in Features Space. Here, we work with the rows of the normalized Data matrix instead of the columns: $\text{Data} = (\mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_M)^T$, where each material i is a F -dimensional vector \mathbf{m}_i of features. In this space, the components of the $M \times M$ distance matrix are

$$D_{ij} = \sqrt{(\mathbf{m}_i - \mathbf{m}_j) \cdot (\mathbf{m}_i - \mathbf{m}_j)}. \quad (5.2)$$

We illustrate these vectors schematically in Fig. 5.1.

We applied the Tsafriar et al. algorithms to the data in both spaces. These algorithms reorder the data so the points in the respective space are closer together than they are in the tables. In Materials Space, they group together features; in Features Space, they group together materials. Is there anything to be gained by viewing the data in these two different ways?

5.2 Theory

The Tsafriar et al. algorithms find a permutation matrix P that minimizes

$$F(P) = \text{Tr}(\text{PDP}^T \mathbf{W}). \quad (5.3)$$

A permutation matrix is a matrix whose elements are all zero except for one element in each row and column that is unity. The permutation matrix P is usually represented as an integer array IP where $IP(i) = j$ which gives the non-zero column j for the i th row.

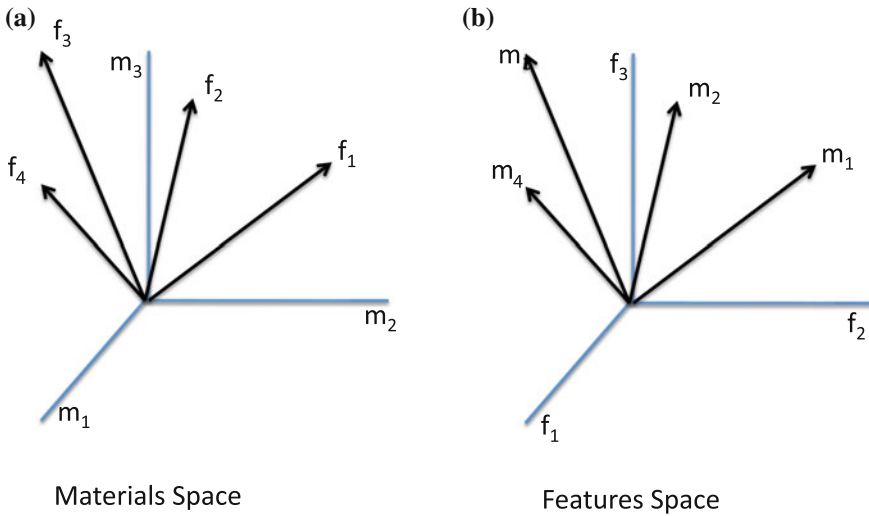


Fig. 5.1 Different spaces in which to represent the data. **a** Space of the materials: The materials point to a feature. **b** The space of the features: The features point to a material

Different W matrices define the two different algorithms [2]. The first algorithm is called “Side-to-Side (STS)”. Here the matrix elements of W are

$$w_{ij} = x_i x_j \quad (5.4)$$

where the x_i are components of any vector \mathbf{X} that satisfy $x_i < x_j$. The W created from this vector pushes apart data separated by large distances. For the results reported here, we used the choice of Tsafirir et al. $\mathbf{X} = (-N/2, -N/2 + 1, \dots, N/2)$, although we found that multiplying this vector by a factor of two, three, or four often gave more satisfying results. Shifting the vector so all its components are positive, typically degraded the results. Their second algorithm is called “Neighborhood (NBRHD)”. Here the matrix elements of W are

$$w_{ij} = \exp(-|i - j|/\sigma^2). \quad (5.5)$$

This choice pulls together data separated by small distances.

Both minimizations are NP -hard problems, meaning obtaining a good (non-unique) local minimum the best that one can expect. As the order of the matrix increases, many good minima can exist. Side-to-Side belongs to a class of problems called quadratic assignment problems; Neighborhood, linear assignment problems. Tsafirir et al. give a numerical procedure to do the minimization for each choice of W . The overall advice was to restart each procedure multiple times, each from a different point, and keep the lowest value. For Neighborhood, they note that the parameter σ could be used as an “annealing” parameter: get a solution for a small value, use that solution for a larger value, and then repeat these two steps ten times or so. Because the

W of Side-to-Side factorizes, the algorithm they proposed for it scales as the square of the order of the matrices; for Neighborhood, their proposed procedure scales as the cube. We found the suggested numerical procedures of Tsafirir et al. gave mixed performance for the datasets under study. For each W , we instead used Algorithm 1.

```

Initialize  $t = 0$ ,  $P^{t-1} = 0$ ,  $P^t = I$ , and  $W^t = W$ .
while  $P^t \neq P^{t-1}$  do
   $t \leftarrow t + 1$ 
  Solve for  $P^t = \arg \min_P \text{Tr} PDW^{t-1}$ 
   $W^{t+1} = [P^t]^T$ 
end while
 $D \leftarrow P^t D [P^t]^T$ 

```

Algorithm 1: Minimization Procedure

This algorithm modestly differs from that of Tsafirir et al.’s Neighborhood algorithms in the following respects: First, we are using it for both the Side-to-Side and Neighborhood W . In other words, we are treating each problem as if it were a linear assignment problem. The most important difference is our step “Solve ...”: Instead of using their suggested procedures to find the permutation matrix, we are using the Hungarian algorithm [5], a standard method for assignment problems. Our stopping criterion also differs: Instead of using $|F(P^{t+1}) - F(P^t)| < \varepsilon$, or something similar, we are iterating until the permutation matrix ceases to change. This criterion in general forced more iteration steps and produced a lower value for the minimum. Restarts for the new algorithm for the cases considered seemed to gain little.

We also tried minimizing by using a greedy Monte Carlo optimization procedure (making many random permutations many times and keeping the best) and a rudimentary simulated annealing optimization. In general, the results from the Hungarian algorithm had the tightest structure in the visualization.

5.3 Results

We now report select results. All were obtained with Algorithm 1. The three datasets studied were used by Balachandran et al. [1]. We call the datasets `piezo.dat`, `pls.dat`, and `tree.dat`. All the materials in the dataset `piezo.dat` are known ferroelectrics. Balachandran et al. used this data in a feature-reduction principal component analysis. The `pls.dat` data was used for a partial least-squares (PLS) analysis of the piezoelectric data, after a further somewhat *ad hoc* feature reduction, to generate an analytic expression for the Curie temperature which they then used to predicted possible Curie temperatures for perovskite chemistries not yet known to exist. All the materials in `tree.dat` had the ABO_3 chemistry but not all had a perovskite crystal structure. This data was used by Balachandran et al. to construct a binary decision tree giving rules for when a perovskite crystal structure should exist.

5.3.1 The Piezo Data

This dataset has 22 materials and 31 features. In Fig. 5.2 are the distance matrices before the re-ordering of the data, and in Figs. 5.3 and 5.4 are these matrices after the re-ordering. Viewing Figs. 5.2 and 5.3 or Figs. 5.2 and 5.4 together, one can see small groupings of materials and features. We can relate these groupings to materials or features having the same A or B atoms. Otherwise, larger clumping of materials or features is not prevalent. The initial distance matrix in Fig. 5.2 shows only minor block structure near the diagonal (zero distance), with a bit more in Features Space (left) than in Materials Space (right). The Side-to-Side ordering produced more distinct clumping in Features Space than in Materials Space. The Neighborhood ordering produced very distinct clumping in Materials Space.




Fig. 5.2 Distance matrices for **piezo.dat** before reordering




Fig. 5.3 Distance matrices after re-ordering with Side-to-Side

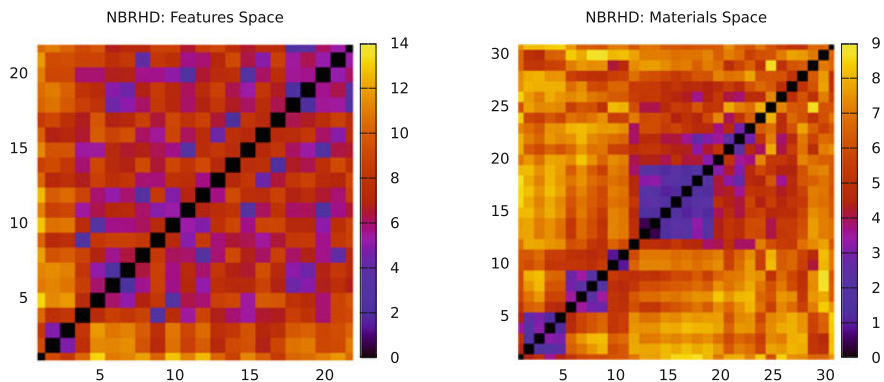


Fig. 5.4 Distance matrices after re-ordering with Neighborhood. $\sigma = 10$




Fig. 5.5 Distance matrices for `pls.dat` before reordering

5.3.2 The Pls Data

This dataset has 21 materials and 7 features. Figure 5.5 is the distance matrices before re-ordering, and Figs. 5.6 and 5.7 are these matrices after re-ordering. Viewing Figs. 5.5 and 5.6 or Figs. 5.5 and 5.7 together, one sees less clumping than seen in Figs. 5.3 or 5.4. Presumably this is caused by simply having fewer features. The initial distance matrices in Fig. 5.5 show little block structure along the diagonal. The Side-to-Side ordering produced distinct clumping in the Features and Materials Spaces. Neighborhood ordering produced virtually identical clumpings.

5.3.3 The Tree Data

Here, there are 355 materials and 13 features. Figure 5.8 is the distance matrices before re-ordering, and Figs. 5.9 and 5.10 are these matrices after re-ordering. As for

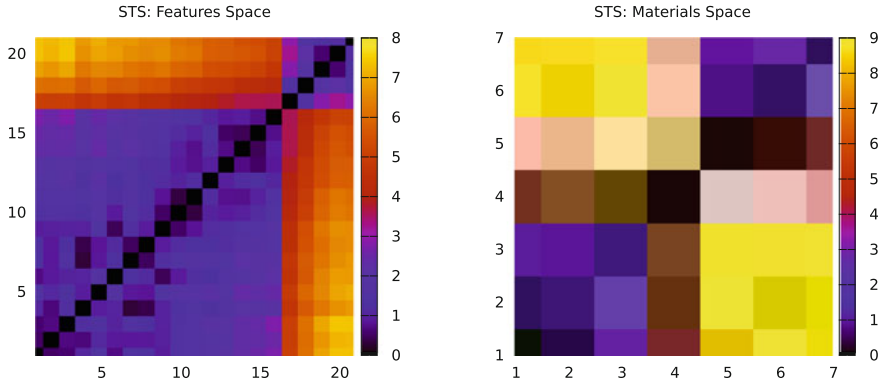


Fig. 5.6 Distance matrices after re-ordering with Side-to-Side




Fig. 5.7 Distance matrices after re-ordering with Neighborhood. $\sigma = 10$

the `pls.dat`, the number of features is smaller than the number of materials. Here, their number is much smaller. In Materials Space, Fig. 5.8 shows some clear block structure along the diagonal. The Side-to-Side ordering tightened the clumping a bit in Materials Space, but it is Neighborhood ordering that produced the most distinct clumping in both spaces.

5.4 Concluding Remarks

This initial study suggests several recommendations and items for future study. First, our findings are consistent those of Tsatfir et al. that the Neighborhood method is generally the most revealing algorithm.

Understudied to date is the potential for using σ to enhance the results. Figures 5.11, 5.12 and 5.13 show a brief study of what happens if the results in




Fig. 5.8 Distance matrices for `tree.dat` before reordering




Fig. 5.9 Distance matrices after re-ordering with Side-to-Side




Fig. 5.10 Distance matrices after re-ordering with Neighborhood. $\sigma = 10$




Fig. 5.11 Distance matrices after re-ordering with Neighborhood. $\sigma = 100$




Fig. 5.12 Distance matrices after re-ordering with Neighborhood. $\sigma = 200$




Fig. 5.13 Distance matrices after re-ordering with Neighborhood. $\sigma = 300$

Fig. 5.10 were extended from $\sigma = 10$ to $\sigma = 100, 200$, and 300 . The changes are mainly exposing more structure in Features Space.¹ In general, finding materials clumping in Features Space for the `tree.dat` was the reason various modifications of the Tsafrir et al. algorithms were attempted and several Monte Carlo optimization methods were explored. Instead of using σ as an annealing parameter, suggested by them, one could consider using it as a tempering parameter: Parallel tempering is generally a more effective Monte Carlo minimization scheme than simulated or quantum annealing. More effective still are the recently proposed partial and infinite swapping methods [6, 7]. The upfront question first needing an answer is, How good of a solution is needed for the intended applications? At this writing, the answer to this question is unestablished.

The parameter σ likely has as more immediate use in setting length scales. The differences in Features Space between Figs. 5.10 and 5.11 illustrate this. Years ago, the connection between a data clustering algorithm and a first-order phase transition was noted [8]. Several physics-based algorithms have exploited this fact to develop successful data clustering algorithms [9, 10]. The algorithms of Tsafrir et al. in a sense, are part of this alternative perspective. In a first-order phase transition, clustering (strong correlations) among interacting particles occurs at various length scales that are the consequences of the distances over which the interaction between particles are attractive or repulsive. The correlations become stronger as the temperature is lowered towards the transition temperature. σ is an analog to the temperature: Varying it here varies a length length scale in the matrix W . Distinguishing the physics-based algorithms from standard machine learning algorithms is the presence of several length scales as opposed to none. Curiously, a seminal paper [11] on the k -means clustering algorithm, one of the most popular machine learning clustering algorithms, proposed a “grouping” algorithm that had two length scales, one for refinement and one for coarsening. Refinement increases the “attraction” of data to a particular mean and coarsening provides a “repulsion” from it. This suggestion captures the “physics” of a clustering method and is an algorithm needing implementation.

We remark that the classification and clustering problems are connected. For classification problems using algorithms that have length scales in them is likely to be highly desirable. Finding the effective scales for either type of problem for the given data is likely more important than trying a suite of machine learning algorithms to find the one that is most effective or a few that are consistent.

A variety of choices for the distance matrix exist. It appears that for whichever one is used, using it with a large number of features, at least with the current choices, has the potential of “washing out” the few features that are most important. For example, the datasets studied all had the tolerance factor as one of the features. By itself, the tolerance factor is traditionally used to separate perovskites from non-perovskites and ferroelectric perovskites from non-ferroelectric perovskites. For the analyses performed here, this feature seemed to have no assertive role.

¹These differences are likely more evident if the **pdf** file of this report is viewed on a monitor with decent resolution than from a printed version of the report.

As part of the feature selection issue, we suggest the following: Clustering and classification methods start with data normalized relative to some fictitious material that has the average features of the given dataset. The majority, not necessarily the optimal, determines the average even though we are seeking materials that lie outside the range of the average. Clustering takes the additional steps of scaling the data to homogenize the range. The ideal perovskite is SrTiO_3 in the sense that it has nearly the ideal cubic crystal structure, but it is not a ferroelectric. PbTiO_3 , which has less than the ideal crystal structure, in another sense is the ideal, except for having Pb, because it is an excellent ferroelectric. It seems that in contrast to current machine learning clustering or classification schemes that define things relative to some average, we would want schemes that define things close to PbTiO_3 but in a direction that points away from SrTiO_3 . It is unclear whether such schemes exist. On the other hand, within the existing visualization/clustering scheme, one can at least start with the data centered relative to PbTiO_3 and then query the results for those cases that are also far from SrTiO_3 .

Generally, it is Features Space in which we want to work as we want to associate new materials with experimentally accessible features. Materials Space reveals features that are close. In some cases working in this space might reveal redundant features, that is, it might provide a means for feature reduction.

This work was supported by the Department of Energy's Laboratory Directed Research and Development Program.

References

1. P.V. Balachandran, S.R. Broderick, K. Rajan, Proc. R. Soc. A (2010). doi:[10.1098/rspa.2010.0543](https://doi.org/10.1098/rspa.2010.0543)
2. D. Tsafrir et al., Bioinformatics **21**, 2301 (2005)
3. D. Filippova, A. Gagni, C. Kingsford, BMC Bioinformatics **13**, 276 (2012)
4. M. Neuditschko, M.S. Khatkar, H.W. Raadsma, PLOS ONE **7**, e48375 (2012)
5. http://en.wikipedia.org/wiki/Hungarian_algorithm
6. N. Plattner et al., J. Chem. Phys. **135**, 134111 (2011)
7. P. Dupuis et al., Multiscale Model Simul. **10**, 986 (2012)
8. K. Rose et al., Phys. Rev. Lett. **65**, 945 (1990)
9. M. Blatt et al., Phys. Rev. Lett. **76**, 3251 (1996)
10. P. Ronhovede, Z. Nussinov, Phys. Rev. E **81**, 046114 (2010)
11. J. B. McQueen, Some methods for classification and analysis of multivariate data. in *Proceedings of 5th Berkeley Symposium on Mathematical Statistics and Probability*, (University of California Press, Berkeley, 1967). p. 281

Chapter 6

Inference of Hidden Structures in Complex Physical Systems by Multi-scale Clustering

Z. Nussinov, P. Ronhovde, Dandan Hu, S. Chakrabarty,
Bo Sun, Nicholas A. Mauro and Kisor K. Sahu

Abstract We survey the application of a relatively new branch of statistical physics—“community detection”—to data mining. In particular, we focus on the diagnosis of materials and automated image segmentation. Community detection describes the quest of partitioning a complex system involving many elements into optimally decoupled subsets or communities of such elements. We review a multiresolution variant which is used to ascertain structures at different spatial and temporal scales. Significant patterns are obtained by examining the correlations between different independent solvers. Similar to other combinatorial optimization problems in the NP

Z. Nussinov (✉) · B. Sun · D. Hu
Washington University in St. Louis, St. Louis, MO 63130, USA
e-mail: zohar@wuphys.wustl.edu

B. Sun
e-mail: bosun@wustl.edu

D. Hu
e-mail: dan1226@gmail.com

Z. Nussinov
Department of Condensed Matter Physics, Weizmann Institute of Science,
76100 Rehovot, Israel

P. Ronhovde
Findlay University, Findlay, OH 45840, USA
e-mail: ronhovde@findlay.edu

S. Chakrabarty
Department of Physics, Indian Institute of Science, Bangalore 560012, India
e-mail: schakrab@go.wustl.edu

N.A. Mauro
North Central College, Naperville, IL 60540, USA
e-mail: Nicholas.mauro@gmail.com

K.K. Sahu
School of Minerals, Metallurgical and Materials Engineering,
Indian Institute of Technology, Bhubaneswar 751007, India
e-mail: kis.sahu@gmail.com

complexity class, community detection exhibits several phases. Typically, illuminating orders are revealed by choosing parameters that lead to extremal information theory correlations.

6.1 The General Problem

A basic question that we wish to discuss in this work is whether machine learning and data mining tools may be applied to the analysis of material properties. Specifically, we will review initial efforts to detect, via statistical mechanics and the tools of information science and network analysis, pertinent structures on all scales in general complex systems. We will describe mapping atomic and other configurations onto graphs. As we will explain, patterns found in these graphs via statistical physics methods may inform us about the structure of the investigated materials. These structures can appear on multiple spatial and temporal scales. In comparison to standard procedures, the advantage of such an approach may be significant.

There are numerous classes of complex systems. One prototypical variety is that of glass forming liquids. “Glasses” have been analyzed with disparate tools [1–16]. Although they have been known for millennia, structural glasses still remain ill understood. It is just over eighty years since the publication of one of the most famous papers concerning the structure of glasses [2]. Much has been learned since the early days of hand-built plastic models and drawings, yet basic questions persist.

Amorphous systems such as glasses strongly contrast with idealized simple solids. In simple crystals, the structure of an atomic unit cell is replicated to span the entire system. Long before scattering and tunneling technologies, prominent figures such as Robert Hooke, Christiaan Huygens, and their contemporaries in the 17th century proposed the existence of sharp facets in single crystals results from recurrent fundamental unit cell configurations. The many years since have seen numerous breakthroughs (including the advent of quantum mechanics and atomic physics) and witnessed a remarkable understanding as to how the quintessential simple periodic structure of crystals accounts for many of their properties. However, while simple solids form a fundamental pillar of current technology (e.g., the transistor whose invention was made possible by an understanding of the electronic properties of nicely ordered periodic crystals and chemical substitution therein), there are many other complex systems whose understanding is extremely important yet still lacking. The discovery of salient features of these materials across all scales is important for both applied and basic science. The recognized significance of this problem engendered the *Materials Genome Initiative* [17]—a broad effort to develop infrastructure for accelerating materials innovation.


This work discusses a path towards solving this problem in complex amorphous materials. The framework that we will principally suggest is that of *multi scale community detection*. This approach does not invoke assumptions as to which system properties are important and construct resulting minimal toy models based on the assumptions. The insightful guess-work that is typically required to describe complex

materials is, in the work that we review, replaced by a computerized variant of the *wisdom of the crowds* phenomena [18]. The key concepts underlying this approach may be applied to general hard problems beyond those concerning the structure of materials or even general data mining. In the next section, we review an “Information theoretic ensemble minimization” method that may be suited for such tasks.


6.2 Ensemble Minimization

Before delving into complex material and network analysis, we first discuss a general strategy for solving hard problems. The concept underlying this approach is perhaps best conveyed by a simple cartoon such as that sketched in Fig. 6.1a. In this illustration, each sphere corresponds to an individual solver (or “replica”) that explores an energy landscape. On its own, each such sphere might get stuck in a local energy minimum. The collective *ensemble* of solvers may, however, thwart such situations more readily as compared to the same single solver algorithm [21]. In Fig. 6.1b, the individual solvers not only roam the energy landscape but also interact amongst themselves as schematically denoted by springs. If a single solver gets stuck in a

Fig. 6.1 The spheres in panel (a) of the figure depict solvers (or “replicas”) independently navigating the energy landscape defined by (6.2). Strong correlations among the replicas indicate stable, well-defined partition. We evaluate agreement among all replica pairs using the information correlations (Sect. 6.4). In panel (b), interactions between the replicas assist the ensemble in finding optimal low energy states



(a) Independent solvers (or “replicas”) on a schematic energy landscape.



(b) Coupled solvers.

false minimum, the other solvers may “pull it out” and explore broader regions of the energy landscape.

This collective evolution of individual solvers is quite natural and has appeared in different guises across many fields. In anthropological contexts, this basic principle is known as “wisdom of the crowds” [18]. That is, the crowd or ensemble of individuals might do far better than a single solver. Unlike ensemble related approaches such as swarm intelligence [22] or genetic [23] algorithms, relevant problems in our context do not focus exclusively on minimizing a given energy function. Rather, we will try to maximize information theory correlations [the effect of the springs in Fig. 6.1b] while simultaneously minimizing a cost function [20]. If all (or many) solvers agree on a particular candidate solution then that solution may naturally arise in many instances and may be of the high importance regardless of whether or not it is the absolute minimum of the energy. In the physical problems that we will consider—that of finding natural structures in materials—these considerations are pertinent.

The above discussion is admittedly abstract and may, in principle, pertain to any general problem. We next briefly explain the basic mathematical framework—the *community detection* problem—in which we will later couch the material structure detection endeavor.

6.3 Community Detection and Data Mining

Community detection pertains to the quest of partitioning a given graph or network into its optimally decoupled subgraphs (or so-called communities), e.g., [24–37]. As the reader may anticipate, given the omnipresence of networks and the generality of this task, this problem appears in disparate arenas including biological systems, computer science, homeland security, and countless others. In what follows, we introduce some of the key elements of community detection. The graphs of interest will be composed of nodes where a *node* is a fundamental element of an abstracted graph. An *edge* in the graph is a defined relationship between two nodes. Edges may be weighted or unweighted where the unweighted case is the one most commonly examined. In our applications, we will need to assign weights to the edges in the graph as we will describe. Similarly, in general applications, edges may be either symmetric or directed.

Now we come to a basic ingredient of community detection. A community corresponds to a subset of nodes that are more cohesively linked (or densely connected for unweighted edges) within their own community than they are to other communities. The above definition might seem a bit loose. Indeed, there are numerous formulations of community detection in the literature. As intuitively one may expect, most of these do, more or less, the same thing. When clear community detection solutions exist, all algorithms quantify the structure of large complex networks in terms of the smaller number of the natural cohesive components. Rather general data structures may be cast in terms of abstract networks. Thus, the community detection problem and other network analysis methods can have direct implications across multiple fields. Indeed, we will elaborate how this occurs for image segmentation and material analysis.




Fig. 6.2 A small network partition where individual communities are represented by different node shapes and colors. “Friendly” or “cooperative” relations are depicted by *solid, black lines*. These are modeled as ferromagnetic interactions in (6.2). “Missing” or “undefined” relations work to break up well-defined communities, so they are modeled with anti-ferromagnetic interactions, meaning they are repulsive in terms of their energy contributions. The physical energy model trivially extends to more general relations including weighted and adversarial relations (not depicted here)

In what follows we will briefly review the rudiments of an “Absolute Potts Model” method for community detection [19] that avoids a “resolution limit” that an insightful earlier Potts model [38] exhibited. To cast things generally, we make a simple observation underlying the “Potts” characterization. Any partition of the numbered nodes $i = 1, 2, 3, \dots, N$ into q different communities (the ultimate objective of any community detection algorithm) is an assignment $i \rightarrow \sigma_i$ where the integer $1 \leq \sigma_i \leq q$ denotes the community number to which node i belongs. With a characterization $\{\sigma_i\}$ in hand, we next construct an energy functional.

To illustrate the basic premise, we first consider an unweighted graph—one in which the link strength A_{ij} between the two nodes i and j is $A_{ij} = 1$ if an edge is present between the two nodes and $A_{ij} = 0$ if there is no link. As Fig. 6.2 demonstrates, for each pair of nodes there are four principal cases to consider. That is, either (i) the two nodes belong to the same community and have an “attraction” between them (i.e., $A_{ij} = 1$), (ii) two nodes in the same community can have a missing link between them ($A_{ij} = 0$), (iii) the two nodes may belong to different communities yet nevertheless exhibit cohesion between themselves ($A_{ij} = 1$), or (iv) nodes i and j may belong to different communities and have no edge connecting them ($A_{ij} = 0$). Situations (i) and (iv) agree with the intuitive expectation that nodes in the same community should be connected to one another while those in different communities ought to be disjoint. We may take these four possibilities as the foundation of an energy function. That is, any given pair of nodes may be examined to see which of these categories it belongs to. Thus, a contending cost function is given by the Potts model Hamiltonian

$$H = -\frac{1}{2} \sum_{i \neq j} [A_{ij} \delta(\sigma_i, \sigma_j) + \gamma(1 - A_{ij})(1 - \delta(\sigma_i, \sigma_j))]. \quad (6.1)$$

In (6.1), $\delta(\sigma_i, \sigma_j)$ is a Kroncker delta (i.e., $\delta(\sigma_i = \sigma_j) = 1$, $\delta(\sigma_i \neq \sigma_j) = 0$) and γ is a “resolution parameter” that will play a notable role in our analysis. Before turning to the origin of the name of this parameter, we observe that, subtracting an innocuous additive constant, (6.1) is trivially

$$H = -\frac{1}{2} \sum_{i \neq j} [A_{ij} - \gamma(1 - A_{ij})] \delta(\sigma_i, \sigma_j). \quad (6.2)$$

As (6.2) makes clear, by virtue of the Kronecker delta $\delta(\sigma_i, \sigma_j)$, the sum is *local*—i.e., the sum only includes intra-community node pairs. The Hamiltonian of (6.2) may be minimized by a host of methods. In practice, when the solution of the problem is easy to find, nearly all viable approaches will yield the same answer. Amongst many others, two approaches are afforded by spectral methods [in which the discrete Potts model spins are effectively replaced by continuous spherical model (or large n) spins] and a conceptually more primitive steepest descent type approach.

A simple incarnation of the relatively successful greedy algorithm [19, 20] that extends certain ideas introduced in [29] is given by the following steps: (a) Initially, each node forms its own community [i.e., if there are N (numbered) nodes then there will be $q = N$ communities].

(b) A node (whose number is i_1) is chosen stochastically and then another edge sharing node i' is picked at random. (c) If it is energetically profitable to move the node i' together into the group formed by i_1 then this is done (otherwise community assignments are unchanged). (d) Yet another node i_2 is next chosen and once again it is asked whether moving yet another node into the community of i_2 lowers the energy. As earlier mentioned, if this change lowers the energy of (6.2), the nodes will be merged. Otherwise no change will be made. (e) In this manner, we cycle through each of the N nodes and repeat as necessary. (f) The process stops and a candidate partition is found once all further possible mergers do not lower the energy further. As the reader can appreciate, such a simple algorithm lowers the energy until the system becomes trapped in a local minimum. To improve the accuracy (i.e., further lower the energy of candidate solutions), one may repeat the above steps a finite number of times for a finite number of *trials*—i.e., repeat the above when vertices i_1, i_2, \dots, i_N are chosen in a different random order to see if a lower energy solution may result.

For the wide range of examined problems, the number of trials for each replica of the system is typically on the order of ten or smaller. When approaching the “hard phase” (to be discussed in Sect. 6.6) with multiple false minima, an increase in the number of trials may likely further increase the accuracy (this rise in the accuracy was termed the “computational susceptibility” in [20, 61]). Typically, elsewhere the improvement in the precision due to a further increase in the number of trials is nearly nonexistent (see, e.g., Fig. 13 in [20]). Further embellishments of the bare algorithm outlined above, include the acceptance of zero energy moves and other refinements [19]. Other illuminating greedy type approaches for the inference of community structure have been advanced, e.g., [39].

6.4 Multi-scale Community Detection

We now turn to “multi-scale” community detection, e.g., [20, 40–45]. In certain notable approaches, e.g., [45], detection of scale is performed without the resolution parameter but rather by examining the effects of thermal fluctuations in a pure ferromagnetic system (one sans the antiferromagnetic interaction present in the second term of (6.2)), and other considerations elsewhere. In what follows, we will build on the ideas introduced in Sect. 6.3 that lead to an accurate determination of structure on diverse pertinent scales. To understand the physical content of the resolution parameter (and the origin of its name) in (6.2), we consider several trivial limits. First, we focus on the case of $\gamma = 0$. In such a situation, the energy of (6.2) is minimized when all nodes belong to a single community. This is the lowest energy solution since each intra-community link lowers the energy [the first term of (6.2)], but there is no energy penalty from any missing links between nodes in the same community since the second term in (6.2) is trivially zero. Thus, in order to maximize the number of internal links it is profitable to assign all nodes to the same community. In the diametrically opposite limit—that of $\gamma \rightarrow \infty$, the energy penalty diverges unless every pair of nodes belonging to the same community share a link. Thus, in this limit, the lowest energy states are those in which the system fragments into (typically) a large number of communities where each node is connected to all other nodes in its community. That is, the communities are “perfect cliques.” As γ is monotonically increased from zero, the ground states of (6.2) lead to communities that veer from the extreme global case ($\gamma = 0$) to the limit of many disparate densely internally connected local communities ($\gamma \rightarrow \infty$). Putting all of the pieces together, the reader can see why γ is inherently related to the intra-community edge density and thus is indeed a “resolution parameter”.

At this stage, it is not yet clear which values γ should be assigned in order to lead to the most physically pertinent solutions. The non-uniqueness of γ is, actually, a virtue of the Potts model based approach of (6.2). That is, in general, there may be several relevant resolution scales that lead to different insightful candidate low energy partitions of this Hamiltonian. This is the situation which is schematically depicted in Fig. 6.3 for a synthetic system that exhibits a hierarchical structure. In such cases as γ is increased, the minima of (6.2) unveil different resolutions in the hierarchy. In practice, the multi-resolution community-detection method [20] systematically infers the pertinent scale(s) by information-theory-based correlations [46–49] between different independent solvers (or “replicas”, as discussed in Sect. 6.2) of the same community detection problem. In most studied systems, the number of replicas used is $s \leq 12$. As alluded to in Sect. 6.3, the lowest energy solution amongst a fixed number of trials is taken for each of the individual replicas. If these solvers (i.e., the replicas) strongly concur with each other about local or global features of the solution [20], then these aspects are likely to be correct. Such an agreement between solvers is manifest in the information correlations. Information theory extrema [50–52] then provide *all relevant system scales*.




Fig. 6.3 A partition of a synthetic network with 256 nodes having three resolution levels [19]. The random edge density (fraction of edges connecting pairs of points in different communities) is 10% on the global scale. At increasing resolution there are five groups with an inter-community edge density of 30%. At the highest resolution, these five groups are further split into small sub clusters (16 in total) each having an internal edge density of 90%. As described in Sect. 6.4, a multi-resolution algorithm may identify different categories of partitions in hierarchical systems. See Fig. 6.4 for a demonstration of how the multiresolution algorithm accurately isolates both levels of the hierarchy

Figure 6.4 shows the results of our analysis as the resolution parameter γ is varied for the synthetic system of Fig. 6.3. Plotted are three information theory correlations between replicas—the average inter-replica variation of information (VI), the mutual information (I), the normalized mutual information (NMI), the total number of communities (q) found for different values of γ , and the Shannon entropy (H) averaged over different replicas. Transitions between viable solutions are evident as jumps in the number of communities q and, most notably, as transitions between crisp information theory measure plateaux. As shown, each of the plateaux in Fig. 6.4 corresponds to a different level of the hierarchy of the synthetic network in Fig. 6.3. Similar to our discussion in Sect. 6.3, in practice the replicas differ from one another in the order in which consecutive vertices are picked and moved so as to minimize the energy of (6.2). Thus, for any given problem has an ensemble of very similar (or nearly identical) viable solutions associated with it. A detailed summary of this approach appears in [20].

In accord with the above explanation, as γ is increased, the associated candidate energy minima partition the system into more local, smaller communities (deeper levels of the hierarchy). The inter-replica information theory correlations further afford a measure of the quality of the viable partitions. High NMI values (i.e., of size close to unity) indicate solutions that are likely to be pertinent. In the spirit of Sect. 6.2, if the different replicas all agree with one another on a putative partition, then that partition is likely to be physically meaningful. The variation of information measures the disparity between candidate solutions; thus the VI values are high between different NMI plateaux and are low within the NMI plateaux.

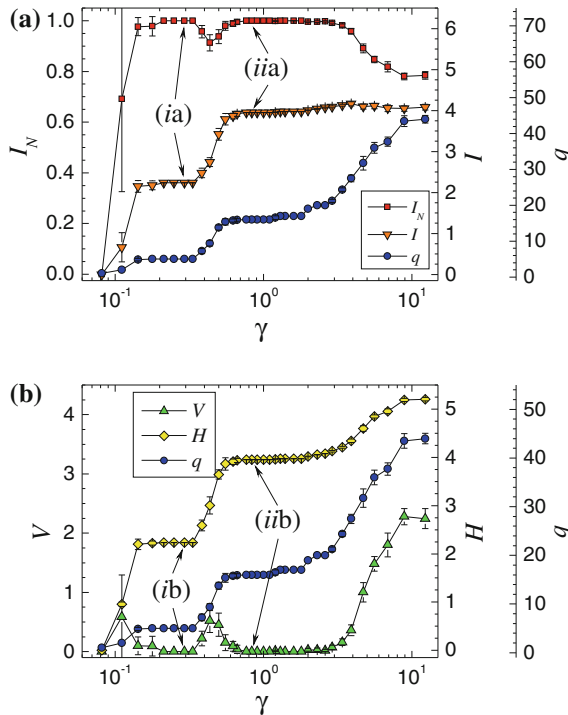


Fig. 6.4 Information theoretic and other metrics of the multiresolution algorithm in Sect. 6.4 as applied to the synthetic partition depicted in Fig. 6.3 [20]. In the *top panel*, the average inter-replica normalized mutual information (I_N), (un-normalized) mutual information (I), and number of clusters (or communities) q are plotted as a function of the resolution parameter γ . In the *bottom panel*, the Shannon entropy (H) and the average inter-replica variation of information (V) are further provided. As described in the text, stable partitions lead to plateaux (or more general local extrema) in the inter-replica information theory and other correlations as a function of the resolution parameter. Two such candidate resolutions (marked (i) and (ii)) are seen in both panels (a) and (b). These plateaux show how the multiresolution algorithm may isolate both level 2 (superclusters) and level 3 (smallest clusters) of the hierarchy of Fig. 6.3

6.5 Image Segmentation

Our goal is to identify structure in materials, but before turning to this endeavor, we first illustrate how patterns may, literally, be revealed by community detection. The ideas underlying this objective will elucidate our approach to material genomics. The aim of image segmentation [52–58] is to divide a given digital image into separate objects (or segments) based on visual characteristics. Two somewhat challenging examples are provided in Fig. 6.5 [59, 60].

To transform the problem into that of community detection, we map a digital image into a network as follows. Each pixel in an image is regarded as a node in a graph. (2) The edge weights between nodes in the graph are determined by the degree



Fig. 6.5 Examples of the image segmentation challenges [59, 60]. *Left* the left image is that of zebra with the a similar “stripe” background. *Right* the image on the right is that of a dalmatian dog. Most people do not initially recognize the dog before given clues as to its presence. Once the dog is seen it is nearly impossible to perceive the image in a meaningless way


of similarity between the additive color RGB (i.e., the Red, Green, and Blue) strength of individual pixels or, more generally, of finite size boxes geometrically centered about a given pixel. The bare edge strengths may be embellished and replaced by weights set by the Fourier weights associated with finite size blocks about a given node. Alternatively, we can use exponential weighting of the inter-node edge strength based on the geometric distance between them (the distance between the centers of the finite size blocks about them) [52]. The edge value assignment is such that if two pixels i and j (or boxes centered about them) have similar RGB values (or absolute Fourier magnitudes), then a function V_{ij} set by these differences will be small. Analogously, if nodes i and j (or boxes centered around them) are dissimilar then V_{ij} will become large.

With such functions V_{ij} at hand, a simple generalization of (6.2) is given by

$$H = \frac{1}{2} \sum_{s=1}^q \sum_{i,j \in C_s} \left[(V_{ij} - \bar{V}) [\Theta(\bar{V} - V_{ij}) + \gamma \Theta(V_{ij} - \bar{V})] \right]. \quad (6.3)$$

Here, $\Theta(z)$ is the Heavyside function ($\Theta(x > 0) = 1$ and $\Theta(x < 0) = 0$) and \bar{V} is an adjustable background value. As the astute reader undoubtedly noticed, the locality constraint imposed by the Kronecker delta in (6.2) has been made explicit in (6.3) by having only intra-community sums for each of the q communities $\{C_s\}$. Details of the construction of the weights V_{ij} are provided in [52]. Following our more colloquial description here, there are four or five adjustable parameters in (6.3): the resolution parameter γ , the background value \bar{V} , the block size L centered about each pixel (or more general rectangular blocks of size $L_x \times L_y$), and the pixel distance ℓ over

Fig. 6.6 The application of multiresolution algorithm for the segmentation of the zebra and dalmatian dog images of Fig. 6.5. The results correspond to typical partitions found with the optimal parameter set. The first and the second rows contain “camouflages” of a similar stype. We are able to detect the boundary of the zebra and discern the body and hind legs of the dog albeit with some “bleeding” [52]



which the pixel interconnection function V_{ij} decays. Once these are set, the earlier community detection algorithm of Sect. 6.3 may be applied. The determination of the optimal value(s) of these parameters may be performed using the same procedure outlined in Sect. 6.4.

While systems such as the synthetic hierarchical network of Fig. 6.4 exhibit well defined plateaux in the information theory and other measures, we found more generally that the optimal values of parameters z correspond to local extrema whereby variations in the parameters do not alter the outcome. That is, if Q is a measured quantity of interest (e.g., information theory correlations, Shannon entropy, the energy associated with the given Hamiltonian) then optimal parameters z are found by the requirement that $\nabla_z Q = 0$. These may lead to multiple viable solutions corresponding to very different meaningful partitions.

In practice, we found that in all but the hardest cases, meaningful solutions are found when arbitrarily setting all parameters to a fixed value and that, similar to Sect. 6.4, the multi-scale solutions may be found by only varying the resolution parameter γ . The results of our method are given in Fig. 6.6; these correspond to typical partitions found with the optimal parameter set. The above image analysis ideas may be applied for the detection of the primitive cells in simple Bravais lattices, the inference of domain walls in spin systems, and hierarchical structures in quasicrystals [52]. For a complete classification of contending partitions and, most notably, a deeper understanding of whether the found solutions are meaningful or not, it is useful to survey the canonical finite temperature phase diagram associated with (6.3) when all of the above parameters, including temperature, are varied. In the current context, by “temperature”, we allude to the finite temperature study of the

Hamiltonian of (6.2) either analytically or via a thermal bath associated with, e.g., the acceptance of the moves in the algorithm outlined at the end of Sect. 6.3 [50, 52, 61–63].

6.6 Community Detection Phase Diagram

As the bare edge weights and additional parameters setting the values of V_{ij} in the Hamiltonian of (6.3) and temperature are modified, quantities such as the system energy, Shannon entropy, the number of communities, and information theory correlations amongst the found ground states generally attest to the presence of multiple phases. Additional metrics including the “computational susceptibility” (the change in the average inter-replica NMI as the number of trials, see Sect. 6.3, is increased [20, 61, 62]), the time required for convergence (when attainable), and the ergodic/non-ergodic character (“chaotic” type feature) of the dynamics all delineate the very same phase diagram boundaries inferred from each of the examined quantities. Information theory measures have been used to study other specific interesting systems, e.g., [64]. The observed phases in the community detection problem naturally extend to finite temperatures (T) when the analysis of the system defined by the Hamiltonian of (6.3) is broadened to include positive temperatures. Finite size systems such as the real networks and images that we discuss cannot exhibit thermodynamic phase transitions and all finite temperature functions are analytic. Nevertheless, practically, sharp changes appear as temperature and other parameters are varied.

Similar to other NP hard [65] combinatorial optimization problems [66–68], three prototypical phases were established in general community detection problems with a distribution of varying community sizes [61]. Subsequently, these have been beautifully explored in depth in several specific graph types—most notably the so-called “stochastic block models”, in which a graph has equal size communities e.g., [69–72] and in other penetrating works, e.g., [73–75]. Earlier signatures of a bona fide transition in stochastic block and power law distributed models [19, 20] and limits on detectability in the stochastic block model via the cavity approximation were suggested [76]. To intuitively highlight the essential character of the prototypical phases with a minimum of jargon, we will colloquially term these the “easily solvable”, the “solvable hard”, and the “unsolvable” phases.

In realistic finite yet very large scale systems [62, 63] various results can be established and these may be further examined in various limits. Of course, bona fide transitions formally occur only in the thermodynamic limit. A trivial behavior results in infinite size graphs when the average number of nodes per community is of finite size [62, 63]. As one would expect, typically all community detection problems are either solvable or unsolvable. In NP hard problems, the solvable phase splinters into an “easy” and a “hard phase”. When the edge weights set by V_{ij} are associated with sharp community detection partitions, then finding a natural solution is rather trivial (and nearly all algorithms, not only the Potts model described here, will readily unearth such an answer). On the other hand, if the couplings V_{ij} are sufficiently




Fig. 6.7 The normalized mutual information I_N as the function of the resolution $\log(\gamma)$ and temperature T for the “bird” image in the upper lefthand panel of Fig. 6.9. We mark the “easy” phase (where I_N is almost 1 as “A”, the “hard” phase where I_N decreases as “B”, the “unsolvable” phase where I_N forms a plateau whose value is less than 1 as “C”. The “easy-hard-unsolvable” phases will be further confirmed by the corresponding image segmentation results in Fig. 6.9, as these appear, respectively, in panels A, B, and C therein)

“noisy” so as to be of, effectively, equally the same strength for edges between nodes in the same putative community as for edges linking nodes belonging to different supposed communities, then no well defined community detection solutions exist. Similarly, at sufficiently high temperatures, in most cases, all traces of structures found in the ground state(s) are lost. The most common variant of the community detection problem has been proven to be NP complete [33].

As in disparate NP problems [68], it was found that in broad classes of the community deception problem (and in its image segmentation variant) [52, 61–63, 69, 71, 73, 75], lying between the extremities of the “easy” and “unsolvable” phases there often exists a “hard phase”; in this phase, solutions exist, but due to the plethora of competing states, they may be extremely hard to find. Information theory measures may be used to delineate phase boundaries [52, 61–63]. Using information theory correlations and the global Shannon entropy, we show, in Figs. 6.7 and 6.8 respectively, the phase diagram associated with the image shown in the upper lefthand side of Fig. 6.9. In the solvable phase(s), typically, all partitions produced by parameters that lie in the same basin, lead to qualitatively similar results. Moderate temperature and/or disorder can lead to order by disorder or annealing effects (similar to those found in other systems, e.g., [77–81]). However, at sufficiently high temperatures and/or the introduction of noise about the initial V_{ij} values, the system will be in the unsolvable phase. By carefully studying the system phase diagram and the character




Fig. 6.8 The Shannon entropy H as the function of the resolution $\log(\gamma)$ and the temperature T for the “bird” image in the upper lefthand panel of Fig. 6.9. The signatures of the three phases “easy”, “hard” and “unsolvable” are easily detected in this phase diagram and agree with those ascertained via the normalized mutual information of Fig. 6.7

and magnitude of the information theory overlaps or thermodynamic functions such as the internal energy and entropy as well as the dynamics, one may assess whether the perceived community detection solutions may be meaningful. When applied to image segmentation, the consistency of this procedure may be inspected visually and intuitively judged sans complicated analysis.

6.7 Casting Complex Materials and Physical Systems as Networks

With all of the above preliminaries, we now finally turn to the ultimate data mining objective of this work: that of the important detection of spatial and temporal structure in complex materials and other systems [50, 51, 82–87]. This problem shares a common conceptual goal with image segmentation yet is, generally, far more daunting for human examination. Similar to the analysis presented thus far, the approach that we wish to discuss casts physical systems as graphs in space or space-time and then employs the above discussed multi-scale community detection to determine meaningful partitions.




Fig. 6.9 The image segmentation results of the “bird” image. The original image is on the *upper left*. The other images denoted as “A”, “B”, and “C” correspond to the image segmentation results with different parameter pairs $(\log(\gamma), T)$ marked in *. Both result A and B are able to distinguish the bird from the “background”. However in panel B, the bird is composed of lots of small clusters. Result C is unable to detect the bird. Thus, the results shown here demonstrate the corresponding “easy-hard-unsolvable” phases in the phase diagram in Figs. 6.7 and 6.8. From [52]

In this case, nodes in the graph code basic physical units of interest (e.g., atoms, electrons, etc.). Multi-particle interactions or experimentally measured correlations in the physical system are then ascribed to edge weights V_{ij} between the nodes (for two-particle interactions or experimentally measured pair correlations [50, 51]), or to three-node triangular weights (for three-particle interactions or correlations) V_{ijk} , and so on. Given these static or time-dependent weights, the graph is then (similar to the discussion in earlier sections) partitioned into “communities” of nodes (e.g., clusters of atoms) that are more tightly linked to or correlated with each other than with nodes in other clusters [19]. As in the earlier examples explored in this work, information theory based multi-scale community detection provides both local structural scales (e.g., primitive lattice cell, nearest neighbor distance, etc.) as well as global scales (such as correlation lengths) and any other additional intermediate scales if and when these are present.

The results of this approach for a two-dimensional Lennard-Jones system with vacancies are shown in Fig. 6.10. When the edge weights between nodes are set equal to the Lennard-Jones strength associated with the distance between them, the multi-scale community detection algorithm recognizes both the typical triangular unit cells as well as larger scale domains (communities) in which the vacancy defects tend, on average, to lie on their boundaries. Partitions in which defects tend to aggregate




Fig. 6.10 A diluted two-dimensional Lennard Jones system with edge weight set equal to the pair interaction energies. The ground state of a two dimensional Lennard Jones model is that of a triangular lattice in which the lattice spacing is equal to the distance at which the Lennard Jones potential attains its minimum. In this figure, the *triangular lattice* is diluted by introducing defects in the form static vacancies (denoted by *white holes*). The found community boundaries are intuitively relegated defects lying on the periphery of these domains [50]

at the domain boundaries is consistent with general expectations for stable domains and is intuitively appealing.

As the reader may envisage, the community detection method may be extended to general many-body systems with different types of species (e.g., disparate ion types in metallic glass formers [50, 51]). One example is depicted in Figs. 6.11, 6.12, and 6.13 corresponding to a ternary system of $\text{Al}_{88}\text{Y}_7\text{Fe}_5$ based on a molecular dynamics simulation of 1600 atoms in which edge weights were set by pair potentials is provided in Figs. 6.12 and 6.13. As seen in the partition of Fig. 6.13, for which the inter-replica information theory were extremal and which lies in the solvable phase, below the liquidus temperature (the temperature at which the system is an equilibrium liquid), large clusters were detected. Along similar lines, clusters may be identified across many problems. In Fig. 6.16 we show typical clusters found in a Kob-Andersen binary system. While for human analysis the complexity of potentially identifying pertinent clusters may grow dramatically with the number of atom types, for the multi-resolution analysis there is no such increase (Figs. 6.14 and 6.15).

Fig. 6.11 From [50, 51]. In order to apply the algorithm in Sect. 6.4 to complex physical systems, we may generally define two types of replica sets. Panel (a) depicts a few nodes as they appear for a static system—i.e., one with no time separation between simulation replicas. Panel (b) depicts a similar set of replicas with each separated by a successive amount of simulation time Δt . In either case, we then generate the replica networks using the potential energy between the atoms as the respective edge weights in the network. Consequently, we minimize (6.1) using a range of γ values in the algorithm described in Sect. 6.4

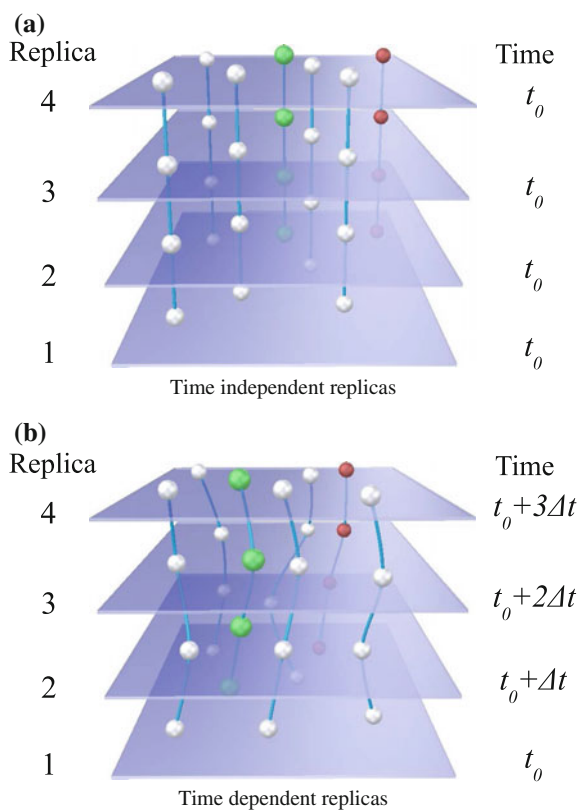



Fig. 6.12 From [50]. A static snapshot from a molecular dynamics simulation of $\text{Al}_{88}\text{Y}_7\text{Fe}_5$ system of 1600 atoms that has been quenched from an initial temperature of 1500 to 300 K and then allowed to partially equilibrate. The atoms are Y, Al, and Fe, respectively, in order of increasing diameters. In this figure, the atoms are color coded—Fe atoms are *red* and Y atoms are marked *green*






Fig. 6.13 The figure shows a static partition of Fig.6.12. Here, different clusters are identified by individual colors. It is also possible to incorporate overlapping nodes in neighboring clusters to account for the possibility of multiple cluster memberships per node, yielding an interlocking system of clusters [50]

In a similar manner, the edge weights can be set by experimentally measured pair correlations. In [50], atomic configurations consistent with the experimentally determined scattering data for quenched $Zr_{80}Pt_{20}$ [3–6] were generated [50, 51] using Reverse Monte Carlo methods [7, 8].

At low temperatures, typically the found structures in all of these cases are far larger than local patterns probed for and detected by current methods [88–92]. Four-point correlations have long been employed to ascertain spatio-temporal scales and the quantify “dynamical heterogeneities”, e.g., [91, 93]. A long-standing challenge is the identification of structures of *general character and scale* in amorphous systems. There is, in fact, a proof that as supercooled liquid falls out of equilibrium to become an amorphous, there must be an accompanying divergent length scale [94]. Methods of characterizing local structures [9–12] center on a given atom or link; as such, they are restricted from detecting general structures. Because of the lack of a simple crystalline reference, the structure of glasses is notoriously difficult to quantify beyond the very local scales. In [50–52], graph weights were determined empirically (potentials in a model system, experimentally measured partial pair density correlations in supercooled fluids, or pixels in a given image)—no theoretical input was invoked as to what the important scales should be or if an exotic order parameter may be concocted. Similarly, in a time dependent analysis for dynamically evolving systems, by employing replicas at different time slices as well as regarding the system as a higher

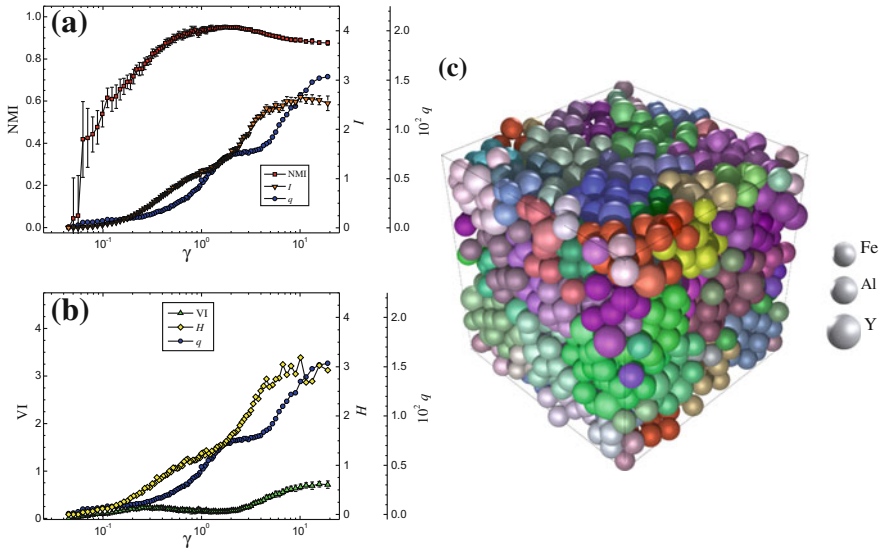


Fig. 6.14 The result of the multiscale community detection applied to a ternary glass former at a simulation temperature of $T = 300$ K [50, 51]. Both panels (a) and (b) on the left depict the information theory correlations between the replicas (as described in Sect. 6.4). In panel (c), each of the communities found is assigned a different color. These structures correspond to the Normalized Mutual Information (NMI) or Variation of Information (VI) extrema. These well-defined structures contrast sharply with the lack of cohesive features in Fig. 6.15

dimensional “image” in space-time, using the inter-replica information theory correlations, spatio-temporal patterns were found and time dependent structures were quantified. In this approach, the data speak for themselves. We remark that notwithstanding the aforementioned difficulties, recently extremely large growth of static structure was observed by far simpler network analysis in certain binary metallic glasses that exhibit crisp icosahedral motifs [96]. Similar to the description above, one may likely find other motifs in other systems. The problem is that guessing and hopefully finding pertinent patterns can be extremely challenging to do by conventional analysis.

6.8 Summary

In this work, we reviewed key features of a statistical-mechanics-based “community detection” approach to find pertinent features and structures (both spatial and temporal) in complex systems. In particular, we illustrated how this method may be applied to image segmentation and the analysis of amorphous materials. The demand for automated data mining approaches may become more acute with ever

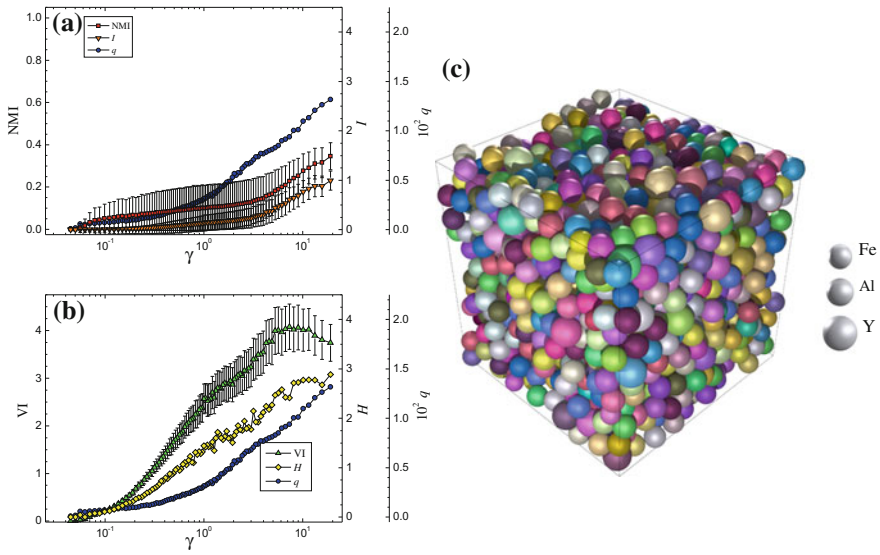



Fig. 6.15 The structure of same ternary glass former of Fig. 6.15 at a simulation temperature of $T = 1500$ K. Inter-replica information theory correlations are provided in panels (a) and (b). As is evident in panel (c), depict the corresponding lack of structure by significantly higher VI or lower NMI as compared to those in Fig. 6.14

Fig. 6.16 From [50]. A set of optimal clusters found in a low temperature Kob-Andersen system [95] in which two types of atoms (color coded *red* and *silver*) appear



increasingly available data on numerous complex systems. The study of complex materials may be extremely challenging to carry out by current conventional means that rely on guessed patterns, simplified models, or brute force human examination.

Acknowledgments We have benefited from interactions with numerous colleagues. In particular, we would like to thank S. Achilefu, S. Bloch, R. Darst, S. Fortunato, V. Gudkov, K.F. Kelton, T. Lookman, M.E.J. Newman, S. Nussinov, D.R. Reichman, and P. Sarder for numerous discussions and collaboration on some of the problems reviewed in this work and their outgrowths. We are further grateful to support by the NSF under Grants No. DMR-1106293 and DMR-1411229. ZN is indebted to the hospitality and support of the Feinberg foundation for visiting faculty program at the Weizmann Institute.

References

1. C.A. Angell, Formation of glasses from liquids and biopolymers. *Science* **267**(5206), 1924–1935 (1995)
2. W.H. Zachariasen, The atomic arrangement in glass. *J. Am. Chem. Soc.* **54**, 3841 (1932)
3. T. Nakamura, E. Matsubara, M. Sakurai, M. Kasai, A. Inoue, Y. Waseda, Structural study in amorphous Zr-noble metal (Pd, Pt and Au) alloys. *J. Non-Cryst. Solids* **312–314**, 517 (2002)
4. J. Saida, K. Itoh, S. Sato, M. Imafuku, T. Sanada, A. Inoue, Evaluation of the local environment for nanoscale quasicrystal formation in $Zr_{80}Pt_{20}$ glassy alloy using Voronoi analysis. *J. Phys. Condens. Matter* **21**, 375104 (2009)
5. D.J. Sordelet, R.T. Ott, M.Z. Li, S.Y. Wang, C.Z. Wang, M.F. Besser, A.C.Y. Liu, M.J. Kramer, Structure of $Zr_x Pt_{100-x}$ ($73 \leq x \leq 77$) metallic glasses. *Metall. Mater. Trans. A* **39A**, 1908–1916 (2008)
6. S.Y. Wang, C.Z. Wang, M.Z. Li, L. Huang, R.T. Ott, M.J. Kramer, D.J. Sordelet, K.M. Ho, Short- and medium-range order in a $Zr_{73}Pt_{27}$ glass: experimental and simulation studies. *Phys. Rev. B* **78**, 184204 (2008)
7. R.L. McGreevy, Understanding liquid structures. *J. Phys. Condens. Matter* **3**, F9 (1991)
8. D.A. Keen, R.L. McGreevy, Structural modelling of glasses using reverse Monte Carlo simulation. *Nature* **344**, 423–5 (1990)
9. H.W. Sheng, W.K. Luo, F.M. Alamgir, J.M. Bai, E. Ma, Atomic packing and short-to-medium-range order in metallic glasses. *Nature* **439**, 419–425 (2006)
10. J.L. Finney, Random packings and the structure of simple liquids. I. The geometry of random close packing. *Proc. R. Soc. Lond. Ser. A* **319**(1539), 479–493 (1970)
11. J. Dana Honeycutt, H.C. Andersen, Molecular dynamics study of melting and freezing of small Lennard-Jones clusters. *J. Phys. Chem.* **91**, 4950–4963 (1987)
12. P.J. Steinhardt, D.R. Nelson, M. Ronchetti, Bond-orientational order in liquids and glasses. *Phys. Rev. B* **28**, 784–805 (1983)
13. T.R. Kirkpatrick, D. Thirumalai, P.G. Wolynes, Scaling concepts for the dynamics of viscous liquids near an ideal glassy state. *Phys. Rev. A* **40**, 1045–1054 (1989)
14. V. Lubchenko, P.G. Wolynes, Theory of structural glasses and supercooled liquids. *Annu. Rev. Phys. Chem.* **58**, 235–266 (2007)
15. G. Tarjus, S.A. Kivelson, Z. Nussinov, P. Viot, The frustration-based approach of supercooled liquids and the glass transition: a review and critical assessment. *J. Phys. Condens. Matter* **17**, R1143–R1182 (2005)
16. Z. Nussinov, Avoided phase transitions and glassy dynamics in geometrically frustrated systems and non-Abelian theories. *Phys. Rev. B* **69**, 014208 (2004)
17. <http://www.whitehouse.gov/mgi>
18. S. James, *The Wisdom of Crowds* (Anchor Books, New York, 2005). ISBN: 0-385-72170-6
19. P. Ronhovde, Z. Nussinov, An improved potts model applied to community detection. *Phys. Rev. E* **81**, 046114 (2010)
20. P. Ronhovde, Z. Nussinov, Multiresolution community detection for megascale networks by information-based replica correlations. *Phys. Rev. E* **80**, 016109 (2009)

21. B. Sun, B. Leonard, P. Ronhovde, Z. Nussinov, An interacting replica approach applied to the traveling salesman problem (2014). [arXiv:1406.7282.pdf](https://arxiv.org/abs/1406.7282)
22. M. Dorigo, T. Sttzle, *Ant Colony Optimization* (MIT Press, Cambridge, 2004) ISBN: 0-262-04219-3
23. M. Mitchell, *An Introduction to Genetic Algorithms* (MIT Press, Cambridge, 1996)
24. S. Fortunato, Community detection in graphs. *Phys. Rep.* **486**, 75–174 (2010)
25. M.E.J. Newman, Modularity and community structure in networks. *Proc. Natl. Acad. Sci. USA* **103**(23), 8577–8582 (2006)
26. M.E.J. Newman, M. Girvan, Finding and evaluating community structure in networks. *Phys. Rev. E* **69**, 026113 (2004)
27. S. Fortunato, M. Barthelemy, Resolution limit in community detection. *Proc. Natl. Acad. Sci. USA* **104**, 36–41 (2007)
28. A. Lancichinetti, S. Fortunato, Community detection algorithms: a comparative analysis. *Phys. Rev. E* **80**, 056117 (2009)
29. V.D. Blondel, J.-L. Guillaume, R. Lambiotte, E. Lefebvre, Fast unfolding of communities in large networks. *J. Stat. Mech.* **10**, 10008 (2008)
30. M.E.J. Newman, Fast algorithm for detecting community structure in networks. *Phys. Rev. E* **69**, 066133 (2004)
31. V. Gudkov, V. Montelaegre, S. Nussinov, Z. Nussinov, Community detection in complex networks by dynamical simplex evolution. *Phys. Rev. E* **78**, 016113 (2008)
32. M. Rosvall, C.T. Bergstrom, Maps of random walks on complex networks reveal community structure. *Proc. Natl. Acad. Sci. USA* **105**, 1118–1123 (2008)
33. U. Brandes, D. Dellng, M. Gaertler, R. Gorke, M. Hoefer, Z. Nikoloski, D. Wagner, On finding graph clusterings with maximum modularity. In *Graph-Theoretic Concepts in Computer Science*. Lecture Notes in Computer Science (Springer, Berlin, 2007). doi:[10.1007/978-3-540-74839-7](https://doi.org/10.1007/978-3-540-74839-7)
34. R.K. Darst, D.R. Reichman, P. Ronhovde, Z. Nussinov, An edge density definition of overlapping and weighted graph communities (2013). [arXiv:1301.3120](https://arxiv.org/abs/1301.3120)
35. M.E.J. Newman, Spectral methods for community detection and graph partitioning. *Phys. Rev. E* **88**, 042822 (2013)
36. M.E.J. Newman, Community detection and graph partitioning. *Europhys. Lett.* **103**, 28003 (2013)
37. R.K. Darst, Z. Nussinov, S. Fortunato, Improving the performance of algorithms to find communities in networks. *Phys. Rev. E* **89**, 032809 (2014)
38. J. Reichardt, S. Bornholdt, Statistical mechanics of community detection. *Phys. Rev. E* **74**, 016110 (2006)
39. P. Tiago Piexoto, Efficient Monte Carlo and greedy heuristic for the inference of stochastic block models., *Phys. Rev. E* **89**, 012804 (2014)
40. J.M. Kumpula, J. Saramaki, K. Kaski, J. Kertesz, Limited resolution in complex network community detection with Potts model approach. *Eur. Phys. J. B* **56**, 41 (2007)
41. P. Ronhovde, Z. Nussinov, Local multi resolution order in community detection. *J. Stat. Mech.* P01001 (2015)
42. L.G.S. Jeub, P. Balachandran, M.A. Porter, P.J. Mucha, M.W. Mahoney, Think locally, act locally: the detection of small, medium-sized, and large communities in large networks (2014). [arXiv:1403.3795.pdf](https://arxiv.org/abs/1403.3795)
43. M. De Domenico, A. Insolia, Entropic approach to multiscale clustering analysis. *Entropy* **14**, 865 (2012)
44. P. Tiago Piexoto, Hierarchical block structures and high-resolution model selection in large networks. *Phys. Rev. X* **4**, 011047 (2014)
45. S. Wiseman, M. Blatt, E. Domany, Superparamagnetic clustering of data. *Phys. Rev. E* **57**, 3767 (1998)
46. A.L.N. Fred, A.K. Jain, Robust data clustering. In *2003 Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2 (2003), pp. 128–133

47. M. Meil, Comparing clusterings—an information based distance. *J. Multivar. Anal.* **98**, 873–895 (2007)
48. L. Danon, A. Diaz-Guilera, J. Duch, A. Arenas, Comparing community structure identification. *J. Stat. Mech. Theory Exp.* **9**, 09008 (2005)
49. G. Bianconi, Statistical mechanics of multiplex networks: entropy and overlap. *Phys. Rev. E* **87**, 062806 (2013)
50. P. Ronhovde, S. Chakrabarty, M. Sahu, K.F. Kelton, N.A. Mauro, K.K. Sahu, Z. Nussinov, Detecting hidden spatial and spatio-temporal structures in glasses and complex physical systems by multiresolution network clustering. *Eur. Phys. J. E* **34**, 105 (2011)
51. P. Ronhovde, S. Chakrabarty, M. Sahu, K.K. Sahu, K.F. Kelton, N. Mauro, Z. Nussinov, Detection of hidden structures on all scales in amorphous materials and complex physical systems: basic notions and applications to networks, lattice systems, and glasses. *Sci. Rep.* **2**, 329 (2012)
52. D. Hu, P. Ronhovde, Z. Nussinov, A replica inference approach to unsupervised multi-scale image segmentation. *Phys. Rev. E* **85**, 016101 (2012)
53. L.G. Shapiro, G.C. Stockman, *Computer Vision* (Prentice-Hall, New Jersey, 2001), pp. 279–325. ISBN: 0-13-030796-3
54. J. Shi, J. Malik, Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **22**(8), 888–905
55. L. Wang, H. Cheng, Z. Liu, C. Zhu, A robust elastic net approach for feature learning. *J. Vi. Commun. Image Represent.* **25**, 313 (2014)
56. A.A. Abin, F. Mahdisoltani, H. Beigy, WISECODE: wise image segmentation based on community detection. *Imaging Sci. J.* **62**, 327 (2014)
57. H. Dandan, P. Sarder, P. Ronhovde, S. Bloch, S. Achilefu, Z. Nussinov, Automatic segmentation of fluorescence lifetime microscopy images of cells using multiresolution community detection: a first study. *J. Microsc.* **253**(1), 54–64 (2014)
58. D. Hu, P. Sarder, P. Ronhovde, S. Bloch, S. Achilefu, Z. Nussinov, Community detection for fluorescent lifetime microscopy image segmentation. In *Proceedings of the SPIE 8949, Three-Dimensional and Multidimensional Microscopy: Image Acquisition and Processing XXI* (2014), p. 89491K. <http://dx.doi.org/10.1117/12.2036875>
59. <http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/>
60. See <http://www.gifford.co.uk/?principia/Illusions/dalmatian.htm>
61. D. Hu, P. Ronhovde, Z. Nussinov, Phase transitions in random Potts systems and the community detection problem: spin-glass type and dynamic perspectives. *Philos. Mag.* **92**(4), 406–445 (2012). [arXiv:1008.2699](https://arxiv.org/abs/1008.2699) (2010)
62. H. Dandan, P. Ronhovde, Z. Nussinov, Stability-to-instability transition in the structure of large-scale networks. *Phys. Rev. E* **86**, 066106 (2012)
63. P. Ronhovde, H. Dandan, Z. Nussinov, Global disorder transition in the community structure of large-q Potts systems. *EPL (Europhys. Lett.)* **99**(3), 38006 (2012)
64. O. Melchert, A.K. Hartmann, Information-theoretic approach to ground-state phase transitions for two- and three-dimensional frustrated spin systems. *Phys. Rev. E* **87**, 022107 (2013)
65. S. Cook, The complexity of theorem-proving procedures. In *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, 1971) pp. 151–158
66. P. Cheeseman, B. Kanefsky, W.M. Taylor, Where the really hard problems are? In *Proceedings of 12th International Joint Conference on AI (IJCAI-91) Automated Reasoning* vol. 1, ed. by J. Mylopoulos, R. Reiter (1991), p. 331
67. R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, Lidror Troyansky, *Nature* **400**, 133 (1999)
68. M. Mezard, G. Parisi, R. Zecchina, Analytic and algorithmic solution of random satisfiability problems. *Science* **297**, 812 (2002)
69. A. Decelle, F. Krzakala, C. Moore, L. Zdeborova, Phase transition in the detection of modules in sparse networks. *Phys. Rev. Lett.* **107**, 065701 (2011). [arXiv:1102.1182](https://arxiv.org/abs/1102.1182)
70. E. Mossel, J. Neeman, A. Sly, Stochastic block models and reconstruction (2012). [arXiv:1202.1499](https://arxiv.org/abs/1202.1499)

71. R.R. Nadakuditi, M.E.J. Newman, Graph spectra and the detectability of community structure in networks. *Phys. Rev. Lett.* **108**, 188701 (2012)
72. R.K. Darst, D.R. Reichman, P. Ronhovde, Z. Nussinov, Algorithm independent bounds on community detection problems and associated transitions in stochastic block model graphs. *J. Complex Netw.* (2014). doi:[10.1093/comnet/cnu042](https://doi.org/10.1093/comnet/cnu042)
73. G. Ver Steeg, C. Moore, A. Galstyan, A. Allahverdyan, Phase transitions in community detection: a solvable toy model. *Europhys. Lett.* **106**, 48004 (2014)
74. A. Montanari, Finding one community in a sparse graph (2015). [arXiv:1502.05680](https://arxiv.org/abs/1502.05680)
75. X. Zhang, R.R. Nadakuditi, M.E.J. Newman, Spectra of random graphs with community structure and arbitrary degrees. *Phys. Rev. E* **89**, 042816 (2014)
76. J. Reichardt, M. Leone, (Un)detectable cluster structure in sparse networks. *Phys. Rev. Lett.* **101**, 78701 (2008)
77. S. Kirkpatrick, C.D. Gelatt Jr, M.P. Vecchi, Optimization by simulated annealing. *Science* **220**, 671 (1983)
78. J. Villain, R. Bidaux, J.P. Carton, R. Conte, Order as an effect of disorder. *J. Physique* **41**, 1263 (1980)
79. C.L. Henley, Ordering due to disorder in a frustrated vector antiferromagnet. *Phys. Rev. Lett.* **62**, 2056 (1989)
80. Z. Nussinov, M. Biskup, L. Chayes, J. van den Brink, Orbital order in classical models of transition-metal compounds. *Europhys. Lett.* **67**, 990 (2004)
81. P.G. Wolynes, Folding funnels and energy landscapes of larger proteins within the capillarity approximation. *Proc. Natl. Acad. Sci. USA* **94**(12), 6170–6175 (1997)
82. D.S. Bassett, E.T. Owens, K.E. Daniels, M.A. Porter, Influence of network topology on sound propagation in granular materials. *Phys. Rev. E* **86**, 041306 (2012)
83. F. Cerina, V. De Leo, M. Barthelemy, A. Chessa, Spatial correlations in attribute communities. *PLoS ONE* **7**(5), e37507 (2012)
84. P. Holme, J. Saramaki, Temporal networks. *Phys. Rep.* **519**, 97 (2012)
85. A. Cardillo, J. Gmez-Gardenes, M. Zanin, M. Romance, D. Papo, F. del Pozo, S. Boccaletti, Emergence of network features from multiplexity. *Sci. Rep.* **3**, 1344 (2013)
86. G. Petri, P. Expert, Temporal stability of network partitions. *Phys. Rev. E* **90**, 022813 (2014)
87. R.L. Jack, A.J. Dunleavy, C. Patrick Royall, Information-theoretic measurements of coupling between structure and dynamics in glass formers. *Phys. Rev. Lett.* **113**, 095703 (2014)
88. J.-P. Bouchaud, G. Biroli, On the Adams-Gibbs-Kirkpatrick-Thirumalai-Wolynes scenario for the viscosity increase in glasses. *J. Chem. Phys.* **121**, 7347 (2004)
89. M. Mosayebi, E.D. Gado, P. Ig, H.C. Ottinger, Probing a critical length at the glass transition. *Phys. Rev. Lett.* **104**, 205704 (2010)
90. L. Berthier, G. Biroli, J.-P. Bouchaud, L. Cipelletti, D. El Masri, D. L'Hote, F. Ladieu, M. Pierno, Direct experimental evidence of a growing length scale accompanying the glass transition. *Science* **310**, 1797 (2005)
91. S. Karmakar, C. Dasgupta, S. Sastry, Growing length and time scales in glass-forming liquids. *Proc. Natl. Acad. Sci. USA* **106**, 3675 (2010)
92. J. Kurchan, D. Levine, Correlation length for amorphous systems (2009). [arXiv:0904.4850](https://arxiv.org/abs/0904.4850)
93. C. Dasgupta, A.V. Indrani, S. Ramaswamy, M.K. Phani, Is there a growing correlation length near the glass transition? *Europhys Lett.* **15**, 307 (1991)
94. A. Montanari, G. Semerjian, Rigorous inequalities between length and time scales in glassy systems. *J. Stat. Phys.* **125**, 23–54 (2006)
95. W. Kob, H.C. Andersen, Testing mode-coupling theory for a supercooled binary Lennard-Jones mixture: the van Hove correlation function. *Phys. Rev. E* **51**, 4626 (1995)
96. R. Soklaski, Z. Nussinov, Z. Markow, K.F. Kelton, L. Yang, Connectivity of icosahedral network and a dramatically growing static length scale in Cu-Zr binary metallic glasses. *Phys. Rev. B* **87**, 184203 (2013)

Part II
Materials Prediction with Data,
Simulations and High-throughput
Calculations

Chapter 7

On the Use of Data Mining Techniques to Build High-Density, Additively-Manufactured Parts

Chandrika Kamath

Abstract The determination of process parameters to build additively-manufactured parts with desired properties remains a challenge, especially as we move from machine to machine or process new materials. In this chapter, we show how we can combine simple simulations and experiments to iteratively constrain the design space of parameters, and quickly and efficiently identify parameters to create parts with >99% density. Our approach is based on techniques from statistics and data mining, including design of physical and computational experiments, feature selection to identify important variables, and data-driven predictive models that can act as surrogates for the simulations.

7.1 Introduction

Additive manufacturing (AM), a process for fabricating parts layer by layer directly from a three-dimensional digital model, presents an opportunity for producing complex, individually-customized parts not possible with traditional manufacturing processes. While AM can reduce both the time to market and material waste, a number of technical issues must still be addressed before widespread use of AM technology becomes a reality. These include gaps in measurement methods, accuracy of AM parts, process optimization to quickly build parts with desired properties, and increased confidence in properties of parts fabricated using this process [16].

In this chapter, we focus on metal AM using selective laser melting (SLM), which is a powder-based AM process where a three-dimensional part is produced layer by layer by using a high-energy laser beam to fuse metallic powder particles. We are interested in developing an approach that can be used to identify process parameters that would result in high-density (>99%) parts. We start by describing the process of laser powder-bed fusion and discuss the current approaches to optimizing AM parts


C. Kamath (✉)

Lawrence Livermore National Laboratory, 7000 East Avenue, Livermore, CA 94551, USA
e-mail: kamath2@llnl.gov

© Springer International Publishing Switzerland 2016
T. Lookman et al. (eds.), *Information Science for Materials
Discovery and Design*, Springer Series in Materials Science 225,
DOI 10.1007/978-3-319-23871-5_7

141

Fig. 7.1 Schematic illustrating the SLM process and some of the process parameters that influence the properties of a part



for high density. We then describe our approach that combines simple simulations and experiments using techniques from data mining and statistics. We illustrate our approach using 316L stainless steel as an example, and show that it is indeed possible to efficiently arrive at process parameters that result in high-density AM parts.

7.1.1 Additive Manufacturing Using Laser Powder-Bed Fusion

In SLM using metal powder-bed fusion, a three-dimensional digital model of the part is first sliced into two-dimensional layers, each of a specified thickness, usually in the range of 30–100 μm . Metal powder is then spread on a base plate and the first layer is created by selectively melting the powder in the locations indicated in the first slice of the part. The next layer of powder is then spread over the first layer and the powder melted in the regions corresponding to the second slice of the part. Thus, the part is built layer by layer, with the power and speed of the laser selected so that the energy density is sufficient to melt the powder and the layer below it, integrating the new layer into the rest of the part.

The design freedom afforded by AM comes with associated complexity. There are a large number of parameters, more than 130 by some estimates [19], that influence the process and thus the final quality of the part. Some of these parameters pertaining to the laser and the powder bed are shown in Fig. 7.1. The large number of parameters and the complex interactions among them make it challenging to determine the values that should be used to create parts with desired properties.

7.2 Optimizing AM Parts for Density: The Current Approach

There has been much work done in finding optimal parameters that result in additively-manufactured parts with >99 % density (see for example, the summary in [10] for the work done in 316L stainless steel). Initially, the approach taken was an

experimental one, where small cubes were built to understand how various process parameters, such as powder quality, layer thickness, laser power, laser speed, and scanning strategies, would influence the density and surface roughness of a part [13, 18]. Other efforts performed a very systematic study, by carefully identifying the factors that influenced the density, surface roughness, and mechanical properties of a part, and using micrographs and various measurements to understand the effects of these factors [21]. Since much of this work was done using systems with relatively low laser powers of 50–100 W, the design space spanned by laser power and speed was not very large, making optimization through experimentation a practical option.

A slightly different approach was taken by Kempen et al. in their study of process optimization for AlSi10Mg. They started with single track experiments [20], where single tracks are made on a layer of powder using a range of laser power and speed values. The resulting melt-pool characteristics were then analyzed to identify a process window for use in optimization. Tracks considered for inclusion in the window were those that met certain constraints, such as track continuity, a large height of the track to build up the part, and a connection angle of near 90° with the previous layer so that the part would be of high density and dimensionally accurate. A similar approach was also taken by Laohaprapanon et al. [14], who used single track experiments to narrow the space of power and speed values to use in building cubes for density optimization.

More recently, with higher-powered lasers and new scan strategies expanding the design space, techniques from statistics, including the design and analysis of experiments [4, 17], have started playing a role in systematic studies to understand the influence of the parameters on various properties of the parts. For example, Delgado et al. [2] used a full factorial experimental design with three factors (layer thickness, scan speed, and build direction) and two levels per factor in their study on part quality for a fixed laser power. The outputs of interest were dimensional accuracy, mechanical properties, and surface roughness. The results of the experiments were analyzed using an ANOVA (ANalysis Of VAriance) approach to understand the effects of various factors on the outputs.

To complement the insight gained into SLM using experiments, scientists are also using computer simulations to understand the relationship between processing parameters and the thermal behavior of the material as it is melted by the laser [5, 7, 11, 15]. When these three-dimensional simulations include various aspects of the physics underlying SLM, they can be quite expensive to run, even on high-performance computer systems.

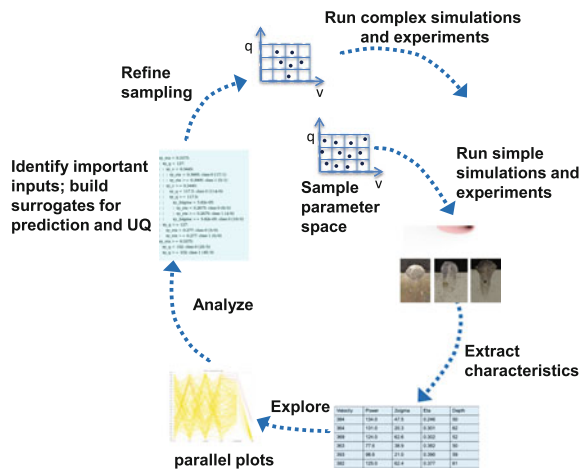
Our approach builds on these ideas and uses both simulations and experiments, combining the insight from each using statistics and data mining techniques. Our goal is to reduce the time it takes to determine the process parameters required to build high-density parts.

7.3 A Data Mining Approach Combining Experiments and Simulations

Despite the wealth of literature on parameters used to create high-density parts with commonly-used materials, such as 316L stainless steel, it is still a challenge to determine the appropriate parameters to use as we move from one machine to another with different power ranges or beam sizes, change powder sizes, or work with new materials. Our work was motivated by the fact that our AM machine, a Concept Laser M2 system, had a relatively narrow beam, with $D4\sigma = 54 \mu\text{m}$, and maximum power of 400 W. As a result, we could not use the parameters for optimal density that were available in the literature as these were for machines with lower powers of $<225 \text{ W}$ and larger beam sizes of $D4\sigma \approx 120 \mu\text{m}$. Given the large range of laser power (0–400 W) for our machine, we realized that a design of experiments approach would require a large number of samples to fully explore the design space, making such an approach prohibitively expensive. We therefore needed an alternative that would help us to determine the optimal parameters for our machine efficiently.

Figure 7.2 illustrates the systematic approach we devised that combines computer simulations and experiments. The approach is an iterative one. Starting with a densely-sampled design space of parameters, we run simple, and relatively inexpensive, simulations and experiments to progressively narrow the space of parameters as we move towards more expensive and accurate simulations and experiments. In each cycle, we have a set of samples that span the space of interest, which is the space of input SLM parameters. We run the experiments and/or simulations at the sample points, extract the characteristics of interest (such as the melt-pool characteristics or the density), and analyze the data that relate the sample points to the characteristics of interest. This analysis could include visualization using scatter plots or parallel-coordinate plots [8], feature selection to identify important parameters, building

Fig. 7.2 Schematic illustrating the iterative process that combines simulations and experiments to reduce the time and costs to determine optimal density parameters



surrogate models for prediction, and uncertainty quantification to find regions that are less sensitive to minor changes in the parameters. As a result of this analysis, we identify a subset of samples that meet our requirements. We then perform more complex simulations and experiments at these sample points, and iterate until we have obtained the desired results.

This iterative approach has several benefits. First, by starting with simple simulations and experiments, we can quickly and efficiently identify which regions of the design space are viable and which are unlikely to result in melt pools that are deep enough so that a part can be built. This is particularly relevant when we are working with materials that may not have been additively manufactured before, or with machines with different process parameters, or with powders with different size distributions. Second, the large number of parameters that have to be set in laser powder-bed fusion implies that we need to identify sample points in a high-dimensional space, where the dimension of the space is the number of parameters. To span a space adequately, the number of samples we need is exponential in the dimension. This makes it prohibitively expensive to start exploring the entire space by building complex parts. Starting with simpler experiments and simulations allows us to lower the cost of exploring the space of parameters more fully, thus increasing the chance of finding all sets of parameters that yield desired properties. Third, the iterative approach enables us to progressively make larger samples and perform more complex simulations, while building on what we have learned from simpler experiments and simulations. Finally, by using data mining techniques to analyze the data from the simulations and experiments at each step, we can fully exploit the data we do collect and better guide the next set of experiments and simulations.

We next describe how we used this approach to identify process parameters for high-density 316L stainless steel. We have also successfully applied this approach to create parts with >99% density for other materials and the ideas can be extended to other properties of a part as well.

7.3.1 Using Simple Simulations to Identify Viable Parameters

To identify the viable range of process parameters, we started with the very simple Eagar-Tsai (E-T) model [3] to determine under what conditions we would obtain melt pools that were deep enough to melt a layer of powder and the substrate below. E-T considers a Gaussian beam on a flat plate to describe conduction-mode laser melting. The resulting temperature distribution is then used to compute the melt-pool width, depth, and length as a function of four parameters—laser power, laser speed, beam size, and laser absorptivity of the powder.

Note that the E-T model does not directly relate the process parameters to the density of a part. Further, it does not consider powder other than the effect of powder on absorptivity, so its results provide only an estimate of the melt-pool characteristics. However, we found that this estimate was sufficient to guide the next steps in our work. In addition, the simplicity of the model made it computationally inexpensive,

taking ≈ 1 min to run on a laptop. This allowed us to use the E-T model to sample the input parameter space rather densely, ensuring that we considered all possible viable cases.

7.3.1.1 Sampling the Design Space

We used a full factorial design of computer experiments [4, 17] to explore the four-parameter input space. This method divides the range of each parameter into several levels, and then randomly selects a point in each cell. We varied the speed from 50 to 2250 mm/s with 10 levels, the power from 50 to 400 W using 7 levels, the beam size ($D4\sigma$) from 50 to 68 μm using 3 levels, and the laser absorptivity from 0.3 to 0.5 using 2 levels. This resulted in 462 parameter combinations that were input to our simulation.

The range of values for each variable was selected as follows. Our CL20 machine had a peak power of 400 W, which determined the upper bound on the power. The lower limit on the speed was set to ensure sufficient melting at the low power values such that the melt-pool depth would be at least 30 μm (the layer thickness selected for our experiments). The upper limit on the speed was estimated at a value that would likely result in a relatively shallow melt pool at the high power value. The lower and upper limits on the beam size were obtained from measurements of the beam size on our machine at focus offsets of 0 and 1 mm. By varying the beam size and the absorptivity, we were able to account for possible variations in these parameters over time or build conditions as we built the parts.

7.3.1.2 Selecting Important Input Parameters

Having identified the sample points in the four-dimensional space of laser power, laser speed, beam size, and laser absorptivity, we then ran the E-T simulations at these sample points and obtained the melt-pool width, depth, and length. This output from the simulations was analyzed in several different ways. In earlier work [10], we showed how we can use parallel-coordinate plots [8] and feature selection methods from data mining [9] to identify input variables that are more relevant to the melt-pool characteristics. We use the term “feature” to refer to variables, such as the input parameters, that describe a simulation. The feature selection methods we used were designed for problems with discrete data, so we first discretized the continuous input and output variables before applying the method. Since the results could potentially depend on the discretization used, in this chapter, we consider two methods that work directly with the continuous variables.

The Correlation-based Feature Selection (CFS) method [6] is a simple approach that calculates a figure of merit for a feature subset of k features as

$$\text{Merit} = \frac{k\bar{r}_{cf}}{\sqrt{k + k(k-1)r_{ff}}} \quad (7.1)$$

where $\overline{r_{cf}}$ is the average feature-output correlation and $\overline{r_{ff}}$ is the average feature-feature correlation. We use the Pearson correlation coefficient between two vectors, X and Y , defined as

$$\frac{Cov(X, Y)}{\sigma_X \sigma_Y} \tag{7.2}$$

where $Cov(X, Y)$ is the covariance between the two vectors and σ_X is the standard deviation of X . A higher value of Merit results when the subset of features is such that they have a high correlation ($\overline{r_{cf}}$) with the output and a low correlation ($\overline{r_{ff}}$) among themselves.

In the second feature selection method, the features are ranked using the mean squared error (MSE) as a measure of the quality of a feature [1]. This metric is used in regression trees (see Sect. 7.3.1.3) to determine which feature to use to split the samples at a node of the tree. Given a numeric feature x , the feature values are first sorted ($x_1 < x_2 < \dots < x_n$). Then, each intermediate value, $(x_i + x_{i+1})/2$, is proposed as a splitting point, and the samples split into two depending on whether the feature value of a sample is less than the splitting point or not. The MSE for a split A is defined as

$$MSE(A) = p_L \cdot s(t_L) + p_R \cdot s(t_R) \tag{7.3}$$

where t_L and t_R are the subset of samples that go to the left and right, respectively, by the split based on A , p_L and p_R are the proportion of samples that go to the left and right, and $s(t)$ is the standard deviation of the $N(t)$ output values, c_i , of samples in the subset t :

$$s(t) = \sqrt{\frac{1}{N(t)} \sum_{i=1}^{N(t)} (c_i - \overline{c(t)})^2} \tag{7.4}$$

For each feature, the minimum MSE across the values of the feature is obtained and the features are rank ordered by increasing values of their minimum. This method considers a feature to be important if it can split the data set into two, such that the standard deviation of the samples on either side of the split is minimized, that is, the output values are relatively similar on each side. Note that unlike CFS that considers subsets of features, this method considers each feature individually.

Table 7.1 presents the ordering of subsets of input features by importance for the melt-pool width, length, and depth obtained using the CFS method. A noise feature was added as another input; this is consistently ranked as the least important variable, as might be expected. The table indicates that for the melt-pool depth and width, the single most important input is the speed, while the top two most important inputs are the speed and power. In contrast, for the length of the melt pool, the most important single input is the power, while the top two most important inputs are power and absorptivity.

Table 7.1 Rank order of subsets of the input parameters to the Eagar-Tsai model using the CFS filter

	Speed	Power	Beam size	Absorptivity	Noise
Melt-pool width	5	4	2	3	1
Melt-pool length	3	5	2	4	1
Melt-pool depth	5	4	2	3	1

A higher rank indicates a more relevant input; to select the best subset of k features, select the k features with the highest ranks

Table 7.2 Rank order of subsets of the input parameters to the Eagar-Tsai model using the MSE filter

	Speed	Power	Beam size	Absorptivity	Noise
Melt-pool width	5	4	2	3	1
Melt-pool length	3	5	2	4	1
Melt-pool depth	5	4	1	3	2

A higher rank indicates a more relevant input

Table 7.2 presents the results for the MSE filter. These are very similar to the CFS filter, with the exception that the beam size is ranked lower than the noise variable for the depth of the melt pool. For all three melt-pool characteristics, the three lowest ranked variables have the MSE value roughly the same, so the corresponding three variables have roughly the same order of importance.

Given these results, since the depth and width are the most important melt-pool characteristics, we decided to investigate the effects of the two most important inputs—laser power and speed—on these characteristics. While our simple simulations relate just four inputs to the melt-pool characteristics, we expect that as we move to more complex simulations, feature selection and other dimension reduction techniques will become more useful in helping us to focus on the important variables, potentially limiting the number of experiments or simulations required to create parts with desired properties.

7.3.1.3 Data-Driven Predictive Modeling

The simulation inputs and outputs can also be used to build a data-driven predictive model that can be used to predict the output values for a given set of inputs. A simple predictive model is a regression tree [1], which is similar to a decision tree, but with a continuous instead of a discrete output.

A regression tree is a structure that is either a leaf, indicating a continuous value, or a decision node that specifies some test to be carried out on a feature, with a branch and sub-tree for each possible outcome of the test. If the feature is continuous, there are two branches, depending on whether the condition being tested is satisfied or not. The decision at each node of the tree is made to reveal the structure in the data. Regression trees tend to be relatively simple to implement, yield results that can be interpreted, and have built-in dimension reduction.

Regression algorithms typically have two phases. In the training phase, the algorithm is “trained” by presenting it with a set of examples with known output values. In the test phase, the model created in the training phase is tested to determine how accurately it performs in predicting the output for known examples. If the results meet expected accuracy, the model can be put into operation to predict the output for a sample point, given its inputs.


The test at each node of a regression tree is determined by examining each feature and finding the split that optimizes an impurity measure. We use the mean-squared error, MSE, as defined in Sect. 7.3.1.2, as the impurity measure. The split at each node of the tree is chosen as the one that minimizes MSE across all features for the samples at that node.

To avoid splitting the tree too finely, we stop the splitting if the number of samples at a node is less than 10 or the standard deviation of the values of the output variable at a node has dropped below 10% of the standard deviation of the output variable of the original data set.

The regression tree acts as a surrogate for the data from the E-T simulations and can be used to predict the width, depth and length of the melt pool for a given set of inputs. The inputs for a sample point are used to traverse the tree, following the decision at each node, until a leaf node is reached; the predicted value assigned to the sample is the mean of the output values of the training data that end up at that leaf node.

Figure 7.3 shows the melt-pool depth for the E-T simulations predicted by the regression tree vs. the actual depth from the simulations. The predicted value for each sample point was obtained by creating a regression tree with all other sample points and using it to predict the melt-pool depth for the given sample point. The

Fig. 7.3 Plot of predicted versus actual melt-pool depth (in micron). The predicted value for each sample point in the E-T simulations was obtained using a regression tree built with the rest of the sample points. The actual depth is obtained from the simulations. The *blue line* is the $y = x$ curve



percentage deviation for the entire data set is 11.2%. This is obtained by taking the average over all sample points of the absolute value of the ratio of the residual to the actual value. The residual is the difference between the actual and predicted values.

The accuracy of the regression tree depends on the number and location of the sample points, as well as the complexity of the function being modeled. If there are too few sample points, or they are in the wrong location, then the prediction can be poor, especially if the function being predicted is quite complex. For our set of simulations, the accuracy obtained is reasonable, though it could be improved further by adding new sample points in appropriate locations or by using ensembles of regression trees.

In comparison with the E-T simulations, where each simulation takes ≈ 1 min on a laptop, it takes a few micro-seconds to build the regression tree surrogate from the 462 simulations and practically no time to generate the melt-pool depth for a set of input variables using the surrogate.


7.3.2 Using Simple Experiments to Evaluate Simulation Results

We next considered some simple single-track experiments [20] to evaluate the findings from our simulations. In these experiments, a single layer of powder is spread on a plate and a single track created at a specific laser power and speed. The powder is then removed, and the plate cut so that the cross-section of the track can be obtained and the melt-pool characteristics measured.

Based on prior work, we had decided to use a powder layer thickness of $30\ \mu\text{m}$ as this had resulted in the highest density in experiments with 316L powder [21]. The layer thickness is the amount by which the build platform is lowered in each layer of the build. Since the powder is porous, its height decreases when it melts. Therefore, the next layer of powder has a depth greater than the set value of layer thickness. Due to the shrinkage on melting, the initial layers of powder are progressively deeper, until the thickness of the powder reaches a steady state that is determined by the amount of shrinkage of the powder on melting.

When we translate the results of the E-T model to single-track experiments, we need to account for the fact that the simulations are just an approximation and there is no powder considered in the model. So the melt-pool depth from the E-T model should be sufficiently large compared with the thickness of the powder in the experiment to ensure that the substrate melts as well. We therefore focused on the simulations that gave a melt-pool depth of two to three-times the set layer thickness. Note that this factor is just an approximation that helps us to constrain the range of parameters. In addition to avoiding process parameters that resulted in relatively shallow melt pools, we also wanted to avoid those that gave very deep melt pools. Not only would this have been wasteful, but a high energy density would have resulted in the process going from conduction-mode melting to keyhole-mode melting, resulting in voids that would have introduced porosity into the part [12].

Fig. 7.4 The 40 mm × 40 mm tilted build plate with the 14 tracks, each generated using a different value of laser power and scan speed, as listed in Table 7.3, where track 1 corresponds to the track at the top of the plate. The layer thickness is near zero at the left edge of the plate, increasing linearly to 200 μm at the right edge. The plate is cut vertically to analyze the melt-pool cross-section at a specific layer thickness



Using the results from the E-T simulations, we identified fourteen power and speed combinations that we used to create tracks on a tilted plate as shown in Fig. 7.4. This 40 mm × 40 mm build plate has a tilt so that the layer thickness is 0 at the left and 200 μm at the right, enabling us to evaluate the effect of the process parameters at different layer thicknesses.

Table 7.3 presents the melt-pool characteristics at a layer thickness of 30 μm. The results for the melt-pool depth are very consistent, with higher laser powers and

Table 7.3 The melt-pool width, height, and depth for the 14 tracks, along with the laser power and scan speed values

Track number	Power (W)	Speed (mm/s)	Width (μm)	Height (μm)	Depth (μm)
1	400	1800	112	32	105
2	400	1500	103	79	119
3	400	1200	83	28	182
4	300	1800	94	57	65
5	300	1500	83	35	94
6	300	1200	111	76	114
7	300	800	118	54	175
8	200	1500	84	26	57
9	200	1200	104	45	68
10	200	800	123	24	116
11	200	500	121	61	195
12	150	1200	79	21	30
13	150	800	109	44	67
14	150	500	115	40	120

Track 1 corresponds to the track at the top of the plate in Fig. 7.4. Powder layer thickness is 30 μm

lower speeds resulting in deeper melt pools. In addition, we observe that as the laser speed reduces, the tracks become more complete, melting more of the powder at the deeper layer thicknesses. This can be clearly seen in the three tracks at the bottom of the plate in Fig. 7.4 where, as the speed reduces from 1200 to 500 mm/s at 150 W, more of the powder melts, resulting in a complete track. These results also indicate that we have several tracks where the depth is between two and three times the layer thickness of 30 μm , making these likely process parameters for further investigation.

7.3.3 Determining Density by Building Small Pillars

Thus far we have shown how we can use simple simulations to constrain the parameter space over which we perform simple experiments. These simple experiments, in turn, enable us to constrain the space over which we perform more complex experiments, which, in our case, are building small pillars to evaluate the density.

There are many factors that control the density of an additively-manufactured part. Factors such as the laser power, speed, beam size, and powder layer thickness, control the density locally. So, we select the values of these parameters to ensure that the powder melts and sticks to the substrate, without leaving any un-melted powder particles that may lead to porosity. There are other factors that could also introduce porosity, such as the scan-line or hatch spacing, which controls the distance between adjacent scan lines, and the scanning strategy. For example, if adjacent scans do not overlap sufficiently, powder will accumulate in the space in-between the tracks, potentially causing porosity in the part if the laser parameters are not sufficient to melt this powder in subsequent scans. The use of island scanning could also result in porosity. Here, instead of creating each layer with a series of continuous scans, the region is divided into small “islands” that are scanned randomly [21, 22]. To ensure that the islands are connected, that is, there are no gaps created in-between adjacent islands, each island is scanned such that the scan vectors slightly overlap the surrounding islands. If the amount of overlap is set too small, this could introduce porosity in the part.

To identify process parameters that would result in high-density parts, we built small pillars, 10 mm \times 10 mm \times 8 mm high using a variety of power and speed combinations. We used island scanning, with 5 mm \times 5 mm islands. All other parameters were set to the default, as summarized in our earlier work [10]. The power and speed values for our initial set of twenty-four pillars were chosen based on the results from the single track experiments. We then evaluated the density of these pillars using the Archimedes method. Based on the results, we built another set of twenty-four pillars at the same power values as the first set, but with the speed values chosen to complete the density curves.

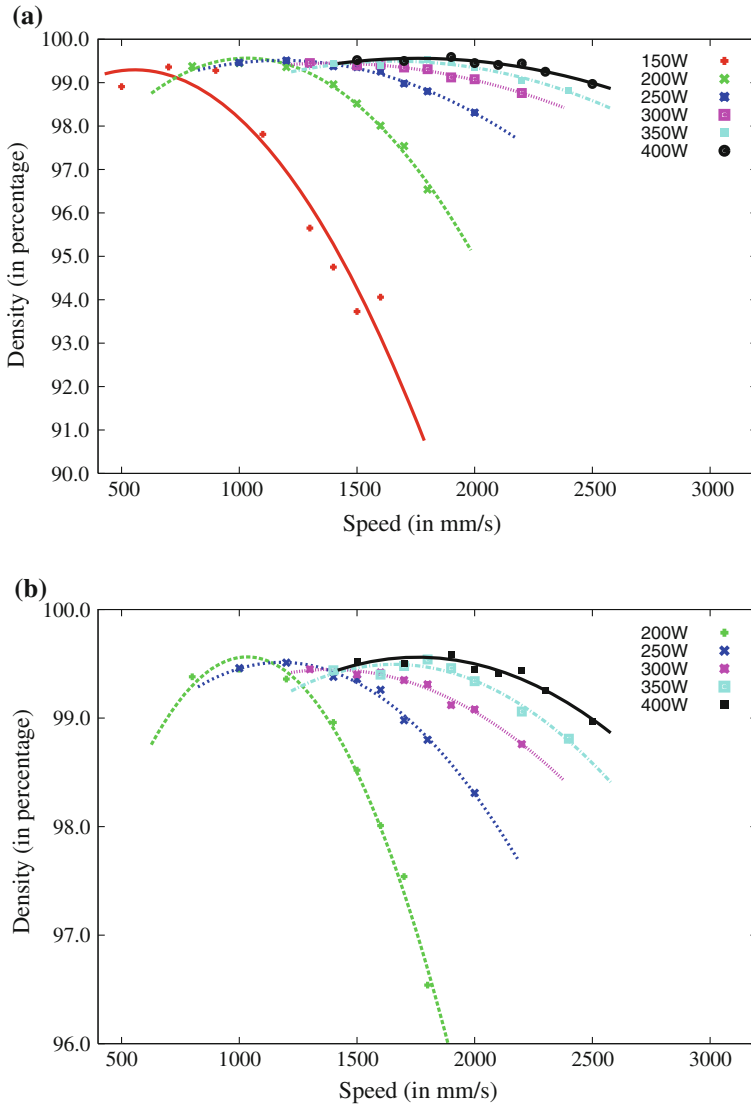


Fig. 7.5 Relative density as a function of laser power and scan speed. Plot (b) excludes the values for power = 150 W to illustrate the variation at high density. A quadratic function is fitted to the points for each power value. **a** Density for 48 316L pillars using CL powder; power 150–400 W. **b** Density for 40 316L pillars using CL powder; power 200–400 W

7.4 Experimental Results

Figure 7.5 shows the relative density of the forty-eight pillars for a range of power and speed values. We make several observations. First, we were able to use our approach to create pillars with >99% relative density for power values ranging from 150 to 400 W. Second, as expected, we found that for a given power value, increasing the speed leads to insufficient melting and lower density. The density also reduces at low speed due to voids resulting from keyhole mode laser melting; this reduction is however not as large as the reduction due to insufficient melting. Finally, we found that at higher powers, the density is high over a wider range of scan speeds, unlike at lower powers. This indicates that higher powers could provide greater flexibility in choosing process parameters that optimize various properties of a manufactured part. However, it remains to be seen if operating at higher powers will have other negative effects on the microstructure or mechanical properties of a part.

7.5 Summary

In this chapter, we showed how we can use techniques from statistics and data mining to reduce the time and cost of determining process parameters that lead to high-density, additively-manufactured parts. Specifically, we used design of computational experiments to understand the design space of input parameters using simple simulations, feature selection to identify important inputs, and data-driven surrogates for predictive modeling. We then built small pillars at various combinations of laser power and speed. Our experiences with 316L stainless steel and other materials indicate that our approach is a viable and cost-effective alternative to finding optimal parameters through extensive experimentation.

Acknowledgments The author acknowledges the contributions of Wayne King (Eagar-Tsai model), Paul Alexander (operation of the Concept Laser M2), Mark Pearson and Cheryl Evans (metallographic preparation, measurement, and data reporting).

LLNL-MI-667267: This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This work was funded by the Laboratory Directed Research and Development Program at LLNL under project tracking code 13-SI-002.

References

1. L. Breiman, J.H. Friedman, R.A. Olshen, C.J. Stone, *Classification and Regression Trees* (CRC Press, Boca Raton, 1984)
2. J. Delgado, J. Ciurana, C.A. Rodriguez, Influence of process parameters on part quality and mechanical properties for DMLS and SLM with iron-based materials. *Int. J. Adv. Manuf. Technol.* **60**, 601–610 (2012)

3. T.W. Eagar, N.S. Tsai, Temperature-fields produced by traveling distributed heat-sources. *Weld. J.* **62**, S346–S355 (1983)
4. K.-T. Fang, R. Li, A. Sudjianto, *Design and Modeling for Computer Experiments* (Chapman and Hall/CRC Press, Boca Raton, 2005)
5. A.V. Gusarov, I. Yadoirtsev, Ph Bertrand, I. Smurov, Model of radiation and heat transfer in laser-powder interaction zone at selective laser melting. *J. Heat Transf.* **131**, 072101 (2009)
6. M.A. Hall, Correlation-based feature selection for discrete and numeric class machine learning. In *Proceedings of the 17th International Conference on Machine Learning* (Morgan Kaufmann, San Francisco, 2000), pp. 359–366
7. N.E. Hodge, R.M. Ferencz, J.M. Solberg, Implementation of a thermomechanical model for the simulation of selective laser melting. *Comput. Mech.* **54**, 33–51 (2014)
8. A. Inselberg, *Parallel Coordinates: Visual Multidimensional Geometry and Its Applications* (Springer, New York, 2009)
9. C. Kamath, *Scientific Data Mining: A Practical Perspective* (Society for Industrial and Applied Mathematics (SIAM), Philadelphia, 2009)
10. C. Kamath, B. El-dasher, G.F. Gallegos, W.E. King, A. Sisto, Density of additively-manufactured, 316L SS parts using laser powder-bed fusion at powers up to 400 W. *Int. J. Adv. Manuf. Technol.* **74**, 65–78 (2014)
11. S.A. Khairallah, A. Anderson, Mesoscopic simulation model of selective laser melting of stainless steel powder. *J. Mater. Process. Technol.* **214**, 2627–2636 (2014)
12. W.E. King, H.D. Barth, V.M. Castillo, G.F. Gallegos, J.W. Gibbs, D.E. Hahn, C. Kamath, A.M. Rubenchik, Observation of keyhole-mode laser melting in laser powder-bed fusion additive manufacturing. *J. Mater. Process. Technol.* **214**, 2915–2925 (2014)
13. J.P. Kruth, M. Badrossamay, E. Yasa, J. Deckers, L. Thijs, J. Van Humbeeck, Part and material properties in selective laser melting of metals. In *Proceedings of the 16th International Symposium on Electromachining (ISEM XVI)*, Shanghai, China, 2010
14. A. Laohaprapanon, P. Jeamwathanachai, M. Wongcumchang, N. Chantarapanich, S. Chantawerod, K. Sitthiseripratip, S. Wisutmethangoon, Optimal scanning condition of selective laser melting processing with stainless steel 316L powder. *Material and Manufacturing Technology II, Pts 1 and 2* (Trans Tech Publications Ltd., Stafa-Zurich, 2012), pp. 816–820
15. Y. Li, D. Gu, Parametric analysis of thermal behavior during selective laser melting additive manufacturing of aluminum alloy powder. *Mater. Des.* **63**, 856–867 (2014)
16. National Institute of Standards and Technology, *Measurement Science Roadmap for Metal-Based Additive Manufacturing*, Technical Report, 2013
17. G.W. Oehlert, *A First Course in Design and Analysis of Experiments*. W.H. Freeman (2000). <http://users.stat.umn.edu/~gary/Book.html>
18. A.B. Spierings and G. Levy, Comparison of density of stainless steel 316L parts produced with selective laser melting using different powder grades. In *Twentieth Annual International Solid Freeform Fabrication Symposium, An Additive Manufacturing Conference*, ed. by D. Bourell (University of Texas at Austin, Austin, 2009), pp. 342–353
19. I. Yadroitsev, *Selective Laser Melting: Direct Manufacturing of 3D-Objects by Selective Laser Melting of Metal Powders* (LAP Lambert Academic Publishing, 2009)
20. I. Yadroitsev, A. Gusarov, I. Yadroitsava, I. Smurov, Single track formation in selective laser melting of metal powders. *J. Mater. Process. Technol.* **210**, 1624–1631 (2010)
21. E. Yasa, Manufacturing by combining selective laser melting and selective laser erosion/laser re-melting. Ph.D. thesis, Faculty of Engineering, Department of Mechanical Engineering, Katholieke Universiteit Leuven, Heverlee (Leuven), 2011. Available from Katholieke Universiteit Leuven
22. E. Yasa, J. Deckers, J.P. Kruth, M. Rombouts, J. Luyten, Investigation of sectoral scanning in selective laser melting. In *Proceedings of the ASME 10th Biennial Conference on Engineering Systems Design and Analysis*, vol. 4 (2010), pp. 695–703

Chapter 8

Optimal Dopant Selection for Water Splitting with Cerium Oxides: Mining and Screening First Principles Data

V. Botu, A.B. Mhadeshwar, S.L. Suib and R. Ramprasad

Abstract We propose a powerful screening procedure, based on first principles computations and data analysis, to systematically identify suitable dopants in an oxide for the thermochemical water splitting process. The screening criteria are inspired by Sabatier's principle, and are based on requirements placed on the thermodynamics of the elementary steps. Ceria was chosen as the parent oxide. Among the 33 dopants across the periodic table considered, Sc, Cr, Y, Zr, Pd and La are identified to be the most promising ones. Experimental evidence exists for the enhanced activity of ceria for water splitting when doped with Sc, Cr and Zr. The surface oxygen vacancy formation energy is revealed as the primary descriptor correlating with enhanced water splitting performance, while the dopant oxidation state in turn primarily governs the surface oxygen vacancy formation energy. The proposed screening strategy can be readily extended for dopant selection in other oxides for different chemical conversion processes (e.g., CO₂ splitting, chemical looping, etc.).

V. Botu

Department of Chemical and Biomolecular Engineering, University of Connecticut,
Storrs, CT 06269, USA
e-mail: venkatesh.botu@uconn.edu

A.B. Mhadeshwar

Center for Clean Energy and Engineering, University of Connecticut,
Storrs, CT 06269, USA

A.B. Mhadeshwar

Present Address: ExxonMobil Research and Engineering, Annandale, NJ 08801, USA

S.L. Suib

Department of Chemistry, University of Connecticut, Storrs, CT 06269, USA

S.L. Suib · R. Ramprasad (✉)

Institute of Materials Science, University of Connecticut, Storrs, CT 06269, USA
e-mail: rampi@ims.uconn.edu

R. Ramprasad

Department of Materials Science and Engineering, University of Connecticut,
Storrs, CT 06269, USA




© Springer International Publishing Switzerland 2016

T. Lookman et al. (eds.), *Information Science for Materials
Discovery and Design*, Springer Series in Materials Science 225,
DOI 10.1007/978-3-319-23871-5_8

8.1 Introduction

Utilizing dopants to optimize, enhance, or fundamentally change the behavior of a parent material has been exploited in many situations ranging from material strengthening to electronics to electrochemistry. The search and identification of suitable dopant candidates has been laborious though, and dominated either by lengthy trial-and-error strategies (guided by intuition) or plain serendipity. We are entering an era where such Edisonian approaches are gradually being augmented (and sometimes, replaced) by rational strategies based on advanced computational screening [1]. Often these strategies rely on first principles methods, that provide a reasonably accurate description of the underlying chemistry [2–4]. More recently, it has been shown that supplementing first principles investigations with data-driven approaches can help identify meaningful correlations within the data [5–13]. In the present contribution, we offer such a prescription for the selection of suitable dopants within cerium oxides in order to enhance the thermochemical splitting of water.

Complete gas phase thermolysis of water is highly endothermic ($\Delta H = +2.53$ eV) requiring temperatures in excess of 4000 K to be thermodynamically favorable, making such reactions unviable for H_2 synthesis [14, 15]. On the other hand, partial thermolysis via a multistep process in the presence of MO catalysts provides an attractive practical alternative [15, 16]. The latter approach is performed at two distinct temperatures (both well below 4000 K): a high-temperature (≈ 2200 K) reduction step that involves creation of O vacancies in the MO (and the consequent evolution of O_2 gas), and lower-temperature (≈ 900 K) oxidation steps in the presence of steam, which lead to the filling up of O vacancy centers (resulting in the evolution of H_2 gas). Owing to this multistep procedure, an additional step to separate the H_2 and O_2 products is eliminated entirely. Equations (8.1)–(8.3) below represent a reordered version (for ease of subsequent discussion) of the multiple steps involved in this process.



The (s) and (g) subscripts represent solid and gas phases, respectively. Equations (8.1) and (8.2) are the low-temperature steps, with $MO-V_o$ and $MO-(H)(H)$ representing, respectively, the oxide containing an O vacancy and the oxide in which the O vacancy is filled up by a H_2O molecule (with ‘(H)(H)’ indicating that the H atoms of H_2O are adsorbed on the oxide surface). Equation (8.3) is the high-temperature activation step that leads to the creation of $MO-V_o$.

Unfortunately, several MOs require temperatures in excess of 2700 K (leading to poor H_2 production efficiencies), leaving only a subset of oxides based on Zn, Fe and Ce to be the most promising [17, 18]. Oxides of Zn and Fe are prone to sintering, phase transformation or volatility due to the proximity of the high temperature step




Fig. 8.1 Reaction pathway and energetics (*red solid line*) for the dissociation of H₂O on an undoped ceria surface. CeO₂-V_o is an oxide with a vacancy, CeO₂-(H)(H) is an oxide with vacancy filled by a H₂O molecule and CeO₂ is a stoichiometric surface. The *green dotted line* shows the minimum energy pathway for dissociation. Ce, O and H are represented by beige, red and white colors respectively

to their melting points [19]. CeO₂, on the other hand, displays high stability and high melting temperature (≈ 2600 K), and is thus overwhelmingly favored [17].

Still, the efficiency of H₂ production with CeO₂ is quite low ($< 1\%$) [18]. This low efficiency is rooted in the high temperatures (> 1900 K) required for the reduction step (8.3), related directly to the large O vacancy formation energy of CeO₂, along with other operational difficulties [18, 20]. Figure 8.1 shows the energies E₁, E₂ and E₃ of (8.1), (8.2) and (8.3), respectively, computed here using density functional theory (DFT) (details below), and helps identify the causes of the low efficiency. The dotted line indicates the uphill nature of the water splitting process. The ideal system should display E₁ and E₂ close to zero (for facile H₂ evolution at low temperatures), and small E₃ values (to alleviate the burden on the reduction step). In the case of CeO₂, E₁ is too negative and E₃ is too positive.

A pathway to circumvent these hurdles is to control the energetics of (8.1)–(8.3) individually by the introduction of dopants (although, of course, the overall energetics of H₂O splitting cannot be altered). For instance, this strategy may be used to destabilize O in CeO₂ (and thus reduce the O vacancy formation energy) [17, 21–27]. Doping CeO₂ with a plethora of elements has been explored in the recent past [28–40], and many dopants (e.g., Zr, Cr, Sc) have been shown to help significantly increase the efficiency of H₂ production by reducing the temperatures required to accomplish (8.3) [32, 34, 35]. *Nevertheless, a clear rationale for why a given dopant is desirable, and a framework for the systematic (non-Edisonian) selection of dopants is currently unavailable.* This work attempts to fill that gap. First, we propose a framework to systematically screen for dopants, based on guidelines inspired by Sabatier’s

principle, then we identify the best candidates using first principles methods, and finally we use data analysis methods, specifically feature selection, to identify the primary factors that make these dopants attractive.

8.2 Screening Framework

In the present first principles/data-driven based work, we consider a host of dopants in CeO₂, including 33 elements spanning the 4th, 5th and 6th period of the Periodic Table (specifically the *alkali*, *alkaline earth* and *d* series elements). Assuming that the energetics of (8.1)–(8.3) determine whether a dopant is favorable or not, we define the following screening criteria to be used in a successive manner:

- *Criterion 1*: $0 \leq E_3^D \leq E_3$
- *Criterion 2*: $0 \leq E_1^D \leq \delta$
- *Criterion 3*: $0 \leq E_1^D + E_2^D \leq \delta$

The superscripts D merely indicate that these are the energetics of doped ceria.

The rationale underlying this specific choice and sequence of screening criteria stems from insights derived from Sabatier’s principle, and may be understood as follows (cf. Fig. 8.1). *Criterion 1* merely states that the O vacancy formation energy (which is what E_3^D represents) should not be too small to prevent further water dissociation nor too large (certainly not larger than that of undoped ceria (E_3)) to mandate higher activation temperatures. This criterion is listed first because E_3^D appears to most strongly control the temperature requirement of the costly high-temperature step, and also because E_3^D is the easiest quantity to compute (as it does not involve the H₂O species at all). *Criterion 2* states that E_1^D should also be bracketed, but by a smaller range. Noting that overall dissociation of water for undoped ceria is too negative (see Fig. 8.1), thus potentially adding an energy penalty to subsequent steps, we generously allow δ to be 1.5 eV, which is a reasonable choice considering energy uncertainties within DFT and the neglect of entropy. *Criterion 3* is specific to thermochemical water splitting and bounds the overall oxidation process within δ , ensuring that E_1^D or E_2^D occur at a lower temperature compared to E_3^D . In the case where this no longer holds, the process fails to fall within the realm of thermochemical water splitting.

8.3 First Principles Studies

8.3.1 Methods and Models

To measure the thermodynamic quantity, E_i^D , where i is 8.1, 8.2 or 8.3, DFT calculations were performed using the VASP code with the semi-local Perdew-Burke-Ernzerhof (PBE) exchange-correlation functional and a cutoff energy of 400 eV to

accurately treat the valence O 2s, 2p and Ce 5s, 5p, 4f, 5d, 6s states [41–43]. The electron-core interactions were captured by projector-augmented (PAW) potentials, and all calculations were spin polarized to ensure the true electronic state of O and reduced Ce was captured [44]. The computed lattice parameter of bulk CeO₂ (5.47 Å) is in good agreement with the corresponding experimental value (5.41 Å) [38]. A 96-atom bulk 2×2×2 supercell model and a 60-atom (2×2) surface model (5 O-Ce-O trilayers) cleaved along the (111) plane were used in all calculations. The bottom 3 trilayers of the slab were fixed to recover the bulk nature of the material, and a vacuum of 15 Å along the c axis ensured minimal spurious interactions between periodic images. A Γ -centered k -point mesh of 3×3×3 and 3×3×1 were used for the bulk and surface calculations, respectively. The Hubbard (U) correction was not applied as no universal U value captures the true electronic state of all elements. Also, given that we consider a dilute vacancy limit, the effect of electron localization is insignificant as shown previously [45, 46].

8.3.2 Enforcing the 3-Step Criteria

Dopants were introduced by replacing a single Ce atom at the center of the bulk model and at the 1st trilayer of the surface model. Our analysis indicated that the majority of the dopants favored the surface site to the bulk by ≈ 0.3 eV. Upon exploring the local coordination environment, a surface dopant was found to be 6-fold coordinated whereas a bulk dopant was 8-fold coordinated. Given the preference of a surface site, all dopants are assumed to occupy the surface unless specified otherwise.

The primary effect of introducing dopants is to induce a local perturbation to disrupt bonding between the metallic and O atoms, thereby altering its ability to form surface O vacancies, as measured by E_3^D (cf. Fig. 8.1), computed here as

$$E_3^D = E_{\text{CeO}_2\text{-V}_o}^D - E_{\text{CeO}_2}^D + \frac{1}{2}\mu_{\text{O}_2} \quad (8.4)$$

where $E_{\text{CeO}_2\text{-V}_o}^D$ and $E_{\text{CeO}_2}^D$ are, respectively, the DFT energies of a doped surface with and without an O vacancy, and μ_{O_2} is the chemical potential of O, taken here to be the DFT energy of an isolated O₂ molecule. In all cases, the O vacancy is created adjacent to the dopant. Figure 8.2 shows E_3^D for various choices of the dopants, with the dot-dashed horizontal line indicating the corresponding value for the undoped case. Dopants adopting a low valence state compared to Ce (e.g., alkali, alkaline earth and late transition series metals) display low O vacancy formation energy, consistent with the observed high O₂ yield by ceria doped with Mn, Fe, Ni and Cu [47]. Conversely, dopants adopting a similar or higher valence state than Ce lead to high E_3^D values (e.g., Mo, Tc, and Ta). These trends are not entirely surprising, and have been noted before in CeO₂ as well as BaTiO₃ [48–50].




Fig. 8.2 Oxygen vacancy formation energy (E_3^D) of doped ceria with elements from the (a) 4th, (b) 5th and (c) 6th period of the Periodic Table. *Dot-dashed maroon line* indicates E_3^D for undoped ceria. *Light green region* indicates dopants that survived *Criterion 1*, while \star identifies dopants that survived the 3 screening criteria

E_1^D helps assess the impact of dopants on the dissociative adsorption of water on the doped surface, and is computed as

$$E_1^D = E_{\text{CeO}_2\text{-(H)(H)}}^D - E_{\text{CeO}_2\text{-V}_o}^D - \mu_{\text{H}_2\text{O}} \quad (8.5)$$

where $E_{\text{CeO}_2\text{-(H)(H)}}^D$ is the DFT energy of a doped surface upon the dissociative adsorption of water at the vacancy site. Upon dissociation, OH fills the vacancy site, while H has two possible adsorption sites; atop an adjacent O or a dopant atom. Interestingly, dopants exhibiting spontaneous vacancy formation ($E_3^D < 0$ eV) fail to accommodate a H atop a dopant, while those dopants that do facilitate H atop a dopant have an alternative lower energy pathway for dissociation. $\mu_{\text{H}_2\text{O}}$ is the chemical potential of water, taken here to be the DFT energy of an isolated H_2O molecule.

With E_1^D and E_3^D at hand (and E_2^D given by $\Delta H - E_1^D - E_3^D$), a plot that is equivalent to Fig. 8.1 but for the case of doped ceria surfaces is shown in Fig. 8.3. We now enforce *Criterion 1*, namely, $0 \leq E_3^D \leq E_3$, with $E_3 = 3.3$ eV (this value is consistent with past work [45]). Of the 33 dopants originally considered, 19 dopants (Sc, Ti, V, Cr, Mn, Co, Y, Zr, Nb, Ru, Rh, Pd, La, Hf, Re, Os, Ir, Pt and Au) satisfy this criterion (given by the dopants within the shaded region in Fig. 8.2). *Criterion 1* picks out those dopants that alter the surface reducibility in just the appropriate manner.

Next, we enforce *Criterion 2*, namely, $0 \leq E_1^D \leq \delta$, with $\delta = 1.5$ eV, on the 19 dopants that pass *Criterion 1*, resulting in the selection of Sc, V, Cr, Co, Y, Zr, Pd, La, Hf and Au. Lastly, enforcing *Criterion 3* on the 10 dopants results in the down selection of 4 promising candidates (Sc, Cr, Zr and La). Inspection of Fig. 8.3 shows that Pd and Y, although they do not pass *Criterion 3*, can be viewed as ‘near misses’. These are hence included in our final list of favored candidates. Figure 8.4 summarizes the list of dopants that passed each stage of the screening process. The 6 dopants identified, namely, Sc, Cr, Zr, La, Pd and Y, lead to desired energetic profiles, with E_1^D and E_2^D low enough to allow for reasonable H_2O dissociation yields at moderate temperatures, and with E_3^D significantly smaller than undoped ceria allowing for low reduction temperatures (c.f., Fig. 8.3). Dopants such as Mn, Fe, Ni, Cu, Sr, Ag, and

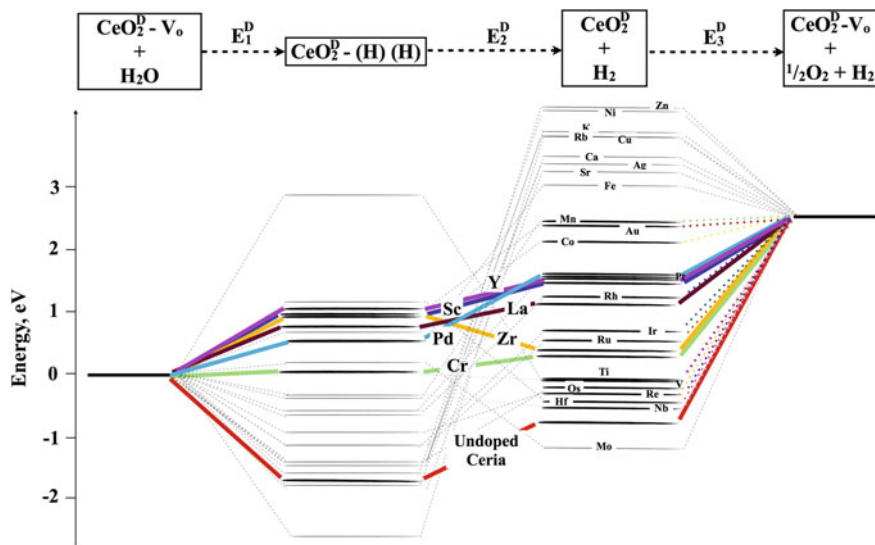


Fig. 8.3 Reaction pathway and energetics for the multistep thermochemical splitting of H₂O on a doped ceria surface. CeO₂^D-V_o is a doped surface with vacancy, CeO₂^D-(H)(H) is a doped surface with vacancy filled by a H₂O molecule and CeO₂^D is a doped stoichiometric surface. Color *solid lines* identify the 4 promising dopants and undoped CeO₂. *Grey dashed lines* identifies the non feasible dopants, while partly colored and *greyed dashed lines* identifies dopants that pass Criterion 1

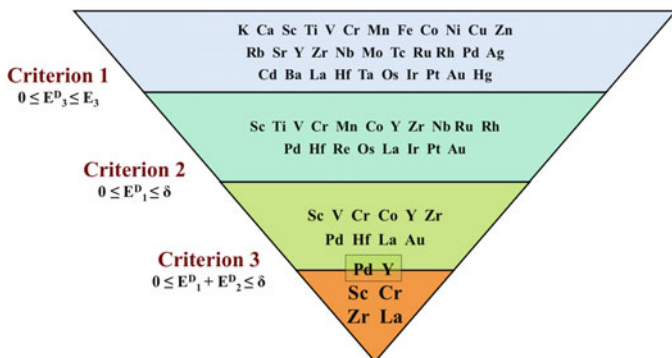


Fig. 8.4 A hierarchical chart showing the list of dopants before and after each stage of the screening process. Sc, Cr, Zr and La were identified as the promising dopant elements, whilst Pd and Y can be viewed as the near miss cases

Ca, which display small or negative E_3^D , do not pass our tests. Although low E_3^D values imply facile surface reduction (this is in fact what is observed experimentally for Mn and Fe) [47], such a tendency would not be appropriate for the multistep thermochemical water splitting process targeted here (lower yields were observed for Ni, Cu and Fe doped CeO₂ compared to undoped CeO₂) [28]. *Criterion 1*, as mentioned above, is imposed precisely to eliminate such candidates. However, dopants that lead

to small or negative E_3^D may be appropriate for photocatalytic water splitting which require surface reduction to occur low temperatures ($\approx 300\text{K}$) [51].

Of the 6 promising dopants identified, experimental evidence exists for the enhanced performance of ceria when doped with Sc, Cr and Zr for the thermochemical water splitting process. Cr doped CeO_2 is known to lower the reduction and oxidation temperature to 750 and 350 K, respectively [35]. Zr and Sc dopants increase the H_2 yield 4-fold and almost 2-fold, respectively, with respect to the undoped situation [28, 29, 38]. Lastly, although not conclusive, La doping appears to improve H_2 yield [39, 52]. The observed performances are strong functions of the synthesis, processing and measurement details. The present work ignores such complexities, and probes only the dominant and primary chemical factors that may control performance.

Irrespective of these difficulties, such a guided screening strategy has led us to some promising candidates, shown as stars in Fig. 8.2. Clearly, the best candidates display an O vacancy formation energy in the 1–2.5 eV range, i.e., neither too high nor too low, thereby respecting Sabatier’s principle. It thus appears that the O vacancy formation energy may be used as a ‘descriptor’ of the activity of doped ceria. This conclusion is consistent with an earlier similar proposal which was based on phase boundaries in surface phase diagrams of ceria exposed to an oxygen reservoir [45].

Thus far, by relying on first principles methods we are able to recognize whether a dopant increases or decreases the O vacancy formation energy, with respect to the undoped material, followed by its corresponding impact on the dissociation of water. However, an understanding of the complex dependence of the chemical attributes of a dopant and the O vacancy formation energy is absent. In the next section, with the help of data analysis methods we attempt to understand the results of the first principles computations for the spectrum of dopants considered.

8.4 Data Analysis

The mining and extraction of information forms the core of the field of data analysis, which lies under a broader umbrella of methods known as machine learning (ML) [53]. Within data analysis a subset of methods, known as feature selection, allows us to unearth correlations between variables [10, 13, 53–56]. In the context of this work, the variables are the chemical factors characterizing a dopant and the corresponding O vacancy formation energy of doped ceria. Given the strong correlation between the O vacancy formation energy and the activity, as discussed above, by identifying the key dopant factors that contribute to the O vacancy formation energy, a more educated guess on its impact on the corresponding thermodynamic activity can be made.

In order to discover such patterns, firstly, each dopant element needs to be represented numerically by a vector of numbers (also referred to as features or fingerprint in the ML community) that uniquely identifies the dopant element. Our choice of features stems from fundamental chemical factors, that are often used to describe

elements in the periodic table. The 7 factors considered in this work are; atomic radius (AR), ionic radius (IR), covalent radius (CR), ionization energy (IE), electronegativity (EN), electron affinity (EA) and oxidation state (OS). To eliminate any bias induced by the spread of the feature values, the dataset was normalized to a mean of 0 and variance of 1. On these set of chemical factors we use two feature selection methods: (i) principal component analysis and (ii) random forests, to narrow down the dominant factors that govern the descriptor (O vacancy formation energy). In the sections to follow we provide a brief overview of these methods and discuss the insights gained. We refer the readers to [53, 57–61] for a more exhaustive description. The data analysis routines used were implemented within the MATLAB statistical toolkit and Scikit-learn python module [62, 63].

8.4.1 Principal Component Analysis

Principal component analysis (PCA) is a common dimensionality reduction technique, often used to identify the dominant subset of features from a larger pool. By transforming the original features into uncorrelated and orthogonal pseudo variables, that are a linear combination of the original features (as done in this work, although non-linear combinations have been recently developed), it allows us to pin point the dominant contributions [10, 55–58]. The new transformed variables are referred to as *principal components* (PCs), which are solutions to the eigen-transformation of the covariance matrix. As with any eigen-transformation problem, the eigenvalues and eigenvectors play a critical role. The eigenvalue of a PC indicates the % of variance captured within the original dataset, whilst the eigenvector provides the coefficients that dictate the linear transformation. We shall make use of this information to down select the dominant chemical factors of a dopant.

First, we plot the transformation coefficient values of the 7 features for the first and second PCs in Fig. 8.5a. Such a plot is referred to as the loadings plot, in which correlated features cluster together. Only the first and second PCs are used as it captures $\approx 80\%$ of the variance within the original dataset (c.f., inset of Fig. 8.5a). Clearly, the dopant's OS is strongly correlated with the O vacancy formation energy. The CR, AR, IE and EN are close to orthogonal to the O vacancy formation energy, suggesting a negligible contribution to the descriptor. On the other hand, the IR and EA are not truly orthogonal, thus their contribution towards the descriptor cannot be ignored. Another interesting phenomena is the congregation of subsets of the 7 features. This isn't entirely surprising, as one would recognize that the AR, CR are similar quantities, and their grouping in the loadings plot further validates this notion. Similarly, the IE and EN group together and appear negatively correlated to the AR and CR, given their $\approx 180^\circ$ separation. By looking at the relative position of all the features in Fig. 8.5a, we can conclude that of the original 7 features considered only 3 are important; OS, IR and EA, in governing the O vacancy formation energy.

Next, we use the linear transformation coefficients of the PCs to transform the original dopant dataset (also referred to as the scores plot) and plot the first and

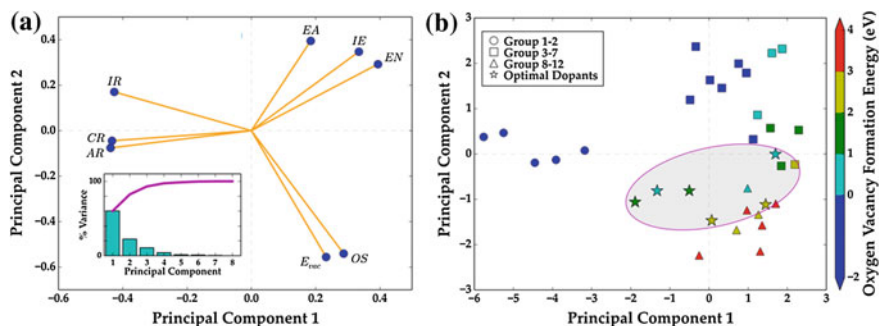


Fig. 8.5 **a** PCA loadings plot showing the correlated dopant features. The features are; atomic radius (AR), ionic radius (IR), covalent radius (CR), ionization energy (IE), electronegativity (EN), electron affinity (EA) and oxidation state (OS). E_{vac} is the O vacancy formation energy. The inset shows the % contribution of each PC to the variance in the dataset. The oxidation state (OS) is the dominant feature governing the O vacancy formation energy. **b** PCA scores plot for the first and second principal components. The dopant elements group together based on their features and the O vacancy formation energy. \star represents the final 6 dopants after the 3 step screening processes. The 6 dopants occupy a sub-space of the scores plot as highlighted by the *grey region*

second PCs in Fig. 8.5b. Each dopant element in Fig. 8.5b has further been classified according to its relative location in the periodic table (as indicated by the different marker type) and the corresponding O vacancy formation energy (marker fill color). Firstly, dopants of similar type, groups 1–2, 3–7 and 8–12 can be seen aggregating together. In particular, dopants that adopt a low valence state lie predominantly in the top/left quadrants, whilst the high valence dopants lie in the bottom/right quadrants, giving rise to an increasing O vacancy formation energy in the direction of the bottom right quadrant. Not surprisingly, amongst the low valence dopants, the alkali and alkaline earth metals further segregate from the late transition series metals, based on their differences of atomic size, amongst others. Now, upon highlighting the location of the 6 promising candidates (Sc, Cr, Y, Zr, Pd and La), as indicated by the stars, they can be seen to occupy only a small subspace of the plot (highlighted by the grey region of Fig. 8.5b). This suggests that in the high dimensional transformation these elements have similar traits, and equivalently a similar thermodynamic activity. Therefore, if one could identify other possible dopants that populate the grey region in Fig. 8.5b, we can further extend the chemical space to achieve improved water dissociation.

8.4.2 Random Forest

Another important class of feature selection algorithms are random forests (RF). Unlike PCA, random forests work by constructing a regression (or classification) model first, in this case between the 7 features and the O vacancy formation energy,




Fig. 8.6 Relative feature importance arranged in descending order for the developed RF model. The features are; atomic radius (AR), ionic radius (IR), covalent radius (CR), ionization energy (IE), electronegativity (EN), electron affinity (EA) and oxidation state (OS). E_{vac} is the O vacancy formation energy. The inset shows a parity plot, comparing the density functional theory (DFT) and RF predicted O vacancy formation energy (E_{vac}). The regression model has an R^2 value of 0.94. The oxidation state (OS) is the dominant feature governing the O vacancy formation energy

following which the important features are then extracted as a by-product. The framework is built upon an ensemble of individual regression models, also known as decision trees [53, 59–61]. The prediction of each individual tree is then averaged across the ensemble, resulting in the final or true predicted value. Given our limited dataset size (based on 33 dopant elements), we selected a 75% split for training, with the remaining kept aside as validation/testing. Each decision tree in the model is then trained on a subset of the original training dataset, a procedure known as bootstrapping. The combination of bootstrapping and ensemble averaging makes RF models robust and devoid of overfitting, a common issue in ML. We generate a forest of 250 trees, based on the 7 dopant features described earlier and the O vacancy formation energy. The final regression model we obtained has an R^2 value of 0.95 (c.f., inset Fig. 8.6), suggesting a good fit. Then by using mean decrease impurity metric, we estimate the relative importance of each feature in the regression model [61].

In Fig. 8.6, we plot the relative importance of the 7 features in descending order. Clearly, the role of a dopant's OS supersedes all others. This observation is consistent with the PCA analysis above. Also, it can be seen that IR and EA rank 2nd and 3rd in feature importance in the regression model, once again suggesting a small contribution towards the descriptor.

Both the PCA and RF methods result in similar conclusions, leading us to believe that the dopant's OS primarily governs the role of the descriptor, i.e., O vacancy formation energy, followed by a much smaller contribution of the IR and the EA. Upon revisiting the OS of the 6 promising dopants, they adopt either a +3 or +4 state. Therefore as a first measure, by understanding the coordination environment of

the dopant within the surface one can hazard a reasonable guess on its corresponding impact on the O vacancy formation energy. Even though many other elements such as Ti, V, Mn, Fe, Nb, Mo, Tc, Ru, Rh, Hf, Ta, Os, Ir adopt a similar OS state, the combination of the OS, IR and EA skews them out of the optimal regime.

8.5 Summary and Outlook

In this work, we considered a host of dopants in cerium oxide, that span the 4th, 5th and 6th period (specifically the *alkali*, *alkaline earth* and *d* series elements) of the Periodic Table, in order to understand the impact on the dissociation of water. Using a screening framework based on a first principles strategy augmented with data analysis methods, we successfully identified 6 promising dopants (Sc, Cr, Y, Zr, Pd and La), consistent with past experimental results, that are worthy of further inquiry. A dopant's oxidation state, ionic radius and electron affinity are found to be the dominant chemical factors that primarily govern the oxygen vacancy formation energy, which in turn governs the activity. The overall framework, we believe, can be easily extended for dopant selection in ceria and other oxides as well as for different chemical conversion processes (e.g., thermochemical CO₂ splitting, chemical looping, etc.).

Nevertheless, some open questions remain on the true measure of activity. First, kinetic factors, such as activation barriers, have been completely ignored in the present work. All the screening criteria were based on the thermodynamic requirements of the elementary steps, and serve as necessary but not sufficient conditions. Second, it is unclear what the impact of non-zero temperatures and gas phase component pressures would be on the computed quantities and final outcomes. Preliminary assessment based on first principles thermodynamics indicates that our main conclusions will be largely unchanged even when such factors are accounted for. However, by incorporating more of such metrics, along with the guidelines from the data analysis methods, we can systematically refine the screening framework.

Acknowledgments This work was supported financially by a grant from the National Science Foundation. Partial computational support through a National Science Foundation Teragrid allocation is also gratefully acknowledged.

References

1. G. Ceder, K. Persson, The stuff of dreams. *Sci. Am.* **309**, 36 (2013)
2. J. Neugebauer, T. Hickel, Density functional theory in materials science. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **3**(5), 438–448 (2013)
3. G. Hautier, A. Jain, S.P. Ong, From the computer to the laboratory: materials discovery and design using first-principles calculations. *J. Mater. Sci.* **47**, 7317 (2012)

4. A.D. Becke, Perspective: fifty years of density-functional theory in chemical physics. *J. Chem. Phys.* **140**, 18A301 (2014)
5. T. Mueller, A.G. Kusne, R. Ramprasad, Machine learning in materials science: recent progress and emerging applications, in *Reviews in Computational Chemistry*, ed. by A.L. Parrill and K.B. Lipkowitz (Wiley, New York, 2016)
6. S. Srinivas, K. Rajan, property phase diagrams for compound semiconductors through data mining. *Materials* **6**, 279 (2013)
7. G. Hautier, C.C. Fisher, A. Jain, T. Mueller, G. Ceder, Finding natures missing ternary oxide compounds using machine learning and density functional theory. *Chem. Mater.* **22**, 3762 (2010)
8. C.C. Fischer, K.J. Tibbetts, D. Morgan, G. Ceder, Predicting crystal structure by merging data mining with quantum mechanics. *Nat. Mat.* **5**, 641 (2006)
9. X. Zhang, L. Yu, A. Zakutayev, A. Zunger, Sorting stable versus unstable hypothetical compounds: the case of multi-functional abx half-heusler filled tetrahedral structures. *Adv. Funct. Mater.* **22**, 1425 (2012)
10. P.V. Balachandran, S.R. Broderick, K. Rajan, Identifying the 'inorganic gene' for high-temperature piezoelectric perovskites through statistical learning. *Proc. R. Soc. A* **467**, 2271 (2011)
11. E.W. Bucholtz, C.S. Kong, K.R. Marchman, W.G. Sawyer, S.R. Phillpot, S.B. Sinnott, K. Rajan, Data-driven model for estimation of friction coefficient via informatics methods. *Tribol. Lett.* **47**, 211 (2012)
12. I.E. Castelli, K.W. Jacobsen, Designing rules and probabilistic weighting for fast materials discovery in the perovskite structure. *Model. Simul. Mater. Sci. Eng.* **22**, 055007 (2014)
13. J. Carrete, W. Li, N. Mingo, S. Wang, S. Curtarolo, Finding unprecedentedly low-thermal-conductivity half-heusler semiconductors via high-throughput materials modeling. *Phys. Rev. X* **4**, 011019 (2014)
14. D.R. Hull, H. Prophet, Janaf thermochemical tables (2014), <http://kinetics.nist.gov/janaf>. Accessed 15 Jan 2014
15. S. Abanades, P. Charvin, G. Flamant, P. Neveu, Screening of water-splitting thermochemical cycles potentially attractive for hydrogen production by concentrated solar energy. *Energy* **31**, 2805 (2006)
16. T. Nakamura, Hydrogen production from water utilizing solar heat at high temperatures. *Sol. Energy* **19**, 467 (1977)
17. S. Abanades, G. Flamant, Solar hydrogen production from the thermal splitting of methane in a high temperature solar chemical reactor. *Sol. Energy* **80**, 1611 (2006)
18. L. D'Souza, Thermochemical hydrogen production from water using reducible oxide materials: a critical review. *Mater. Renew. Sust. Energy* **2**, 1 (2013)
19. W.C. Chueh, S.M. Haile, A thermochemical study of ceria: exploiting an old material for new modes of energy conversion and CO₂ mitigation. *Philos. Trans. R. Soc. A* **368**, 3269 (2010)
20. W.C. Chueh, S.M. Haile, Ceria as a thermochemical reaction medium for selectively generating syngas or methane from H₂O and CO₂. *Chem. Sus. Chem.* **2**, 735 (2009)
21. W.C. Chueh, C. Falter, M. Abbott, D. Scipio, P. Furler, S.M. Haile, A. Steinfeld, High-flux solar-driven thermochemical dissociation of CO₂ and H₂O using nonstoichiometric ceria. *Science* **330**, 1797 (2010)
22. A. Trovarelli, *Catalysis by Ceria and Related Materials* (World Scientific, London, 2002)
23. S. Kumar, P.K. Schelling, Density functional theory study of water adsorption at reduced and stoichiometric ceria (111) surfaces. *J. Chem. Phys.* **125**, 204704 (2006)
24. H.T. Chen, Y.M. Choi, M. Liu, M.C. Lin, A theoretical study of surface reduction mechanisms of CeO₂(111) and (110) by H₂. *Chem. Phys. Chem.* **8**, 849 (2007)
25. Z. Yang, Q. Wang, S. Wei, D. Ma, Q. Sun, The effect of environment on the reaction of water on the ceria(111) surface: a DFT+U study. *J. Phys. Chem. C* **114**, 14891 (2010)
26. M. Fronzi, S. Piccinin, B. Delley, E. Traversa, C. Stampfl, Water adsorption on the stoichiometric and reduced CeO₂ (111) surface: a first-principles investigation. *Phys. Chem. Chem. Phys.* **11**, 9188 (2009)

27. M. Molinari, S.C. Parker, D.C. Sayle, M.S. Islam, Water adsorption and its effect on the stability of low index stoichiometric and reduced surfaces of ceria. *J. Phys. Chem. C* **116**, 7073 (2012)
28. Q.L. Meng, C. Lee, T. Ishihara, H. Kaneko, Y. Tamaura, Reactivity of CeO₂-based ceramics for solar hydrogen production via a two-step water-splitting cycle with concentrated solar energy. *Int. J. Hydrog. Energy* **36**, 13435 (2011)
29. C. Lee, Q. Meng, H. Kaneko, Y. Tamaura, Solar hydrogen productivity of ceria-scandia solid solution using two-step water-splitting cycle. *J. Sol. Energy Eng.* **1135**, 011062 (2013)
30. C. Lee, Q. Meng, H. Kaneko, Y. Tamaura, Dopant effect on hydrogen generation in two-step water splitting with CeO₂-ZrO₂MO_x reactive ceramics. *Int. J. Hydrog. Energy* **38**, 15934 (2013)
31. R. Bader, L.J. Venstrom, J.H. Davidson, W. Lipinski, Thermodynamic analysis of isothermal redox cycling of ceria for solar fuel production. *Energy Fuels* **27**, 5533 (2013)
32. L.J. Venstrom, N. Petkovich, S. Rudisill, A. Stein, J.H. Davidson, The effects of morphology on the oxidation of ceria by water and carbon dioxide. *J. Sol. Energy Eng.* **134**, 011005 (2012)
33. G. Hua, L. Zhang, G. Fei, M. Fang, Enhanced catalytic activity induced by defects in mesoporous ceria nanotubes. *J. Mater. Chem.* **22**, 6851 (2012)
34. J. Rossmeisl, W.G. Bessler, Trends in catalytic activity for SOFC anode materials. *Solid State Ionics* **178**, 1694 (2008)
35. P. Singh, M.S. Hegde, Ce_{0.67}Cr_{0.33}O₂: a new low-temperature O₂ evolution material and H₂ generation catalyst by thermochemical splitting of water. *Chem. Mater.* **22**, 762 (2010)
36. Y. An, M. Shen, J. Wang, Comparison of the microstructure and oxygen storage capacity modification of Ce_{0.67}. *J. Alloy Compd.* **441**, 305 (2007)
37. M. Zhao, M. Shen, X. Wen, J. Wang, Ce-Zr-Sr ternary mixed oxides structural characteristics and oxygen storage capacity. *J. Alloy Compd.* **457**, 578 (2008)
38. A.L. Gal, S. Abanades, N. Bion, T.L. Mercier, V. Harle, Reactivity of doped ceria-based mixed oxides for solar thermochemical hydrogen generation via two-step water-splitting cycles. *Energy Fuels* **27**, 6068 (2013)
39. A.L. Gal, S. Abanades, Dopant incorporation in ceria for enhanced water-splitting activity during solar thermochemical hydrogen generation. *J. Phys. Chem. C* **116**, 13516 (2012)
40. S. Abanades, A.L. Gal, CO₂ splitting by thermo-chemical looping based on Zr_xCe_{1-x}O₂ oxygen carriers for synthetic fuel generation. *Fuel* **102**, 180 (2012)
41. G. Kresse, J. Furthmüller, Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B* **54**, 11169 (1996)
42. G. Kresse, D. Joubert, From ultrasoft pseudopotentials to the projector augmented-wave method. *Phys. Rev. B* **59**, 1758 (1999)
43. J.P. Perdew, K. Burke, Y. Wang, Generalized gradient approximation for the exchange-correlation hole of a many-electron system. *Phys. Rev. B* **54**, 16533 (1996)
44. P.E. Blöchl, Projector augmented-wave method. *Phys. Rev. B* **50**, 17953 (1994)
45. V. Botu, R. Ramprasad, A.B. Mhadeshwar, Ceria in an oxygen environment: surface phase equilibria and its descriptors. *Surf. Sci.* **619**, 49 (2014)
46. M.B. Watkins, A.S. Foster, A.L. Shluger, Hydrogen cycle on CeO₂ (111) surfaces: density functional theory calculations. *J. Phys. Chem. C* **111**, 15337 (2007)
47. H. Kaneko, T. Miura, H. Ishihara, S. Taku, T. Yokoyama, H. Nakajima, Y. Tamaura, Reactive ceramics of CeO₂-MO_x (M = Mn, Fe, Ni, Cu) for H₂ generation by two-step water splitting using concentrated solar thermal energy. *Energy* **32**, 656 (2007)
48. M. Krcha, A.D. Mayernick, M.J. Janik, Periodic trends of oxygen vacancy formation and c-h bond activation over transition metal-doped CeO₂ (111) surfaces. *J. Catal.* **293**, 103 (2012)
49. Z. Hu, H. Metiu, Effects of dopants on the energy of oxygen-vacancy formation at the surface of ceria: local or global. *J. Phys. Chem. C* **115**, 17898 (2011)
50. V. Sharma, G. Pilania, G.A. Rossetti, K. Slenes, R. Ramprasad, Comprehensive examination of dopants and defects in BaTiO₃. *Phys. Rev. B* **87**, 134109 (2013)
51. D. Channei, B. Inceesungvorn, N. Wetchakun, S. Phanichphant, A. Nakaruk, P. Koshy, C.C. Sorrell, Photocatalytic activity under visible light of Fe-nanoparticles synthesized by flame spray pyrolysis. *Ceram. Int.* **39**, 3129 (2013)

52. T. Miki, T. Ogawa, M. Haneda, N. Kakuta, A. Ueno, S. Tateishi, S. Matsuura, M. Sato, Enhanced oxygen storage capacity of cerium oxides in $\text{CeO}_2/\text{La}_2\text{O}_3/\text{Al}_2\text{O}_3$ containing precious metals. *J. Phys. Chem.* **94**, 6464 (1990)
53. T. Hastie, R. Tibshirani, J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd edn. (Springer, New York, 2009)
54. I. Guyon, A. Elisseeff, An introduction to variable and feature selection. *J. Mach. Learn. Res.* **3**, 1157–1182 (2003)
55. E.W. Bucholz, C.S. Kong, K.R. Marchman, W.G. Sawyer, S.R. Phillpot, S.B. Sinnott, K. Rajan, Data-driven model for estimation of friction coefficient via informatics methods. *Tribol. Lett.* **47**(2), 211–221 (2012)
56. S.C. Sieg, C. Suh, T. Schmidt, M. Stukowski, K. Rajan, W.F. Maier, Principal component analysis of catalytic functions in the composition space of heterogeneous catalysts. *QSAR Comb. Sci.* **26**(4), 528–535 (2007)
57. J.E. Jackson, *A User's Guide to Principal Components* (Wiley, New York, 1991)
58. I.T. Jolliffe, *Principal Component Analysis* (Springer, New York, 2002)
59. J. Shotton A. Criminisi, E. Konukoglu, Decision forests for classification, regression, density estimation, manifold learning and semi-supervised learning. Technical Report 114, Microsoft Research Technical Report (2011)
60. L. Breiman, Random forests. *Mach. Learn.* **45**, 5–32 (2001)
61. L. Breiman, J. Friedman, C.J. Stone, R.A. Olshen, *Classification and Regression Trees*, The Wadsworth and Brooks-Cole statistics-probability series (Taylor & Francis, Boca Raton, 1984)
62. MATLAB, *version 8.0.0.783 (R2012b)*. The MathWorks Inc., Natick, Massachusetts (2012)
63. F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011)

Chapter 9

Toward Materials Discovery with First-Principles Datasets and Learning Methods

Isao Tanaka and Atsuto Seko

Abstract When the rule to determine the target property is known *a priori*, and the computational cost for the predictors with the DFT accuracy is not too high to cover the whole library within the practical time frame, “high throughput screening” of first principles (DFT) database is a straightforward strategy for materials discovery. Otherwise we need to adopt learning methods using predictors that can cover the whole library. The learning techniques make a model to estimate the target property, which can be used for “virtual screening” of the library. Here, we show a few examples how such techniques have been used for materials discovery.

9.1 Introduction

Historically materials discovery for a particular application was achieved by chance after lengthy trial-and-error iterations, neither by rational exploration of chemical compositional space, nor on the basis of clear design principles. The situation is changing because of the emergence of two important tools: One is the establishment of efficient first principles calculations with predictive performance. Thanks to the recent progress of computational power and techniques, a large number of density functional theory (DFT) calculations can be performed and the results are stored as big databases. Such databases are available for public uses now, such as Materials Project Database (MPD) [1], Automatic Flow of Materials Discovery Library (aflowlib) [2], and Open Quantum Materials Database(OQMD) [3]. The other important progress can be seen on techniques capable of efficient data mining. Combining DFT database and the data mining techniques, accelerated discovery of materials can be expected. Information techniques to solve chemistry problems have been

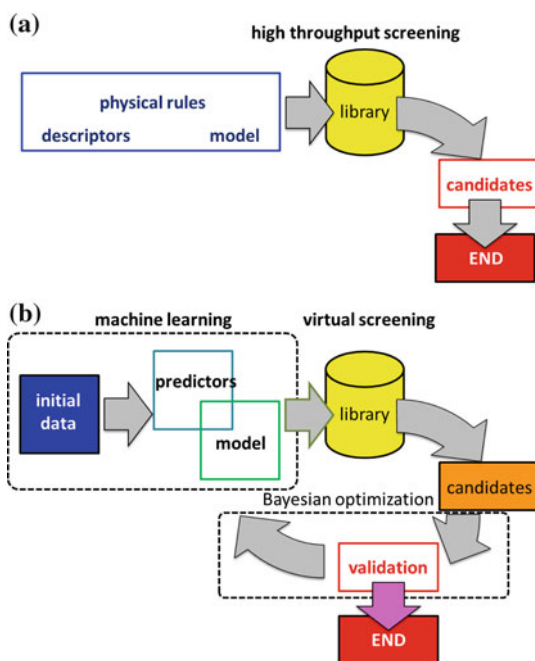
I. Tanaka (✉) · A. Seko
Department of Materials Science and Engineering, Kyoto University,
Kyoto 606-8501, Japan
e-mail: tanaka@cms.mtl.kyoto-u.ac.jp

A. Seko
e-mail: seko@cms.mtl.kyoto-u.ac.jp

called “cheminformatics” (or “chemoinformatics”). A part of the cheminformatics aiming at quantitative estimation of chemical or biological activities of chemicals from physico-chemical and structural database is called “quantitative structure-activity relationship (QSAR)” technique. The term “quantitative structure-property relationship (QSPR)” is used for similar context. Such techniques have been successful in the fields of drug discovery and organic chemistry. As for inorganic materials, however, the use of informational techniques has started just recently. This can be ascribed to the diversity of chemical elements, crystal structures and target properties of inorganic materials. Their structure-property relationships are often more complicated. Before the emergence of the DFT database, quantitative description of materials properties was very difficult.

Strategies of materials exploration with the DFT calculations should be different depending upon many factors such as (1) availability of experts knowledge as a physical or phenomenological rule, (2) abundance of experimental data, (3) computational cost for estimation of physical quantities with DFT accuracy, and (4) extent of the exploration space. Figure 9.1 shows two extreme cases of materials exploration with DFT calculations. When physical rule and descriptors for the target property are well established, and all descriptors can be easily computed by an ordinary DFT method, it is possible to perform DFT calculations of all compounds in a library in order to make “high-throughput screening” as shown in Fig. 9.1a. Candidates can be “discovered” in a straightforward manner. In another extreme case shown in Fig. 9.1b, the rule to determine the target property is not known *a priori*. Therefore, we should consider predictors that can cover the whole exploration space. Then, learning

Fig. 9.1 **a** Scheme of high-throughput screening with DFT datasets. **b** Scheme of virtual screening by combination of DFT datasets and learning methods



techniques should be used to select predictors for making a model to estimate the target property. A library can be used for “virtual screening” to find candidates. Verification process may be required to examine the predictive power of the model when “virtual screening” is made. After receiving the verification results, the model can be revised. Models and the quality of the screening can be improved iteratively through Bayesian optimization process. The virtual screening is also useful when high-throughput screening is not realistic, i.e. when the computational cost for the descriptors is too high to cover the whole library within the practical time frame. It is the same if one needs to explore too large space to cover. In this article, some of recent examples on the materials discovery with DFT dataset and learning methods are given.

9.2 High Throughput Screening of DFT Data—Cathode Materials of Lithium ion Batteries

When physics behind the target property is simple and major ingredients of the physical rule are computable by ordinary DFT methods, high throughput screening (HTS) of the DFT database is a straightforward strategy for materials discovery as shown in Fig. 9.1a. Phenomenological or empirical rules and experts knowledge can be used instead of physical rules. In such cases, selection of “good” descriptors is the critical step for the success of HTS.

HTS with DFT data has been used for materials discovery of lithium ion batteries (LIB). Ceder [4] made a pioneering work to perform HTS of cathode materials for LIB from dual viewpoints, namely charge/discharge capacity and safety. Average battery voltage of a cathode material for fully delithated (charged) and fully lithiated (discharged) conditions is given by the difference of chemical potentials between fully delithated and fully lithiated conditions. Safety of oxygen-containing cathodes (e.g., oxides, phosphates, silicates etc.) can be related to the equilibrium oxygen chemical potential of the delithated state. If the oxygen chemical potential is lower, the cathode is less susceptible to burn the coexisting electrolyte in the battery, which is expected to increase the safety of the battery. The capacity-safety diagram made by a set of DFT calculations was used for the HTS of cathode materials.

Many other properties are required for selecting cathode materials. When LIB is designed for large-scale energy storage for efficiently utilizing renewable energy power sources from solar and wind, a long cycle life is critical. This is different from batteries for portable devices. Cycle life is typically defined as the number of complete charge/discharge cycles before its capacity falls down to a certain level, say 70%, of its initial value. The target of the long-life battery is more than 70% capacity retention after 10,000 cycles. If this target is met then LIB can be practically used over 30 years with a daily charge/discharge cycle. The work by Nishijima and coworkers [5] aimed to develop cathode materials which exhibit prolonged cycle lives by substituting a range of solute elements onto the different cation sites of

the LiFePO_4 (LFP) material. LFP was chosen because it has advantages in cost, safety and cycle life among the range of LIB cathode materials [6]. When renewable energy-storage applications are considered, however, the cycle life of LFP needs to be further improved. The cycle life of battery cathodes is not a quantity that can be derived by a simple physical model. It is determined by the degradation rate during repeated charge/discharge cycles which is influenced by many different factors. The charge and discharge process for LFP proceeds via a two-phase reaction, and which inevitably produces interphase boundaries with different lattice parameters [7]. The volume change of crystalline lattice between LFP and fully delithiated FP is 6.5 % [6]. Micro-cracks are often formed due to the stress inside the LFP cathodes during the repeated charge/discharge cycle, which is widely accepted as the major degradation mechanism of the LFP cathode [8]. The degradation could therefore be retarded by reducing the volume change of the crystalline lattice during the charge/discharge cycle. Nishijima et al. [5] assumed that the relative volume change (RVC) of a compound between fully lithiated and delithiated conditions can be used as the descriptor for the cycle life. Then they explored a wide chemical compositional space in order to optimize solute atoms in LFP cathode materials for prolonged cycle life by systematic DFT calculations. Based upon the results of the screening, synthesis of selected materials was targeted. The strategy is similar to that in Fig. 9.1a, although the rule is based upon intuition or empirical knowledge.

A large set of DFT calculations were systematically made for many different kinds of solute elements that were substituted onto three possible cation sites of LFP. Co-substitution of aliovalent elements were made to maintain the charge neutrality by assuming that the formal ionic charges were unchanged. For example when Zr^{4+} and Si^{4+} were incorporated and were located respectively at Fe^{2+} and P^{5+} sites, two Si atoms and one Zr atom were put into the supercell of the DFT calculation. The situation can be expressed as $(\text{Zr}_{\text{Fe}} + 2 \text{Si}_{\text{P}})$. DFT calculations were thoroughly made for all possible solute arrangements within the unit cell composed of four formula units of LFP (i.e., 28 atoms). The lowest energy structure among them was adopted as the one representing the given chemical composition. Relative volume-change (RVC) obtained for $(\text{A}_{\text{Li}}, \text{M}_{\text{Fe}}, \text{X}_{\text{P}})$ with $\text{X} = \text{Si}$ is shown in Fig. 9.2a. RVC was defined by $100 \cdot (V_L - V_D) / V_L$ (%), where V_L and V_D denote lattice volumes of lithiated and delithiated materials, respectively. As can be seen in Fig. 9.2a, RVC is notably small when $\text{M} = \text{Zr}$. Since substitution of Li-sites by other elements reduces the battery capacity, Nishijima et al. decided to focus their efforts on the $(\text{Zr}_{\text{Fe}} + 2 \text{Si}_{\text{P}})$ -system, which was called Z2S. Its chemical formula is $\text{Li}(\text{Fe}_{1-x}\text{Zr}_x)(\text{P}_{1-2x}\text{Si}_{2x})\text{O}_4$. DFT calculations for Z2S with supercells composed of 8 and 16 formula units were additionally made, which corresponds to $x = 0.125$ and 0.0625 , respectively. Results of the RVC are shown in Fig. 9.2b. RVC decreases linearly with the solute concentration. Synthesis experiments were then performed for Z2S with varying x based on the results of HTS. By optimizing the processing parameters, single phase solid-solution samples were successfully synthesized. Structural analysis by the powder x-ray diffraction (XRD) showed that the Z2S samples were single phase up to $x = 0.125$. The sample was then supplied to electrochemical experiments. The experimental RVC is shown in Fig. 9.2b to compare with the computed values. Satisfactory




Fig. 9.2 **a** Relative volume change (RVC) between lithiated and delithiated co-substituted LFP for (A_{Li}, M_{Fe}, Si_P) by DFT calculations. **b** Comparison of experimental and DFT-RVC for Z2S samples [5]

agreements between experiments and computed results can be seen. The experimental RVC decreased linearly with x from 6.3% ($x = 0$) to 3.7% ($x = 0.125$).

Finally the cycle life performance was examined for the Z2S cathode in a laminated pouch cell using a natural graphite anode. A cell with a pristine LFP cathode was prepared for comparison. The cycle life with 80% capacity retention was 10,000 cycles for the cell with Z2S ($x = 0.050$) cathode, whereas it was 1,800 cycles for the cell with pristine cathode. The significant increase in cycle life was ascribed to the difference in the cathodes, since all other components of the cell and cell testing were the same. The cycle life for 70% capacity retention was estimated to be 25,000 cycles for Z2S ($x = 0.050$) cathode, which corresponds to the lifetime of 68 years by daily charge/discharge cycles.

HTS works using DFT database have been reported for many other applications. Curtarolo and coworkers [9] made a review article of such works and gave a list of descriptors for several problems such as nano-sintered thermoelectrics, topological insulators, non-proportionality in scintillators, and so on.

9.3 Combination of DFT Data and Machine Learning I—Melting Temperatures

In the previous section, examples of HTS with DFT data were described. As described in the Introduction, HTS is not realistic when the computational cost for the predictors with the DFT accuracy is too high to cover the whole library within the practical time frame and/or when the rule to determine the target property is not known *a priori*. Then we have two choices. One is to limit the exploration space: fixing to a certain crystal structure and limiting chemical compositions are typical strategies. The other is to adopt learning techniques using predictors that can cover the whole library. The learning techniques make a model to estimate the target property, which can be used

for “virtual screening” of a materials library. Here, we demonstrate applications of combination of DFT data and machine learning for three kinds of target properties. They are the melting temperature, the ionic conductivity in solid-state electrolytes, and the lattice thermal conductivity in thermoelectric materials.

Experimental data for inorganic substances have been well collected for thermal properties. Let us take an example of the melting temperature for which experimental data are abundant. It may be also important that melting temperature is not keen sensitive to microstructure and sub-percent level impurities. Scattering of the experimental data by different experimental groups is therefore expected to be much smaller than that for other structure or impurity-sensitive properties.

Lindemann rule [10] is often quoted as a model for explaining the melting temperature. It is based on a naive idea that melting occurs when the amplitude of thermal vibration of atoms in a substance exceeds a certain critical fraction of the interatomic distance. Although several modifications of Lindemann rule have been proposed [11–13], it is still far from predicting the melting temperature quantitatively for an arbitrary selected material. Trials to find other rules to determine melting temperatures were proposed for certain classes of materials, i.e., elemental metals [14], covalent crystals [15] and intermetallic compounds [16]. Meanwhile, a machine learning technique was applied to the prediction of the melting temperature for AB suboctet compounds [17].

Seko et al. [18] made a combined study of DFT calculations and regression techniques for prediction of the melting temperature for single and binary compounds. Experimental dataset was obtained from a standard physics and chemistry handbook [19]. Melting temperatures of the 248 compounds ranging from room temperature to 3273 K were used. The set of compounds did not contain transition metals and their compounds to avoid complexity in the DFT calculations. Two sets of predictors as shown in Table 9.1 were used for the regression. One is a set of 4 predictors, x_1 to x_4 , such as crystalline volume and cohesive energy, which were obtained by DFT calculations. DFT calculations were made for all polymorph structures that were given in the Inorganic Crystal Structure Database (ICSD). The physical properties of the lowest energy crystal structure were then adopted as predictors [19]. The other set of predictors, x_5 to x_{23} , is raw or primitive information taken from the Periodic Table and the handbook, such as atomic number, atomic mass and electronegativity. Ten variables were made to be symmetric with respect to the exchange of atomic species in binary compounds to obtain 19 predictors. Note that the sum form of the composition is always unity and it was not used as a predictor. Firstly all of these 23 predictors, which were selected without much intuition, were used for modelling by regressions. These predictors were divided into two sets. Predictor set (1) is composed only of symmetric predictors of the primitive information, x_5 to x_{23} , which does not contain information by the DFT calculations. Predictor set (2) is composed of all 23 variables, x_1 to x_{23} , including 4 variables by the DFT calculations. In order to estimate the prediction error, the data set was divided into training and test data. A randomly selected quarter of the data set and the rest of the data set were regarded as the test and training data, respectively. This was repeated 30 times and then averages of 10-fold cross-validation (CV) scores and the root-mean-square

Table 9.1 Predictors used for a model of the melting temperatures

Volume	$V (x_1)$	
Nearest-neighbor pair distance	$r_{\text{NN}} (x_2)$	
Cohesive energy	$E_{\text{coh}} (x_3)$	
Bulk modulus	$B (x_4)$	
	Sum form	Product form
Composition, c		$c_A c_B (x_5)$
Atomic number, Z	$Z_A + Z_B (x_6)$	$Z_A Z_B (x_7)$
Atomic mass, m	$m_A + m_B (x_8)$	$m_A m_B (x_9)$
Number of valence electrons, n	$n_A + n_B (x_{10})$	$n_A n_B (x_{11})$
Group, g	$g_A + g_B (x_{12})$	$g_A g_B (x_{13})$
Period, p	$p_A + p_B (x_{14})$	$p_A p_B (x_{15})$
van der Waals radius, r^{vdw}	$r_A^{\text{vdw}} + r_B^{\text{vdw}} (x_{16})$	$r_A^{\text{vdw}} r_B^{\text{vdw}} (x_{17})$
Covalent radius, r^{cov}	$r_A^{\text{cov}} + r_B^{\text{cov}} (x_{18})$	$r_A^{\text{cov}} r_B^{\text{cov}} (x_{19})$
Electronegativity, χ	$\chi_A + \chi_B (x_{20})$	$\chi_A \chi_B (x_{21})$
First ionization energy, I	$I_A + I_B (x_{22})$	$I_A I_B (x_{23})$

A set of 4 predictors, x_1 to x_4 , was obtained by DFT calculations. The other set of predictors, x_5 to x_{23} , is raw or primitive information taken from the Periodic Table and the handbook [19], which is made to be symmetric with respect to the exchange of atomic species in binary compounds [18]

(RMS) errors between predicted and experimental melting temperatures for test data were evaluated.

Figure 9.3 summarizes results by both ordinary least square regression (OLSR) and support vector regression (SVR) with predictor sets (1) and (2). CV scores and RMS errors are shown together. The figures were taken from one of the 30 trials of random divisions of the data set. The use of the predictor set (2) was found to significantly improve the model. At the same time it can be pointed out that SVR effectively reduced the error even when the predictor set without DFT results was used. Systematic deviation of the predicted values from the experimental ones can be seen for OLSR with the predictor set (1) in the high temperature region at above 1500 K. This can be ascribed to the difficulty in representing the high melting temperature simply by linear combination of 19 predictors included in the set (1). The situation was improved by the use of non-linear regression model of SVR with the same predictor set (1). The fitting of the high temperature part by OLSR was much improved when 4 additional DFT predictors were included as in the set (2).

All results shown in Fig. 9.3 were obtained with all predictors either in the set (1) or (2). Let us think of selection of “good” predictors among them. For this purpose a stepwise regression method with bidirectional elimination [20] based on the minimization of the Akaike information criterion (AIC) [21] was adopted. The best prediction model with the minimum AIC was found to be composed of 10 predictors and has a RMS error of 295 K by the OLSR, which is smaller than the RMS error by the OLSR with all of 23 predictors. Figure 9.4a shows that the RMS error decreased rapidly and almost converged at 5 predictors. The prediction model

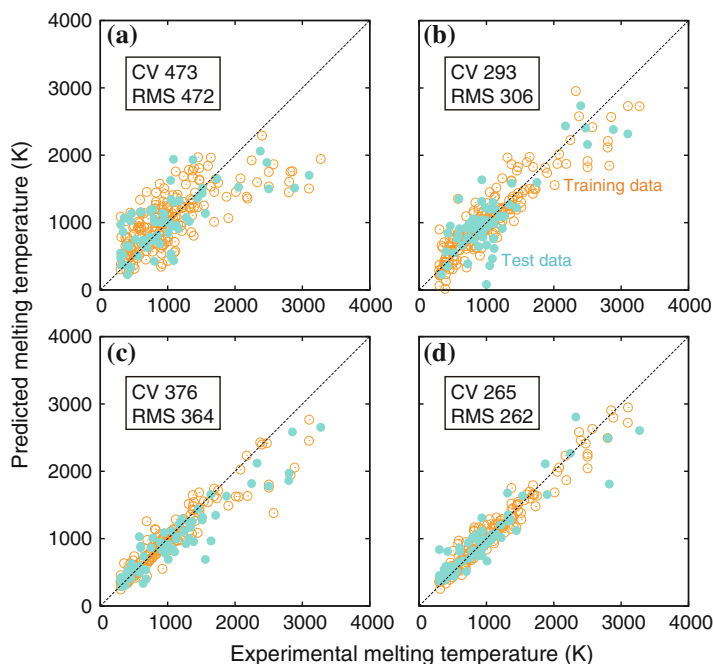


Fig. 9.3 Results by ordinary least square regression (OLSR) and support vector regression (SVR) with predictor sets (1) without DFT datasets and (2) with DFT datasets. CV scores and RMS errors in the unit of K are shown in the corresponding boxes [18]. **a** OLSR (Without DFT). **b** OLSR (With DFT). **c** SVR (Without DFT). **d** SVR (With DFT)

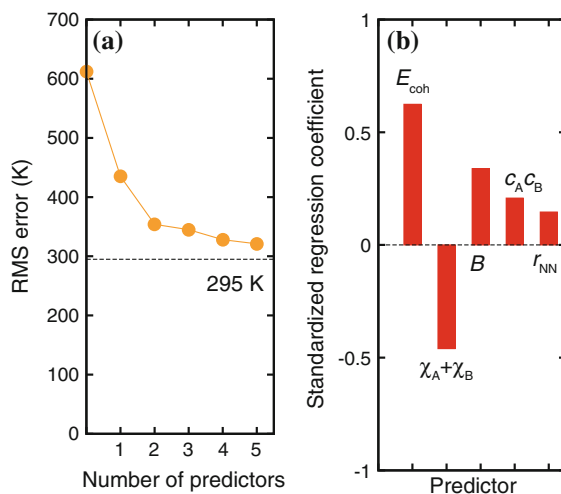


Fig. 9.4 **a** Variation of RMS error for the prediction model of the melting temperature with the number of descriptors selected according to AIC. **b** The standardized regression coefficients of the prediction model with the 5 predictors [18]

with 5 predictors showed the RMS error of 320K. The selected 5 predictors were E_{coh} , $\chi_A + \chi_B$, B , c_{ACB} , and r_{NN} . 3 of the 5 predictors were those computed by the DFT calculation. Figure 9.4b shows the standardized regression coefficients of the prediction model with the 5 predictors. The absolute value of the standardized regression coefficient for E_{coh} , which is the first selected by the stepwise regression, is the largest among the coefficients for the 5 predictors. Hence, E_{coh} contributes the most to the prediction of the melting temperature. It may sound natural to find a good correlation between the melting temperature and E_{coh} . Actually Guinea et al. [14] proposed a linear relationship model between the melting temperature and E_{coh} for metals and alloys. Recently, a linear relationship between the melting temperature and the bulk modulus, B , was also proposed by Lejaeghere et al. [22] for elemental crystals. However, the prediction only with E_{coh} provided poor prediction with the RMS error exceeding 430K for the 248 compounds in the work by Seko et al. [18]. The error was even larger only with B . The facts imply that the models only with E_{coh} or B cannot be universally applicable for predicting the melting temperature, but only useful for elemental crystals and alloys.

Once the model is made by the machine learning process, it can be supplied for virtual screening as shown in Fig. 9.1b. The process will then be followed by the Bayesian optimization procedure. Here we show an example of the optimization for finding the compound with highest melting temperature by kriging by Seko et al. [18]. Kriging was built on Gaussian processes. Figure 9.5a shows a typical situation where several sample points are available. In the kriging, a next sampling point is searched where the chance of getting beyond the current best target property is optimal. To this aim, a Bayesian regression method such as a Gaussian process is applied, and the probability distribution of the target property at all possible parameter values can

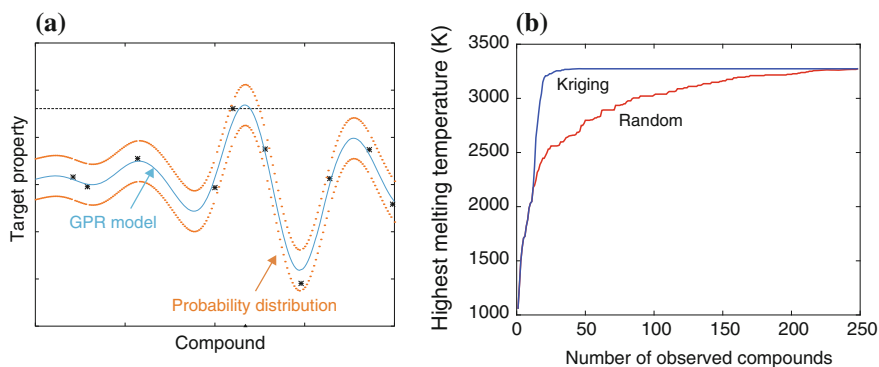


Fig. 9.5 **a** A typical situation of kriging. Gaussian process regression (GPR) is applied to the available samples shown by asterisks to make a prediction model is shown by the blue line. The probability distribution of the target property for all possible compounds is shown by orange closed circles. **b** Highest melting temperature among the observed compounds in simulations for finding the compound with the highest melting temperature based on kriging and random compound selections [18]

be obtained as illustrated in Fig. 9.5a. Then the next sampling point is determined as the one with the highest probability of improvement. Here the kriging was applied to find the compound with the highest melting temperature from a pool of compounds. The procedure can be organized as follows: (1) An initial training set is first prepared by randomly choosing compounds. (2) A compound is selected based on GPR. The compound is chosen as the one with the largest probability of getting beyond the current best value f_{best} . Since the probability is a monotonically increasing function of the z score, $z = [f(x^*) - f_{\text{best}}] / \sqrt{v(x^*)}$, the compound with the highest z score is chosen from the pool of unobserved materials. (3) The melting temperature of the selected compound is observed. (4) The selected compound is added into the training data set. Then the simulation goes back to step (2). Steps (2)–(4) are repeated until all data of melting temperatures are included in the training set.

Here the kriging of the melting temperature was started from a data set of 12 compounds. For comparison, a simulation based on the random selection of compounds was also performed. Both the kriging and random simulations were repeated 30 times and the average number of compounds required for finding the compound with the highest melting temperature was observed. Figure 9.5b shows the highest melting temperature among observed compounds during one of the 30 kriging and random trials. As can be seen in Fig. 9.5b, the compound with the highest melting temperature was found much more efficiently using the kriging. The average number of observed compounds required for finding the compounds with the highest melting temperature over 30 trials using the kriging and random compound selections were 16.1 and 133.4, respectively; hence kriging substantially improved the efficiency of discovery.

9.4 Combination of DFT Data and Machine Learning II—Lithium ion Conducting Oxides

The lithium-ion conducting oxides in the system $\text{LiO}_{1/2}\text{--AO}_{m/2}\text{--BO}_{n/2}$ (where m and n denote the formal valences of cations A and B, respectively) is known as LISICON (Lithium Super Ionic CONductors) [23]. It has general formula of $\text{Li}_{8-c}\text{A}_a\text{B}_b\text{O}_4$ (where $c = ma + nb$). Although the conducting properties of many different LISICONS have been intensively studied since the 1970s, there are still many compositions that have not been reported experimentally. In some cases, results from different groups vary considerably [24–28]. Arrhenius plots of Li-ion conductivities by previous experimental data are shown in Fig. 9.6a. The conductivity changes considerably depending on chemical compositions.

First principles molecular dynamics (FPMD) calculations can be used to estimate the atomic diffusivity. However, high computational costs hinder their use for the purpose of HTS over a wide range of materials. Typically FPMD can be done for less than 100 ps (or 10^5 MD steps), which limits the lowest accessible diffusivity by FPMD to be the order of $D = 10^{-10}$ m²/s. FPMD results alone cannot be used

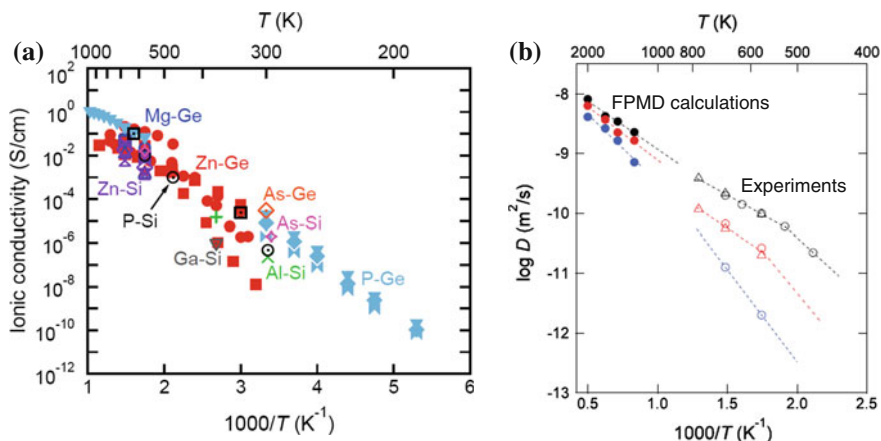


Fig. 9.6 **a** Summary of Arrhenius plots of experimental Li-ion conductivities of LISICON compounds in literature. **b** Comparison of experimental and FPMD results for $\text{Li}_{2+2x}\text{Zn}_{1-x}\text{GeO}_4$ ($x = 0.25, 0.50$ and 0.75) [29]

as predictors of lower diffusivity events. Fujimura et al. made systematic FPMD calculations of LISICON materials at above 1000 K [29]. Arrhenius plots of the calculated Li-ion diffusion coefficients by FPMD calculations for Li-ion conductors, $\text{Li}_{2+2x}\text{Zn}_{1-x}\text{GeO}_4$ ($x = 0.25, 0.50$ and 0.75) were compared with experimental data shown as open circles [24] and open triangles [30], in Fig. 9.6b. The extrapolation of the FPMD results to lower temperatures, < 800 K, where the FPMD were not practical, showed satisfactory agreements with the FPMD results for all three compositions. At the same time, one can point out the presence of deflection points in the experimental conductivity for $x = 0.75$ and 0.50 , which is by no means reproduced by the extrapolation from the high temperature FPMD results. Fujimura et al. [29] assumed that the deflection point corresponds to the order/disorder transition temperature of Li ions on octahedral sites within the LISICON structure. The transition temperature, T_c , was then estimated by a systematic set of DFT calculations and cluster expansion analyses. The estimated T_c were 380, 750 and 1150 K for $x = 0.75, 0.50$ and 0.25 of $\text{Li}_{2+2x}\text{Zn}_{1-x}\text{GeO}_4$, respectively. The tendency to increase the deflection point with decrease of x was reproduced by the estimated T_c .

Although FPMD results can be well extrapolated to lower temperatures at above T_c , prediction of the conductivity at below T_c is difficult. In order to estimate the diffusivity near the room temperature, which is typically below T_c , one needs to use additional predictors on top of FPMD diffusivity and T_c . In order to select the predictors and examine the prediction error, the raw experimental data-points shown in Fig. 9.6a were used for machine learning. Experimental diffusivities, $D(T)$, at temperature T were “learned” to make prediction models of the diffusivity at a given temperature, T_0 . Fujimura et al. [29] used T , D_{1600} (D at 1600 K), T_c and crystalline volume of disordered structure, V_{dis} , as predictors for the ionic conductivity at $T_0 = 373$ K, σ_{373} using the support-vector regression (SVR) method with a Gaussian kernel [31].

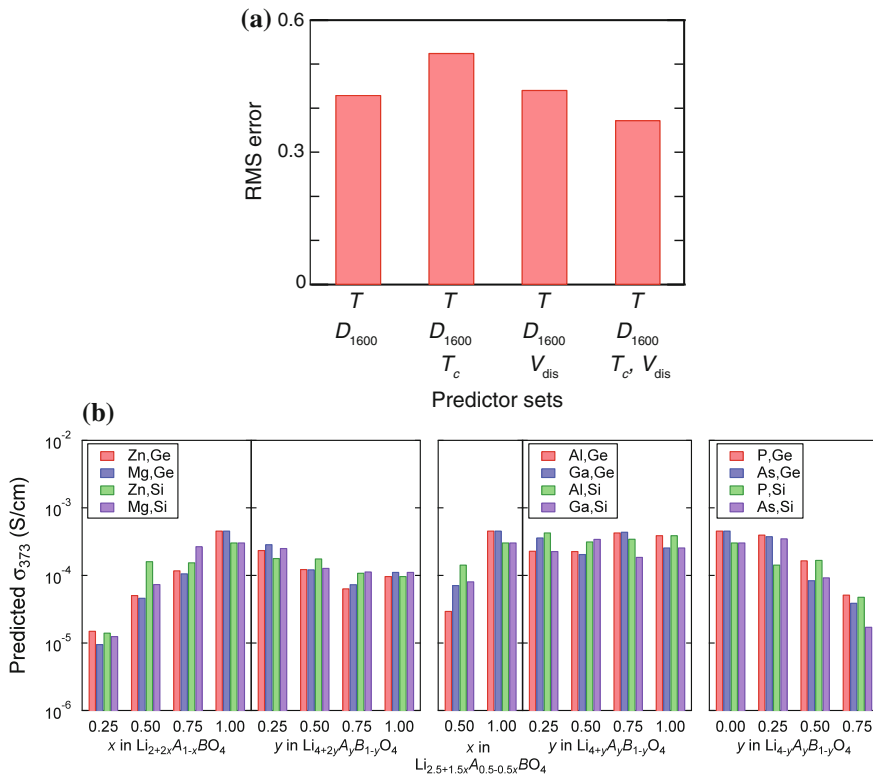


Fig. 9.7 **a** Variation of RMS error for the Li-ion conductivities at 373 K, σ_{373} , with four different sets of predictors. **b** Predicted σ_{373} for 72 compositions of LISICON compounds with the model using all of four predictors [29]

The variance of the Gaussian kernel, the regularization constant and forms of independent variables were optimized by minimizing the prediction error estimated by the bootstrapping method [32]. V_{dis} was calculated by averaging volumes calculated for a few structures with randomly selected Li ion-arrangements on octahedral sites by the DFT method.

Prediction errors of the SVR prediction models for σ_{373} with four different sets of predictors and are compared in Fig. 9.7a. The error increased when only T_c was added on T and D_{1600} . This sounds odd from the physical mechanism viewpoint. However, this can be understood by looking at the experimental data shown in Fig. 9.6a which does not always exhibit two stages separated by the deflection point of the conductivity. The error was lower when V_{dis} was included in the predictor set. The lowest error was obtained when all of T , D_{1600} , T_c and V_{dis} were used. Figure 9.7b shows the predicted σ_{373} for 72 compositions. Even though the theoretical datasets do not contain information about the activation energies explicitly, systems with high D_{1600} and low T_c tend to have high σ_{373} as expected. The conductivity of compounds with

low Zn content such as $\text{Li}_{2+2x}\text{Zn}_{1-x}\text{GeO}_4$ ($x = 0.75$) with high D_{1600} and low T_c were greater than those with high Zn content, such as $\text{Li}_{2+2x}\text{Zn}_{1-x}\text{GeO}_4$ ($x = 0.25$), and high T_c . This result explained the trend observed by experimentalists, namely that the original LISICON composition $\text{Li}_{3.5}\text{Zn}_{0.25}\text{GeO}_4$ has one of the highest Li-ion conductivity. In this study, Li_4GeO_4 was predicted to have the highest σ_{373} of all 72 compounds. However, it has not yet been synthesized because it generally crystallizes to a different crystal structure.

9.5 Combination of DFT Data and Machine Learning III—Thermoelectric Materials

Thermoelectric generators are essential for utilizing waste heat. In order to increase the conversion efficiency, the thermoelectric figure of merit should be increased. Since the figure of merit is inversely proportional to thermal conductivity, many efforts have been placed to decrease the thermal conductivity, especially the lattice thermal conductivity (LTC). In order to evaluate LTC with the accuracy comparable to experimental data, a method that is far beyond the ordinary density functional theory (DFT) calculations is required. Since one needs to treat multiple interactions among phonons, or anharmonic lattice dynamics, the computational cost is many orders of magnitudes higher than the ordinary DFT calculations of primitive cells. Such expensive calculations are practically possible only for a small number of simple compounds. HTS of a large DFT database of LTC is not a realistic approach unless the exploration space is narrowly confined. Carrete and coworkers concentrated their efforts to search low LTC materials within half-Heusler compounds [33]. They made HTS of wide variety of half-Heusler compounds by examination of thermodynamical stability via DFT results. Then LTC was estimated either by full first principles calculations or by a machine-learning algorithm for a selected small number of compounds. HTS of low LTC using a quasiharmonic Debye model was also reported [34]. Efficient prediction of LTC through compressive sensing of lattice dynamics was demonstrated as well [35].

Very recently, Togo et al. [36] reported a method to systematically obtain theoretical LTC through first principles anharmonic lattice dynamics calculations. Results were in quantitative agreement to available experimental data. Using these theoretical data, Seko et al. [37] performed the virtual screening of a library containing 54,779 compounds was by Bayesian optimization using kriging method based on the Gaussian process regressions (see Sect. 9.3). First principles anharmonic lattice dynamics calculations were then performed for highly ranked compounds, which actually showed very low LTC. The strategy is in the category given in Fig. 9.1b. This type of method should be useful for searching materials for many different applications in which chemistry of materials need to be optimized.

References

1. A. Jain, S.P. Ong, G. Hautier, W. Chen, W.D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder et al., *APL Mater.* **1**(1), 011002 (2013)
2. S. Curtarolo, W. Setyawan, S. Wang, J. Xue, K. Yang, R.H. Taylor, L.J. Nelson, G.L. Hart, S. Sanvito, M. Buongiorno-Nardelli et al., *Comput. Mater. Sci.* **58**, 227 (2012)
3. J.E. Saal, S. Kirklin, M. Aykol, B. Meredig, C. Wolverton, *JOM* **65**(11), 1501 (2013)
4. G. Ceder, *MRS bull.* **35**(9), 693 (2010)
5. M. Nishijima, T. Ootani, Y. Kamimura, T. Sueki, S. Esaki, S. Murai, K. Fujita, K. Tanaka, K. Ohira, Y. Koyama, et al., *Nat. Commun.* **5** (2014)
6. A.K. Padhi, K. Nanjundaswamy, J. Goodenough, *J. Electrochem. Soc.* **144**(4), 1188 (1997)
7. C. Delmas, M. Maccario, L. Croguennec, F. Le Cras, F. Weill, *Nat. Mater.* **7**(8), 665 (2008)
8. D. Wang, X. Wu, Z. Wang, L. Chen, *J. Power Sources* **140**(1), 125 (2005)
9. S. Curtarolo, G.L. Hart, M.B. Nardelli, N. Mingo, S. Sanvito, O. Levy, *Nat. Mater.* **12**(3), 191 (2013)
10. F.A. Lindemann, *Phys. Z.* **11**, 609 (1910)
11. A. Lawson, *Phil. Mag.* **81**(3), 255 (2001)
12. A.C. Lawson, *Phil. Mag.* **89**(22–24), 1757 (2009)
13. A. Granato, D. Joncich, V. Khonik, *Appl. Phys. Lett.* **97**(17), 171911 (2010)
14. F. Guinea, J.H. Rose, J.R. Smith, J. Ferrante, *Appl. Phys. Lett.* **44**, 53 (1984)
15. J.A. Van Vechten, *Phys. Rev. Lett.* **29**, 769 (1972)
16. J.R. Chelikowsky, K.E. Anderson, *J. Phys. Chem. Solids* **48**, 197 (1987)
17. Y. Saad, D. Gao, T. Ngo, S. Bobbitt, J.R. Chelikowsky, W. Andreoni, *Phys. Rev. B* **85**, 104104 (2012)
18. A. Seko, T. Maekawa, K. Tsuda, I. Tanaka, *Phys. Rev. B* **89**(5), 054303 (2014)
19. W.M. Haynes, *CRC Handbook of Chemistry and Physics*, 92nd edn. (CRC Press, Boca Raton, 2012)
20. W.N. Venables, B.D. Ripley, *Modern Applied Statistics with S*, 4th edn. (Springer, New York, 2002)
21. H. Akaike, in *Second International Symposium on Information Theory*, (Akademina Kiado, 1973), pp. 267–281
22. K. Lejaeghere, J. Jaeken, V. Van Speybroeck, S. Cottenier, *Phys. Rev. B* **89**(1), 014304 (2014)
23. A. Robertson, A. West, A. Ritchie, *Solid State Ionics* **104**(1), 1 (1997)
24. H.P. Hong, *Mater. Res. Bull.* **13**(2), 117 (1978)
25. U. Alpen, M. Bell, W. Michelhaus, K. Cheung, G. Dudley, *Electrochim. Acta* **23**(12), 1395 (1978)
26. D. Mazumdar, D. Bose, M. Mukherjee, *Solid state Ionics* **14**(2), 143 (1984)
27. P. Bruce, A. West, *J. Solid State Chem.* **44**(3), 354 (1982)
28. P. Bruce, I. Abrahams, *J. Solid State Chem.* **95**(1), 74 (1991)
29. K. Fujimura, A. Seko, Y. Koyama, A. Kuwabara, I. Kishida, K. Shitara, C.A.J. Fisher, H. Moriwake, I. Tanaka, *Adv. Energy Mater.* **3**(8), 980 (2013)
30. S. Takai, K. Kurihara, K. Yoneda, S. Fujine, Y. Kawabata, T. Esaka, *Solid State Ionics* **171**(1), 107 (2004)
31. C.C. Chang, C.J. Lin, A.C.M. Trans, *Intell. Syst. Tech.* **2**(3), 27 (2011)
32. B. Efron, R.J. Tibshirani, *An Introduction to the Bootstrap* (CRC press, Boca Raton, 1994)
33. J. Carrete, W. Li, N. Mingo, S. Wang, S. Curtarolo, *Phys. Rev. X* **4**(1), 011019 (2014)
34. C. Toher, J.J. Plata, O. Levy, M. de Jong, M. Asta, M.B. Nardelli, S. Curtarolo, *Phys. Rev. B* **90**(17), 174107 (2014)
35. F. Zhou, W. Nielson, Y. Xia, V. Ozoliņš, *Phys. Rev. Lett.* **113**(18), 185501 (2014)
36. A. Togo, L. Chaput, I. Tanaka, *Phys. Rev. B* **91**(9), 094306 (2015)
37. A. Seko, A. Togo, H. Hayashi, K. Tsuda, L. Chaput, I. Tanaka, arXiv preprint [arXiv:1506.06439](https://arxiv.org/abs/1506.06439) (2015)

Chapter 10

Materials Informatics Using *Ab initio* Data: Application to MAX Phases

Wai-Yim Ching

Abstract We use a database constructed for a very unique class of laminated inter-metallic compounds, the MAX ($M_{n+1}AX_n$) phase, to show how materials informatics can be used to predict the existence of new, hitherto unexplored phases. The focus of this Chapter is the correlation between seemingly disconnected descriptors and the importance of high quality, computationally derived data. An extension of this approach to other specific materials systems is discussed.

10.1 Introduction

In recent years, information gathering, analysis, and interpretation has emerged as an interdisciplinary research skill involving computer science, information science, and other various domains of science such as physics, chemistry, biology, medicine, materials engineering, design technology, education and social science, etc. [1]. In particular, materials informatics has developed into a flourishing field of study [2]. It aims to find more efficient ways of solving scientific problems related to all kinds on materials using large databases. This started at the initiation of the materials genome project at the federal level and it follows the same approach that the Human Genome Project in the biomedical community did decades ago, which resulted in the now mature discipline of bioinformatics. Creative software, genetic algorithms, and visualization tools have been developed to do statistical analysis of data and to explore the data via data mining aided by powerful high performance computers [3]. There are many examples of highly successful applications for identifying and understanding the structure-properties correlations and to formulate design rules for better materials for specific applications. The information obtained from high

W.-Y. Ching (✉)

Curators Professor of Physics, 250C Robert H. Flarsheim Hall,
5100 Rockhill Road, Kansas City, MO 64110-2499, USA
e-mail: ChingW@umkc.edu
URL: <http://cas.umkc.edu/physics/ching/index.htm>

© Springer International Publishing Switzerland 2016
T. Lookman et al. (eds.), *Information Science for Materials
Discovery and Design*, Springer Series in Materials Science 225,
DOI 10.1007/978-3-319-23871-5_10

throughput materials informatics greatly reduces the time that it takes to go from frontier research to real applications.

There are many different ways of collecting large data and of building powerful databases for applications. Traditionally, the data for materials properties are collected from experimentally measured values published in open literature such as: crystal structures, density, heat of formations, melting temperature, electric conductivity, thermal conductivity, refractive index, bulk modulus, hardness, phase diagrams, and much more. These data cover all kinds of material systems regardless of the source or the reliability of the data. Such database are usually not vetted and they are of varying quality. However, the argument is that in a statistical sense, any invalid data that appears does so as noise and will not make much of a difference, as long as the database is large enough and the method for analysis is carefully designed. This *modus operandi* is more common in the biomedical arena when dealing with experimental or clinical trials with large data collected over a long period of time while looking for small effects [4]. Contrary to some approaches which aim to reduce or avoid accurate atomistic simulations by instead relying purely on statistical predictions, is another approach that is based on the design of a specific data base with high accuracy using computational genomics. This difference simply reflects the emphasis on a different spectrum of the field of materials informatics with different strategies for different systems although both are data driven.

More recently, large amounts of data may be obtained through calculations using different computational methods and packages based on different theories. The trend is usually to cover a focused groups of materials that are categorized either in their structure, composition, functionality, or some specific materials property. Examples for such recent endeavors include: piezoelectric perovskites [5], battery materials comprising oxides, phosphates, borates, silicates sulphates, etc. [3, 6], Half-Heusler semiconductors with low thermoconductivity [7], binary compounds [8], polymer physics materials genome [9], and isotope substitution on phonons in graphene [10] just to name a few. It is also possible to combine the measured data and the calculated data into a bigger database.

In this Chapter, we present a specific case to illustrate the application of materials informatics using a large database of a unique class of materials, the MAX ($M_{n+1}AX_n$) phase [11]. Our approach is to select a specific material system with well-defined structures and compositions for a focused study and then apply state-of-the-art computational tools to systematically generate a large amount of data on their physical properties, and the analysis of correlations amongst them. We then use this database to test the efficacy of exiting data mining and machine learning algorithms. Simultaneously, this enables us to predict the existence of new MAX phases that have not yet been synthesized or studied in the laboratory but which may have outstanding properties. The identification of outliers that clearly do not follow general trends helps to obtain deeper insights and reveal the fundamental reasons behind such deviations. The predictive capability of the data mining is substantially controlled by the quality of the assigned descriptors. At the same time, use of theoretical-based descriptors that demand a large computational time will be impractical. Thus, delegating such types of descriptors to a combinations of less

time-demanding descriptors will be the goal. This approach is certainly different from other approaches which depend on collecting data from various sources but it puts the data under better control with increased reliability in the interpretations.

Another important issue which is less frequently discussed in materials informatics is the way that the data are presented. Many believe that materials informatics relies on massive data collections and their statistical analysis. Everything is numerical and machine-based. On the other hand, we find that creative and insightful graphical representations of the data can allow one to grasp some of the most important points without laborious analysis. This will be amply demonstrated for the materials presented in this Chapter.

The MAX phase is used as an example to illustrate various aspects of materials properties and correlations between different descriptors. We have identified one such descriptor in particular, the total bond order density (TBOD) that plays a dominating role. We articulate on some other materials systems for which the application of the same approach and the use of TBOD can be very fruitful.


10.2 MAX Phases: A Unique Class of Material

MAX phases or $M_{n+1}AX_n$ are transition metal ternary compounds with layered structures where “M” is an early transition metal, “A” is a metalloid element, “X” is either carbon or nitrogen, and n is the layer index. MAX compounds have attracted a great deal of attention in recent decades due to many of their fascinating properties and wide range of potential applications. Up to now, only about 70 of these phases are confirmed or have been synthesized [12]. The majority of these confirmed phases are 211 carbides with $n = 1$ or 2 and with $M = \text{Ti}$ and Zr , and $A = \text{Al}$ and Ga . It has also been demonstrated that the formation of composite phases and solid solutions in MAX phases between different “M” elements, “A” elements and C and N are possible. Such possibilities have greatly extended the range of compositions beyond the ternary phases.

The MAX phases are layered hexagonal crystals (space group: $P6_3/mmc$ No. 194). Figure 10.1 displays the crystal structure of MAX for $n = 1, 2, 3, 4$, which are usually referred to as the 211, 312, 413, and 514 phases. An important feature is that in MAX compounds, layer “A” remains constant whereas layers of “M” and “X” increase with n . The “X” layers are always in between the “M” layers or blocks of MX layers connected by single “A” layer which can significantly affect the properties of a MAX phase. The physical properties of MAX phases vary over a wide range depending on “M”, “A”, “X” and n . MAX phases with $n \geq 5$ are known to exist but are very rare. Most of the existing experimental work on the MAX phases has been on the 211 and 312 carbides.

MAX phases behave both like ceramics and metals with some very desirable properties such as machinability, thermal shock resistance, damage tolerance, fatigue, creep, and oxidation resistance and elastic stiffness. They are also good thermal and electrical conductors [12]. More recently, MAX phases have been considered for

Fig. 10.1 Sketch of crystal structures of four MAX phases M_2AX , M_3AX_2 , M_4AX_3 , M_5AX_4 (i.e. with $n = 1, 2, 3, 4$)



high-temperature structural applications. Other applications include porous exhaust gas filters for automobiles, heat exchangers, heating elements, wear and corrosion protective surface coatings, electrodes, resistors, capacitors, rotating electrical contacts, nuclear applications, as bio-compatible materials, cutting tools, nozzles, tools for die pressing, impact-resistant materials, projectile proof armor, and much more. Some of these applications already have products on the market. The physical properties of MAX phases have been investigated by many groups both experimentally and computationally (see extensive references in Aryal et al. [11]). We have been focusing mostly on the mechanical properties and electronic structure of MAX phases. The elastic coefficients and mechanical parameters such as the bulk modulus (K), shear modulus (G), Young's modulus (E), and Poisson's ratio (ν) are derived from the calculated elastic coefficients under the VRH polycrystalline approximation and were obtained. The G/K ratio, also known as the Pugh moduli ratio is a good indicator of the ductility or brittleness of the alloy based on an analysis in pure metals but it has also been quite effective when applied to metallic alloys [13]. The other physical properties investigated are the optical conductivities in 20 MAX phases [14] and the core-level excitations in some of the compounds [15]. More recently, we also estimated the high temperature lattice thermal conductivities of MAX phases (see Sect. 10.4).

The electronic structure and bonding is the basic information needed to understand the properties in any materials. It has been well studied for MAX phases using density functional theory-based methods by many groups over the last 15 years. Most of the discussion tends to be on the band structure and the density of states (DOS) and partial density of states (PDOS). In MAX phases, interatomic bonding

are fairly complicated involving metallic, partly covalent and partly ionic types of bonding which may extend beyond nearest neighbors. The structural complexity and variations in chemical species make characterization of interatomic bonding in MAX phases particularly challenging. We advocate the use of total bond order (TBO), total bond order density (TBOD) and their partial components (PBOD) as useful metrics to delineate the observed physical properties. TBO is the sum of all bond order pairs in the crystal and when normalized by its volume, we get the TBOD. This will be illustrated more in the following sections.

It is worth mentioning that in addition to the canonical MAX phases, there are related materials derived from the MAX phases such as the solid solutions with different “M” or “A” elements and with mixtures of C and N. The MAX solid solutions can expand the list of such compounds enormously and some of them may have optimized compositions that enhance their properties. This provided a great opportunity to apply the techniques of materials informatics for facilitating the processing of large amounts of data. Other related systems include the 2-dimensional $M_{n+1}X_n$ system called Maxenes by extracting the “A” layer from MAX by exfoliating in solution which offers a variety of new applications. Last but not least, there are quite a few layered intermetallic compounds with different types of stacking layers but involving more or less similar chemical species that have not been fully exploited.

10.3 Applications of Materials Informatics to MAX Phases

10.3.1 Initial Screening and Construction of the MAX Database

We first construct a database consisting of as many MAX phases in accordance with the general guideline suggested by Barsoum (Barsoum’s Book, page 2, Fig. 1.2 [12]). We chose 9 “M” elements (Sc, Ti, Zr, H, V, Nb, Ta, Cr, Mo), 11 “A” elements (Al, Ga, In, Tl, Si, Ge, Sn, Pb, P, As, S), $X = C$ and N and with the layer index $n = 1, 2, 3, 4$. This gives us a total of 792 possible MAX ($M_{n+1}AX_n$) phases. We used the Vienna *Ab initio* Simulation Package (VASP) [16] to optimize the structure and obtain the elastic constants for each crystal. However, not all of these phases will be stable. We therefore screen these by using two stability criteria: the Cauchy-Born elastic stability criteria for hexagonal crystals [17] which eliminated 71 crystals. Next, we calculated the heat of formation (HoF) for the same 792 crystals based on the relative stability of each MAX phase to the formation energy of its elements in their most stable ground state structure. As a result, 45 additional phases with positive HoF were eliminated, resulting in 665 viable MAX phases for a more focused study. The use of these two criteria instead of a more rigorous but far more time consuming one based on thermodynamic assessment on all potential competing phases in the M-A-X ternary phase diagrams is a reasonable compromise. In principle, we can consider these two sets of criteria employed as two descriptors in the data mining

approach. This represents a substantial savings in computational time. The calculated elastic and mechanical properties of the 665 MAX phases are tabulated as illustrated in Table 10.1 for 20 such phases. The electronic structure and bonding of the MAX phases are calculated using the orthogonalized linear combination of atomic orbital method (OLCAO) [18]. This is an extremely efficient and well-tested method using atomic orbitals in the basis expansion. The main descriptors for electronic properties are summarized in tabular form as illustrated in Table 10.2 for 14 such crystals. Both sets of data for the 665 MAX phases are publically available [11].

10.3.2 Representative Results on Mechanical Properties and Electronic Structure of MAX

We selectively present some of the calculated results from the database for the 665 MAX phases. Figure 10.2 shows a scattered plot of shear modulus G versus bulk modulus K for all screened 665 MAX phases. To have a broader perspective, we used different colors for index n , and open or closed symbols for carbides and nitrides respectively. We also include similar data for some metallic compounds and selected binary MX compounds [19]. We note that the MAX phases cover a wide region of bulk and shear moduli, overlapping with those of the common metals and alloys. The dashed lines show the G/K ratios for these data which range from a minimum of 0.12 to a maximum of 0.8. The maximum G/K ratio is close to those of MN binary compounds and the low G/K values are mostly from MAX nitrides. Figure 10.2 illustrates a conventional graphical presentation in materials informatics to provide an overview of the data from a large database.

The data for G/K values for all MAX phases shown in Fig. 10.2 as scattered plot data are presented in Fig. 10.3 in a different way in the form of an innovative map resembling the Periodic Table. For this plot we used the original 792 hypothetical MAX phase data. This enables us to clearly see the locations of those phases that have been screened out relative to those that have not. Here the “M” elements are plotted on the Y-axes and the “A” elements are along the X-axes. The color of each square cell represents the G/K value of that particular MAX phase along with other information such as whether the phase has been synthesized or not. The phases that have been eliminated by the Cauchy-Born criterion or the HoF criteria are marked with a + or a × respectively. The experimentally confirmed phases are marked with a white star. As can be seen, none of the experimentally confirmed phases are among the ones judged to be unstable and screened out. There are many boxes of different colors without the white star, suggesting the existence of a myriad of possible MAX phases not yet explored. While the G/K ratio of MAX phases can vary over a wide range as indicated by the variations in color for the different squares in Fig. 10.3, we can delineate the boundaries of MAX’s materials properties within which optimized functionalities can be further explored. Similar maps for other mechanical properties for the MAX phases can be found in [11].

Table 10.1 Samples of descriptors for mechanical properties in the database (Units in GPA)

Crystals	C ₁₁	C ₁₂	C ₁₃	C ₃₃	C ₄₄	C ₆₆	K	G	E	η	G/K
Ti ₃ AlC ₂	355.8	81.4	75.3	293.4	120.3	137.2	162.5	126.7	301.7	0.191	0.78
Ti ₃ SiC ₂	369.6	96.2	107.6	358.3	155.0	136.7	191.1	141.3	340.0	0.204	0.74
Ti ₃ GeC ₂	362.0	97.2	97.7	332.0	137.3	132.4	182.2	132.2	319.3	0.208	0.73
Ti ₂ AlC	301.9	68.0	63.0	267.9	105.1	117.0	139.7	110.5	262.3	0.187	0.79
Ti ₂ GaC	300.8	79.2	63.8	246.5	92.4	110.8	139.3	101.4	244.9	0.207	0.73
Ti ₂ InC	284.4	69.3	55.2	235.5	83.9	107.5	128.6	96.0	230.5	0.201	0.75
Ti ₂ SiC	312.9	82.1	110.4	329.2	149.6	115.4	173.0	124.9	302.0	0.209	0.72
Ti ₂ GeC	296.6	85.7	96.8	297.1	121.5	105.5	161.0	110.0	268.8	0.222	0.68
Ti ₂ SnC	262.6	88.6	73.1	255.2	96.8	87.0	138.8	92.4	226.8	0.228	0.67
Ti ₂ PC	256.8	144.8	155.0	339.5	166.3	56.0	191.8	93.1	240.4	0.291	0.49
Ti ₂ AsC	212.9	180.4	123.7	289.5	146.3	16.2	150.7	57.2	152.3	0.332	0.38
Ti ₂ SC	339.8	101.4	109.7	361.9	159.5	119.2	186.8	134.4	325.2	0.210	0.72
Ti ₂ AlN	312.9	73.0	95.5	290.7	126.1	120.0	160.5	117.4	283.1	0.206	0.73
V ₂ AlC	334.4	71.5	106.0	320.8	149.8	131.5	172.9	132.1	315.9	0.196	0.76
Nb ₂ AlC	316.6	86.3	117.0	288.6	137.6	115.2	173.6	116.4	285.5	0.226	0.67
Cr ₂ AlC	366.3	85.8	111.3	356.9	142.9	140.2	189.6	137.0	331.2	0.209	0.72
Ta ₂ AlC	344.5	112.2	137.1	327.9	152.3	116.1	198.8	124.1	308.1	0.242	0.62
α -Ta ₃ AlC ₂	453.6	130.5	135.6	388.4	175.0	161.5	232.8	161.1	392.8	0.219	0.69
α -Ta ₄ AlC ₃	459.2	149.1	148.7	383.1	170.5	155.0	243.0	155.3	384.1	0.237	0.64
Ta ₅ AlC ₄	481.5	149.6	158.1	423.6	188.8	165.9	257.2	169.1	416.0	0.231	0.66

Table 10.2 Samples of descriptors from electronic structure in the database

MAX	$\Delta Q^*(M)$	$\Delta Q^*(X)$	$\Delta Q^*(A)$	TBO	BO(M-X)	BO(M-M)	BO(M-A)	BO(A-A)	$N(E_F)$
Ti₂AlC	-0.330	-0.043	0.703	23.510	10.258	4.512	7.231	1.508	11.052
Ti₂GaC	-0.485	0.269	0.701	22.680	10.289	4.060	6.986	1.340	10.572
Ti₂InC	-0.424	0.148	0.700	22.750	10.238	4.396	6.482	1.636	9.260
Ti₂SiC	-0.393	0.097	0.688	22.820	10.344	3.583	8.153	0.742	12.921
Ti₂GeC	-0.509	0.324	0.694	21.750	10.337	3.541	7.111	0.758	14.720
Ti₂SnC	-0.381	0.069	0.693	22.320	10.294	3.926	7.110	0.993	15.084
Ti₂PC	-0.454	0.210	0.699	22.740	10.366	2.802	9.571	0.000	21.762
Ti₂AsC	-0.505	0.316	0.695	21.360	10.382	2.893	8.086	0.000	19.697
Ti₂SC	-0.447	0.189	0.705	21.340	10.380	2.944	8.018	0.000	7.301
Ti₂AlN	-0.295	-0.087	0.679	22.150	8.702	4.646	7.217	1.585	15.502
V₂AlC	-0.277	-0.101	0.655	22.820	10.017	4.192	6.905	1.704	21.663
Nb₂AlC	-0.493	0.245	0.741	15.410	7.319	1.253	5.354	1.399	13.338
Cr₂AlC	-0.098	-0.324	0.521	21.250	9.559	2.837	7.080	1.769	24.384
Ta₂AlC	-0.324	-0.044	0.692	24.810	10.130	5.724	7.561	1.397	11.126

 ΔQ^* and bond orders (electrons) and $N(E_F)$ (states/eV-cell)




Fig. 10.2 Shear modulus versus bulk modulus for 665 screened MAX phases in the database. *Solid circles and open circles* are for carbides and nitrides respectively. Different color is used for different n in $M_{n+1}AX_n$. Also shown are the locations of other metals and binary MC and MN compounds

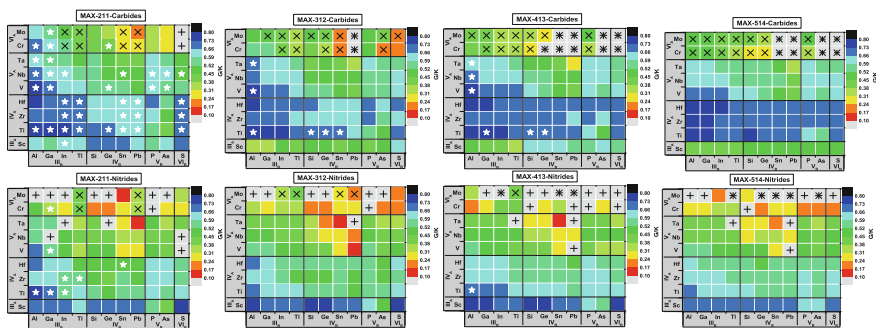


Fig. 10.3 G/K ratio map for 792 MAX phases according to “M” (Y-axis) and “A” (X-axis) elements. *Top panel* for carbides and *lower panel* for nitrides. Color in each cell represents the calculated G/K value as indicated in the *color bar*. A *star* in the box indicates that this phase has been confirmed. “+” means the phases is eliminated for elastic instability and “x” means the phase is screened out for thermodynamic instability or positive HoF

We now present some of the results related to electronic structure. The density of states (DOS) at the Fermi level (E_F) or $N(E_F)$ is one of the important electronic parameter for metallic systems. In MAX phases $N(E_F)$ is a strong function of composition. Some values are close to zero, whereas others are quite large, depending on whether E_F is located in the vicinity of the 3d or 4d orbitals of “M”. The calculated

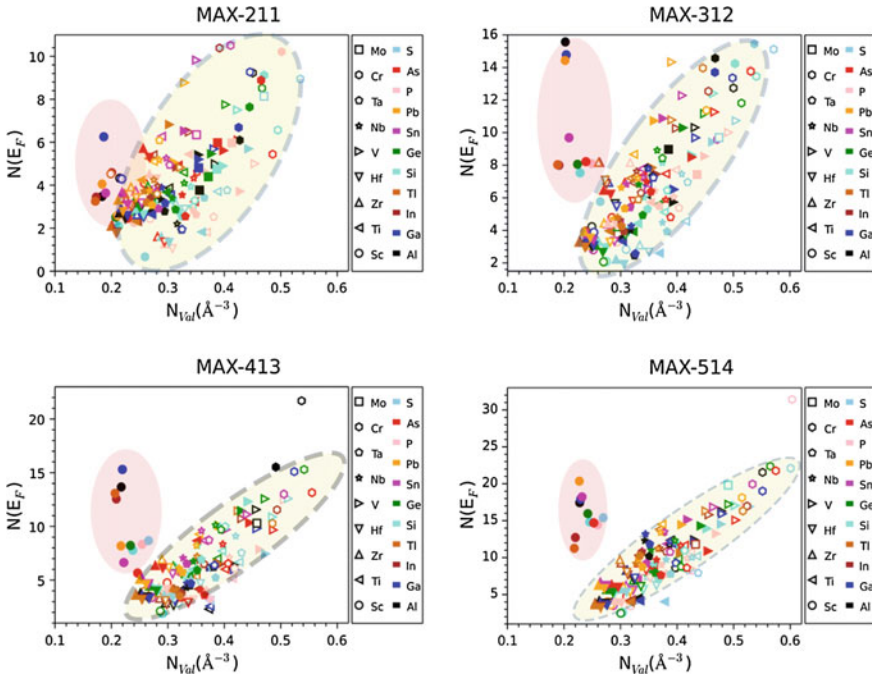


Fig. 10.4 Plot of DOS at Fermi level $N(E_F)$ against total number of valence electrons per unit volume for the 665 MAX phases in the database. *Solid symbols* for carbides and *open symbols* for nitrides. Note the outlying nature of the data for the $M = \text{Sc}$ with $X = \text{C}$ MAX phases

$N(E_F)$ per unit cell is found to be reasonably correlated with the total valence electron number per volume ($N_{\text{val}}(\text{\AA}^{-3})$) as shown in Fig. 10.4. The total valence electron number is the sum of the formal valence electrons of individual atoms in the crystal. In general, larger $N_{\text{val}}(\text{\AA}^{-3})$ corresponds to larger $N(E_F)$ as expected. Also, as n increases, ($N_{\text{val}}(\text{\AA}^{-3})$) increases, the slope of the data distribution decreases. Traditionally, it has been speculated but not rigorously proved that the existence of a local minimum (or pseudo gap) at the Fermi level in a metal or alloy signifies its structural stability [20]. While all the DOS for the MAX phases are available, it is not practical to present the DOS figures for all the phases. However, the relative magnitudes of $N(E_F)$ and its decomposition into different atomic components for each phase is a valid descriptor for the electronic structure.

Figure 10.4 shows that nitrides have larger $N(E_F)$ values than the carbides. The Sc-based carbides, however, are a notable exception. They have significantly higher $N(E_F)$ than their nitride counterparts. The Sc-based carbides (but not the nitrides) also show a marked deviation from the general trend of $N(E_F)$ versus ($N_{\text{val}}(\text{\AA}^{-3})$) with increasing n . The approximate positive linear correlation between the two properties becomes more pronounced with increasing n which can be attributed to the increased M atoms as n increases. One can relate this linear trend to a similar behavior observed

in the binary mono-carbides or nitrides [21]. However, a real distinction here with respect to the MAX phases is the profound role of “A” which is not present in the binary mono-carbides/nitrides. In general, the presence of the “A” element appears to significantly lessen the degree of the linear correlation or increase in the scattering of the data (see Fig. 10.4 for MAX 211). It is only at higher n values, where the “A” content is reduced and consequently the bonding characteristics are less influenced by “A”, that a stronger correlation between the two properties emerges, mimicking those of binary mono-carbides/nitrides.

10.3.3 Classification of Descriptors from the Database and Correlation Among Them

The MAX database consists of the calculated quantities of all MAX phases as illustrated in Tables 10.1 and 10.2 shown earlier. These numerical quantities can be classified into descriptors or controlling factors to be used in data mining algorithms for materials informatics and to explore their correlations. A simple flow chart is shown in Fig. 10.5. For the MAX phases, we classify the descriptors into three categories based on their level complexity and/or computational time required to obtain the data:

- (1) Basic chemistry descriptors: They are (i) number of valence electrons, (ii) atomic number Z of the elements, and (iii) volume of the unit cell.
- (2) Descriptors from the electronic structure and bonding: They are: (i) Total bond order density (TBOD) (normalized to crystal volume), (ii) total bond order from

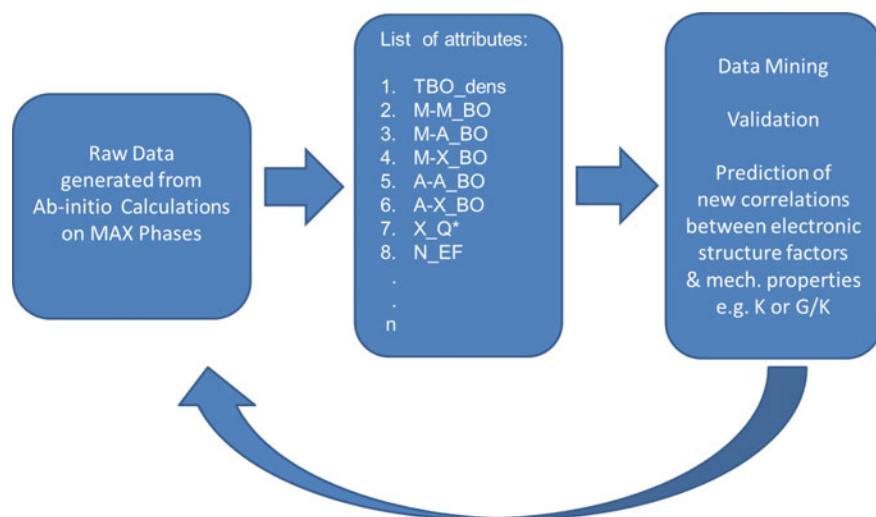


Fig. 10.5 Flow chart of the approach used for data mining in materials informatics for MAX phases

different atomic pairs, or the M–M, M–A, M–X, A–X and X–X pairs, (iii) density of states at the Fermi level $N(E_F)$.

(3) Descriptors for the elastic constants C_{ij} and the bulk mechanical parameters, K , G , E , η and G/K ratio.

We then seek to establish the correlations between these descriptors which are interrelated. More specifically, correlations between elastic and mechanical descriptors, correlations between electronic descriptors and correlation between mechanical descriptors and electronic descriptors [11]. We have been able to demonstrate that over 90% correlation can be achieved using a simple linear regression method implying that the mechanical properties descriptors can be adequately represented by the other two types of descriptors. This is illustrated in the following section.

10.3.4 Verification of the Efficacy of the Materials Informatics Tools

The success in linking the electronic structure factors to complex bulk elastic properties has enabled us to advance the utility of the data mining approach for expanding the materials database for MAX phases. We have also extended our analysis into the components of the second order elastic constants of the MAX phases. Figure 10.6 shows an example of such an analysis as applied to the 2-1-1 MAX carbides, a large subset of our database. The three main elastic constants in the database, C_{11} , C_{33} , C_{13} , which were calculated using an *ab initio* method are compared to those same values as predicted by a combination of electronic structure factors and valence electron information with reasonably high correlation coefficients of 0.83, 0.93, 0.95 respectively. This suggests that such a method is robust enough to probe orientation dependent second-order elastic constants with a high accuracy.




Fig. 10.6 Comparison of **a** C_{11} , **b** C_{33} , and **c** C_{13} of 211 MAX carbides obtained from *ab initio* calculations (x-axis) and those from the data mining prediction (y-axis)

Result of linear regression of C_{11} , C_{33} , and C_{13} with chemical and electronic descriptors shown in Fig. 10.6 are as follows:

$$C_{11} = 0.6235 \times Z_M + 8.1344 \times (\text{GN})_M - 0.8737 \times Z_X + 32.0051 \times Q_A - 144.2461 \times Q_X + 10.9223 \times (\text{BO})_{M-X} + 9.2461 \times (\text{BO})_{M-A} - 7.8791 \times (\text{BO})_{M-M} + 11.8688 \times (\text{BO})_{A-A} + 470.6772 \times (\text{BO})_{A-X} - 2.7405 \times N(E_F) + 243.5997.$$

Correlation coefficient = 0.83

$$C_{33} = 0.7155 \times Z_M + 19.8291 \times (\text{GN})_M - 1.085 \times Z_X + 18.9992 \times (\text{GN})_X + 18.4127 \times (\text{BO})_{M-A} - 8.9407 \times (\text{BO})_{M-M} + 16.2802 \times (\text{BO})_{A-A} - 1.0634 \times N(E_F) + 38.8039.$$

Correlation coefficient = 0.9264.

$$C_{13} = 36.994 \times (\text{GN})_M - 0.402 \times Z_X - 7.3952 \times (\text{GN})_X + 67.8729 \times Q_A - 12.8243 \times (\text{BO})_{M-X} + 15.916 \times (\text{BO})_{M-A} + 7.6037 \times (\text{BO})_{M-M} - 33.8152 \times (\text{BO})_{A-A} - 11.1306.$$

Correlation coefficient = 0.9541.

where Z_M = atomic number of M, $(\text{GN})_M$ = group number of M from Periodic Table, Z_X = atomic number of X, $(\text{GN})_X$ = group number of X, $(\text{BO})_{M-A}$ = total bond order of M-A pairs, $(\text{BO})_{M-X}$ = total bond order of M-X pairs, $(\text{BO})_{M-M}$ = total bond order of M-M pairs, $(\text{BO})_{A-X}$ = total bond order of A-X pairs, $(\text{BO})_{A-A}$ = total bond order of A-A pairs, $N(E_F)$ = DOS at the Fermi level. Q_A and Q_X are the effective charges on A and X respectively (see [11] for more details).

Figures 10.7 and 10.8 show a different way to test the efficacy of the data mining algorithm. Here, we use 50% of randomly chosen data from the database as a training set and use the existing algorithm WEKA [22] to predict the properties of the other 50% of the MAX phases by comparing predicted values with the *ab initio* data in the database. Figure 10.7 (top panel) shows the comparison between K obtained from *ab initio* calculations versus those obtained from the formulas derived from the data mining algorithm for the other 50% for the 211, 312, 413 and 514 MAX phases. An excellent correlation with over 90% correlation factor for each type of the MAX phase is obtained. The lower panel of Fig. 10.7 also shows pie charts of the relative

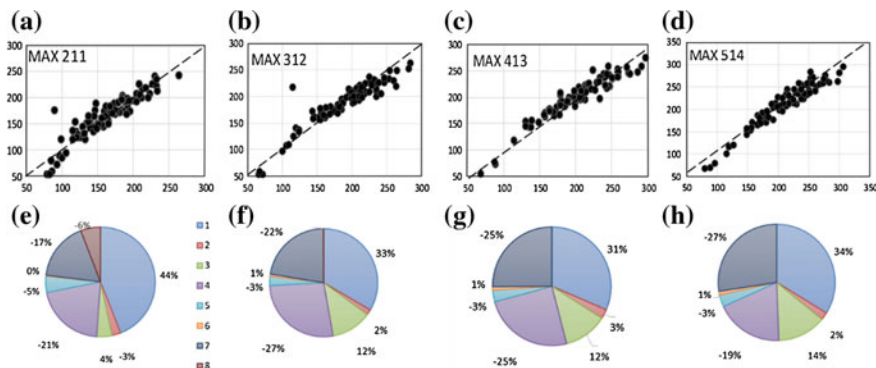


Fig. 10.7 Top panel Use of 50% of MAX data for bulk modulus K as training set to predict the other 50% by comparing with *ab initio* data for 665 MAX phases. Lower panel relative contribution from different electronic structure descriptors




Fig. 10.8 Same as Fig. 10.7 but for the shear modulus

contribution from each type of the electronic structure descriptors used to predict K . The four most important factors are the total bond order density (TBOD), BOD of the M–A pairs (M–A BOD), BOD of the M–X pairs (M–X BOD) and charge transfer for the X elements. TBOD clearly stands out as the most important factor in determining K for all MAX phases. Figure 10.8 shows a similar prediction for G/K ratio using the same procedure. Although the correlation factor is less impressive than for K , the prediction from data mining is still reasonably good, with correlation coefficients around 80 % or higher. The prediction for Poisson’s ratio is at the same level as for the G/K ratio. In both cases, they are strongly affected by the TBOD although it is a negative correlation instead of the positive correlation exhibited by K . The linear correlation is less definitive in this case, which is probably due to the fact that the G/K ratio or Poisson’s ratio are more influenced by the nature of the “A” element. This is an effect that apparently is not fully represented by the BO parameters. Nevertheless, a reasonably good estimate for these two properties can be established solely from a linear combination of electronic structure factors. Furthermore, the TBDO emerges as a significant descriptor that controls the mechanical properties. This data mining approach also demonstrates that a simple correlation can be used to link elastic

parameters such as Poisson's ratio or the Pugh ratio to a series of electronic structure indicators. The use of only 50% of the data as a training set gives credence to the particular machine learning software and the philosophy lying behind it.

10.4 Further Applications of MAX Data

Since the generation of the MAX database less than a year ago, additional results that use this data to estimate the lattice thermal conductivity of MAX phases at high temperature and the calculation of universal elastic anisotropy based on a recently developed theory have been obtained. These are prime examples of the utility of large database to easily get new information without lengthy calculations consistent with the spirit of materials informatics. They are briefly described below.

10.4.1 Lattice Thermal Conductivity at High Temperature

A systematic calculation of the lattice thermal conductivity κ_{ph} and minimum thermal conductivity κ_{min} for the 211, 312, and 413 MAX phases using Slack's equation and the Clarke formula respectively has been carried out [23]. The parameters used in these simplified calculations are extracted from the elastic coefficients C_{ij} , bulk mechanical properties, and equilibrium volume of all stable MAX phase compounds in the database. Essentially, the calculation of κ_{ph} , follows the equation derived by Slack [24],

$$\kappa_{ph} = A \frac{\bar{M}\theta_D^3\delta}{\gamma^2 n^{2/3} T} \quad (10.1)$$

where \bar{M} is the average atomic weight (in units of kg/mol), δ is the average volume of the mass equivalent to one atom in the primitive cell (in units of m^3), T is the absolute temperature, n is the number of atoms per unit cell, γ the Grüneisen constant derived from Poisson's ratio (ν) and A is a coefficient (in units of $W \text{ mol/kg/m}^2/\text{K}^3$) that depends on γ as determined by Julian [25]. These parameters can be obtained from the database for the MAX phases.

Figure 10.9 shows the calculated κ_{ph} for the 211, 312, and 413 phases of MAX carbides and nitrides presented in two separate ways in order to trace the trends associated with variations in the atomic numbers of the "M" and "A" elements. The top panel in Fig. 10.9 is for "M"-based plots and the lower panel is for "A"-based plots. The x-axis lists the 9 "M" elements and 11 "A" elements for the top and bottom panels respectively. To grasp the variations and overall trends in κ_{ph} more easily, we employ the following strategy: (1) the data for the carbides (solid circles) and nitrides (open circles) are plotted on the same figure. (2) The horizontal x-axis is arranged in the order of increasing atomic number Z . They are (Sc, Ti, V, Cr, Zr, Nb, Mo,

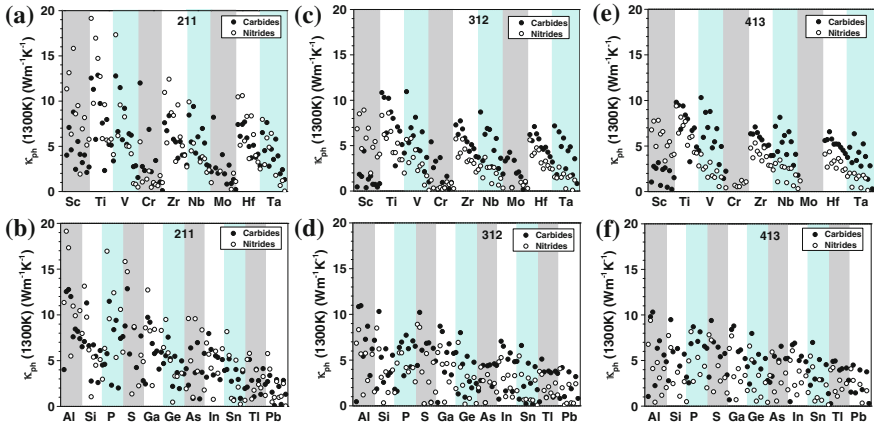


Fig. 10.9 Scatter plots of calculated phonon thermal conductivity (κ_{ph}) at 1300 K of MAX phases: **a** 211 in “M” trend; **b** 211 in “A” trend; **c** 312 in “M” trend; **d** 312 in “A” trend; **e** 413 in “M” trend; **f** 413 in “A” trend. The trend for “M” elements (Sc, Ti, V, Cr, Zr, Nb, Mo, Hf and Ta) and “A” elements (Al, Si, P, S, Ga, Ge, As, In, Sn, Tl and Pb) are along the x-axis in *upper* and *lower* panels, respectively. Each differently colored subpanel contains 22 and 18 MAX phases for the *top* and *bottom* respectively

Hf, Ta) for “M” and (Al, Si, P, S, Ga, Ge, As, In, Sn, Tl, Pb) for “A”. Further, each column is separated by vertical blocks of differently shaded colors. Each colored area encloses 22 MAX phases with different “A” in the upper panel and 16 MAX phases with different ‘M’ in the lower panel. (3) The ordering of both ‘A’ and ‘M’ in each block is in the order of increasing Z. The two panels contain the same number of data points but they are plotted in different ways to facilitate the observation of trends. (4) The vertical scale (0 to 20 $\text{Wm}^{-1}\text{K}^{-1}$) for κ_{ph} is kept the same for easy comparison. Despite the overwhelming amount of data, the following trends and observations can be easily discerned with this creative display: (a) The MAX phases with five highest κ_{ph} at 1300 K are all nitrides. (b) The Sc-based MAX phases (first panel to the left in the upper panel) in 211 phases are more widely dispersed than those from the 312 and 413 phases. The κ_{ph} for carbides are much smaller than those of nitrides. For the other M panels, nitrides have a lower κ_{ph} than carbides except in the 211 MAX phases where they are mixed. There is also the obvious trend of reduced κ_{ph} as Z increases. This observation is much more pronounced in the variation of “A” for a given “M” than on the variation of “M” for a given “A”. Thus, variation in “A” is the major controlling factor for κ_{ph} . They also show more distinctive separations between carbides and nitrides. The trend of reduced κ_{ph} as the Z value of “M” increases is much less pronounced. This is contrary to the notion that “M” should be more influential than “A” since there are more “M” elements in a given MAX phase. (c) In the panels for different “A”, data for carbides and nitrides are rather scattered. Within each panel for a fixed “A”, the trends of decreasing κ_{ph} with increasing Z value of “M” is much less pronounced in striking contrast with

that data for the variation of “A” for a fixed “M”. (d) Data for κ_{ph} are more widely distributed and have larger values in 211 phases. They generally scale inversely with layer index n . The calculated data at 1300K shown in Fig. 10.9 are in reasonable agreement with the only available experimental data on eight MAX phases (Ti_2AlC , Nb_4AlC_3 , Ta_4AlC_3 , Nb_2AlC , Nb_2SnC , Ta_2AlC , Cr_2AlC , and Ti_3SiC_2) [12].

We used the same data to estimate the intrinsic minimum thermal conductivity κ_{min} in MAX phases which is the lowest value of thermal conductivity for perfect crystals at high temperature when phonons are completely uncoupled and energy is transferred between neighboring atoms over the Debye temperature [24]. According to the simple theory advanced by Clarke, κ_{min} is given by [26]:

$$\kappa_{min} = k_B v_m \Lambda_{min} = k_B v_m \left(\frac{M}{n \rho N_A} \right)^{-2/3} \quad (10.2)$$

where Λ_{min} , N_A , and ρ are the phonon mean free path, Avogadro’s constant, and crystal density respectively. The calculated κ_{min} values are consistent with those using (10.1) at $T = 2000$ K.

10.4.2 Universal Elastic Anisotropy in MAX Phases

Recently, Ranganathan and Ostoja-Starzewski [27] developed a new theory for the universal elastic anisotropy A^U for all types of crystals. This gives a single parameter to quantify the crystal anisotropy similar to the Zener anisotropy index [28] which is applicable only to cubic crystals. A^U is given by:

$$A^U = 5(G^V/G^R) + (K^V/K^R) - 6. \quad (10.3)$$

Here K and G are the bulk and shear moduli and the superscripts V and R stand for the Voigt and Reuss approximations [29, 30], respectively. The Voigt (Reuss) approximation assumed a uniform strain (stress) distribution throughout the structure. These two assumptions give the upper and lower limits of bulk mechanical properties and the averaged value of these two limits is the Hill approximation [31] which is usually the value used to compare with measured data. A^U must be positive, and is usually much less than 2.0, and seldom goes beyond 4.0 [27]. $A^U = 0$ implies zero anisotropy. The large data base of elastic coefficients for MAX phases is ideally suited to evaluate A^U , test the new theory, and to ascertain its efficacy when applied to a single class of ternary hexagonal compounds, the MAX phases.

We have recently calculated the universal elastic anisotropy of the 665 MAX phases according to (10.3) [32]. Figure 10.10 shows the scattered plot of A^U versus the total bond order density (TBOD) which we advocate as the single most important metric for the electronic part of the properties. Here, A^U is supposed to be a single parameter that describe the anisotropy in mechanical properties for a given crystal. As can be seen, the majority of the MAX phases have a low A^U of less than 0.5 although




Fig. 10.10 Universal elastic anisotropy (A^U) versus total bond order density (TBOD) for 665 MAX carbides and nitrides in the database. There is evidence of a bimodal distribution with a minimum in A^U corresponding to a TBOD near 0.035

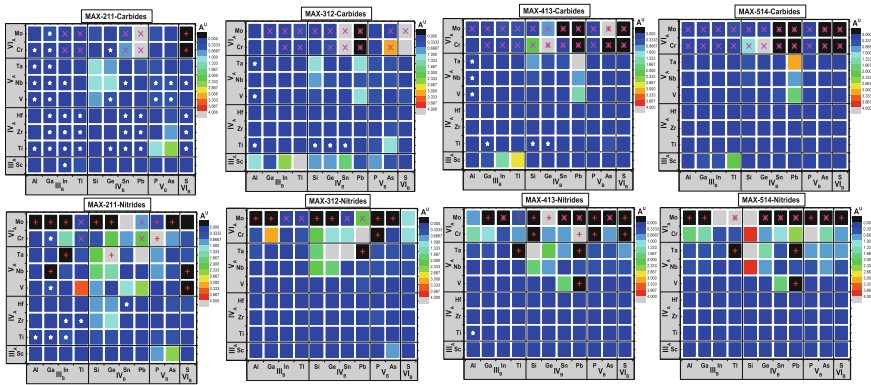


Fig. 10.11 Universal elastic anisotropy (A^U) maps for 792 MAX phases according to “M” (y-axis) and “A” (x-axis) elements. The description is the same as for Fig. 10.4 for the G/K map

some phases have A^U greater than 1.0. There is no apparent difference between MAX carbides and MAX nitrides in the A^U distribution, but there is evidence of a bimodal distribution of the data, i.e. there is a broad minimum in the middle range of the TBOD. The implication of this interesting result has yet to be explored.

Figure 10.11 shows the A^U map for all the MAX phases according to “M” (Y-axis) and “A” (X-axis) elements similar to Fig. 10.4. Color in the square cell represents calculated A^U values as indicated in the color bar. Again, a star in the box indicates that the phase has been confirmed. The symbol “+” stands for elastic instability and “x” indicates that the phase is screened out for positive heat of formation. This map

clearly shows with a glance that most of the MAX phases have low A^U close to 0.1–0.5 and that all of the confirmed phases have low A^U . Few of the phases with high A^U are easily identified. Once more, we used this innovative map of A^U to easily and clearly show the general trends in universal elastic anisotropy according to this new theory and to identify some isolated MAX phases as outliers.

10.5 Extension to Other Materials Systems

10.5.1 MAX-Related Systems, MXenes, MAX Solid Solutions, and Similar Layered Structures

A newly discovered class of 2D materials, labeled MXenes, presents a unique opportunity for the development of exceptional functional properties with diverse applications [33, 34]. MXenes are anisotropic laminated transition metal compounds derived from predecessor MAX phases. Very recently, a family of 2D materials was derived from MAX phases by extracting the A element in MAX. These so-called MXenes ($M_{n+1}X_nT_x$), have their surface terminated by T_x (O, OH, or F). Only a few MXenes out of a large number of possibilities have been reported. Some of these MXenes have high electric conductivities with hydrophilic surfaces making them ideal for applications as electrodes in Li-ion batteries or Li-ion capacitors. Demonstration of spontaneous intercalation of cations of different sizes and charges (Li^+ , Na^+ , K^+ , Mg^{2+} , Al^{3+} , $(NH_4)^+$, N_2H_4) between 2D $Ti_3C_2T_x$ surfaces in various salt solutions presents a great opportunity for material tunability. Such a wide range of choices of cation intercalations and the rich chemistry of the functionalized surfaces offer truly unique functional applications beyond the limits of even the most advanced 2D structures. The underlying factors governing the design and syntheses of MXenes however is largely unknown at atomistic-scales. Thus, MXenes are an ideal system to apply the techniques used in materials informatics similar to the ones we used for the MAX phases in creating a database for $M_{n+1}C_n$, $M_{n+1}N_n$, and $M_{n+1}AT_x$.

Another area is to consider extending MAX phases to their solid solutions. This offers a much larger variation in composition range and in the fine tuning of the desired properties. MAX solid solutions can be formed by partial substitutions in “M”, “A” elements or between C and N. There have not been many accurate calculations on the MAX solid solutions because such calculations require significant computational resources. Solid solutions are no longer crystalline phases with well-defined long range order. They are essentially a class of disordered solids with random site substitutions. A large number of supercells must be used to properly describe the structure and property variations with composition x . Thus MAX solid solutions offer another great opportunity to apply the methods of materials informatics for extensive studies. The database for MAX phases and their various applications in data mining schemes demonstrated above can facilitate the effort to investigate MAX solid solutions. We have recently carried out a detailed investigation on one of the most

important MAX solid solution phases, $\text{Ti}_2\text{Al}(\text{C}_x\text{N}_{1-x})$ [35]. For this solid solution, the mechanical properties vary continuously with x between Ti_2AlN and Ti_2AlC with no evidence of improved mechanical parameters beyond the end members. They do have subtle variations for $x > 0.5$ which are supported by some existing experimental observations. This does not rule out the possibility of strengthening MAX phases in other solid solutions via substitutions in “M” or “A” elements.

In addition to MAX solid solutions, another route to significantly enlarge the database of MAX or MAX-like compounds is to consider the quaternary alloys by adding another metal element to create a new crystal structure with different crystal symmetry. A recent example of achieving this goal is the theoretical suggestion of a new compound $(\text{Cr}_2\text{Hf})_2\text{Al}_3\text{C}_3$ [36]. The crystal structure and elastic properties of this MAX-like compound are studied using similar computational methods as in the MAX phases. Unlike MAX phases which have a hexagonal symmetry (space group: $P6_3/mmc$, #194), $(\text{Cr}_2\text{Hf})_2\text{Al}_3\text{C}_3$ crystallizes in the monoclinic structure with a space group of $P2_1/m$ (#11) with lattice parameters $a = 5.1739 \text{ \AA}$, $b = 5.1974 \text{ \AA}$, $c = 12.8019 \text{ \AA}$; $\alpha = \beta = 90^\circ$, $\gamma = 119.8509^\circ$. The calculated total energy per unit cell for this crystal is found to be energetically more favorable than potential competing phases. The calculated total energy per formula unit of -102.11 eV is significantly lower than those of the allotropic segregation (-100.05 eV) and solid solution phases (-100.13 eV). Calculations using a stress versus strain approach and the VRH approximation for polycrystals show that $(\text{Cr}_2\text{Hf})_2\text{Al}_3\text{C}_3$ has outstanding elastic moduli, better than Cr_2AlC or Hf_2AlC . Obviously, this approach can be used to explore many more new phases with some exotic properties. It is probably premature to apply techniques of materials informatics to study quaternary MAX-like compounds at this stage unless new creative algorithms can be designed.

10.5.2 CSH-Cement Crystals

Cement materials represent another system that is highly amenable for materials informatics. They are very complicated in both composition and structure but have well defined industrial standards for desirable attributes in properties. Most of all, they are of great importance and relevance to the construction industry, the environment, and the world economy. Calcium silicate hydrate (CSH) is the main binding phase of Portland cement, the single most important structural material in use worldwide. Due to the complex structure and chemistry of CSH, accurate computational studies at the atomic level are almost non-existent. Recently, we studied the electronic structure and bonding of a large subset of the known CSH minerals [37]. Table 10.3 lists the 20 CS and CSH crystal phases with well-documented atomic positions used in this study. They are divided into four groups according to the Strunz scheme [38]: **a**, Clinker and hydroxide phase; **b**, nesosubsilicates; **c**, sorosilicates; and **d**, ionosilicates. Each group in Table 10.3 is arranged in ascending order of calcium to silicon (C/S) ratio. The clinker phases (a.1 and a.2 with no H) and the Portlandite (a.3 with no Si) are

Table 10.3 List of 20 CS and CSH crystals divided into four groups based upon Strunz classification as discussed in our previous study [37]

Mineral Name	Sym/space group	Chemical formula	Ca/Si	ρ (g/cc)	TBOD
(a) Clinker/Hydroxide					
a.1 Belite	O: <i>P 1 21/m 1</i>	2(CaO) SiO ₂	2.00	3.316	0.0226
a.2 Alite	M: <i>P-1</i>	3(CaO) SiO ₂	3.00	3.084	0.0197
a.3 Portlandite	M: <i>P-3 m 1</i>	Ca(OH) ₂	N/A	2.668	0.0177
(b) Nesosilicates					
b.1 Afwillite	M: <i>P1</i>	Ca ₃ (SiO ₃ OH) ₂ ·2H ₂ O	1.50	2.590	0.0243
b.2 α -C2SH	O: <i>P 21/b 21/c 21/a</i>	Ca ₂ (HSiO ₄)(OH)	2.00	2.693	0.0216
b.3 Dellite	Tc: <i>P-1</i>	Ca ₆ (Si ₂ O ₇)(SiO ₄)(OH) ₂	2.00	2.929	0.0210
b.4 Ca Chondrodite	M: <i>P 1 1 21/b</i>	Ca ₅ [SiO ₄] ₂ (OH) ₂	2.50	2.828	0.0194
(c) Sorosilicates					
c.1 Rosenhahnite	Tc: <i>P-1</i>	Ca ₃ Si ₃ O ₈ (OH) ₂	1.00	2.874	0.0241
c.2 Suolunite	O: <i>F d 2 d</i>	Ca ₂ [Si ₂ O ₅ (OH) ₂] ₂ H ₂ O	1.00	2.649	0.0255
c.3 Kilehoanite	O: <i>I 2 c m</i>	Ca ₆ (SiO ₄)(Si ₃ O ₁₀)	1.50	2.937	0.0209
c.4 Killalaite	M: <i>P 1 21/m 1</i>	Ca _{6.4} (H _{0.6} Si ₂ O ₇) ₂ (OH) ₂	1.60	2.838	0.0217
c.5 Jaffeite	Tg: <i>P 3</i>	Ca ₆ [Si ₂ O ₇](OH) ₆	3.00	2.595	0.0192
(d) Inosilicates					
d.1 Nekoite	Tc: <i>P1</i>	Ca ₃ Si ₆ O ₁₅ ·7H ₂ O	0.50	2.204	0.0240
d.2 T11 Å	M: <i>B 1 1 m</i>	Ca ₄ Si ₆ O ₁₅ (OH) ₂ ·5H ₂ O	0.67	2.396	0.0248
d.3 T14 Å	M: <i>B 1 1 b</i>	Ca ₅ Si ₆ O ₁₆ (OH) ₂ ·7H ₂ O	0.83	2.187	0.0226
d.4 T 9 Å	Tc: <i>C1</i>	Ca ₅ Si ₆ O ₁₇ ·5H ₂ O	0.83	2.579	0.0247

(continued)

Table 10.3 (continued)

	Mineral Name	Sym/space group	Chemical formula	Ca/Si	ρ (g/cc)	TBOD
d.5	Wollastonite	Tc: P-1	$\text{Ca}_3\text{Si}_3\text{O}_9$	1.00	2.899	0.0224
d.6	Xonotlite	Tc: A-1	$\text{Ca}_6\text{Si}_6\text{O}_{17}(\text{OH})_2$	1.00	2.655	0.0214
d.7	Foshagite	Tc: P-1	$\text{Ca}_4(\text{Si}_3\text{O}_9)(\text{OH})_2$	1.33	2.713	0.0211
d.8	Jennite	Tc:P-1	$\text{Ca}_9\text{Si}_6\text{O}_{18}(\text{OH})_6 \cdot 8\text{H}_2\text{O}$	1.50	2.310	0.0223

placed in **group a**. Portlandite is included in this group because it forms the basis for hydration of cement.

Our results reveal a wide range of contributions from each type of bonding, especially the hydrogen bonding. We find that the total bond order density (TBOD) is again an ideal overall metric for assessing crystal cohesion of these complex materials and should replace the conventionally used Ca/Si ratio. A rarely known orthorhombic phase Suolunite is found to have higher cohesion (TBOD) in comparison to Jennite and Tobermorite, which are considered to be the backbone of hydrated Portland cement [37, 39, 40].

Obviously the crystalline CSH phases listed in Table 10.3 can be greatly expanded to include additional elements such as Al in the Ca-Si-Al-hydrates. A large database of cement crystals similar to the MAX phases can be built for materials informatics to design new construction materials which are more economical, environmentally friendly, and durable. This is another example of using the TBOD as a proper descriptor for materials design.

10.5.3 Extension to Other Materials Systems: Bulk Metallic Glasses and High Entropy Alloys

Another promising system for materials informatics are bulk metallic glasses (BMGs) [41] and the related high entropy alloys (HEAs) [42]. Metallic glasses are a special class of non-crystalline solid that are completely different from crystalline metals due to the lack of long-range order. They have many excellent properties and significant potential as next-generation structural materials. However, there is a lack of fundamental understanding about the structure and dynamics of BMGs at the atomic and electronic level despite many years of intense research. Many of the fundamental issues in BMGs require accurate data that can only be obtained by first-principles calculations. Detailed information about the atomic-scale interactions and their implications on the short-range and medium range orders are still missing. Current research efforts appear to focus mostly on the geometrical analysis of structures to explain the mechanical properties, deformation behavior, glass forming ability, etc. We again advocate for the use of TBOD from high quality electronic structure calculations as a useful theoretical metric to characterize the overall properties of a BMG which can be correlated with glass forming ability and other physical properties. The challenges we face are the requirement for both the accuracy and the size of BMG models and the large number of models that are needed to reach valid conclusions. Most conventional BMGs are either binary (e.g. Zr_xCu_{1-x} and Ni_xNb_{1-x}) or ternary alloys such as $Zr_xCu_yAl_z$. However, there are BMGs with more than 3 or 4 components such as $Zr_{41.2}Ti_{13.8}Cu_{12.5}Ni_{10.0}Be_{22.5}$ (Vitreloy) [43]. In these multi component BMGs, accurate *ab initio* modeling is *sine qua non* because any classical molecular dynamic simulation are infeasible due to lack or inadequacy

of appropriate potentials. The dependence of BMGs on the specific composition requires a large number of calculations to validate any hypothesis.

High-entropy alloys (HEAs) represent another class of systems which are ideal for using a materials informatics approach. Unlike the traditional alloys based on the principal elements (Fe, Ni, Cu, Ti, Zr, Al etc.) as the matrix, HEA is essentially an n -component alloy system with $5 \leq n \leq 13$. The % of major (minor) component X_i (X_s) where we have $5\% \leq X_i \leq 35\%$ and $X_s \leq 5\%$. High-entropy implies high n . The compositional possibilities for HEAs are almost unlimited. They have attracted a great deal of attention in recent years as replacements for traditional alloys such as Ni_3Al , Ti_3Al which have reached their ultimate limit of materials performance. Many new applications in different industrial and medical areas require alloys with special properties such as high hardness and strength at high temperature, resistance to wear and oxidation, low thermal conductivity, special magnetic properties, and easy formation of nanoparticles. The main effects offered by HEA are thermodynamics (the high entropy effect), the dynamic effect, the lattice distortion effect due to different sizes of the elements, and the effect due to interatomic interactions (the so-called cocktail effect).

A major difference between HEAs and BMGs is that the underlying structure for HEAs is crystalline, mostly in fcc lattice or a mixture of fcc and bcc lattice even though both HEAs and BMGs are disordered alloys. HEAs are more suitable for the systematic application of materials informatics tools because the structural part of the alloy is much simpler to model than in BMGs. On the other hand, the challenge is the enormous number of compositional possibilities which will make the database extremely large.

10.6 Conclusions

In this Chapter, we have discussed the construction and analysis of a large database for a unique class of materials, MAX phases, and we have articulated a specific approach for using *ab initio* data for materials informatics. What we have learned is that materials informatics is extremely useful but also that it faces a lot of challenges. Our approach will need a large amount of computational resources depending on the systems to be studied, but creative planning and targeted application together with the ways and means the data are presented are very important. A data mining approach can be very effective for accelerating the database generation as exemplified by the MAX phase study. The selective process of establishing internal links amongst the potential descriptors is the key. We have also found that the total bond order density (TBOD) is a very useful descriptor for analyzing a variety of properties and in their interpretations. We also described several other material systems that can employ the similar approach for materials informatics research because they share some common attributes with the MAX phase and they also have well-defined descriptors.

Acknowledgments I acknowledge with thanks the contributions and assistance from Drs. Sitaram Aryal, Yuxiang Mo and Liaoyuan Wang; Professors Michel W. Barsoum, Ridwan Sakidja, and Paul Rulis; Mr. Chamila C. Dharmawardhana, and Mr. Chandra Dhakal. This work was supported by the National Energy Technology Laboratory (NETL) of the U.S. Department of Energy (DOE) under Grant No. DE-FE0005865. This research used the resources of the National Energy Research Scientific Computing Center (NERSC) supported by the Office of Basic Science of DOE under Contract No. DE-AC03-76SF00098.

References

1. V. Vapnik, *The Nature of Statistical Learning Theory* (Springer Science & Business Media, New York, 2000)
2. K. Rajan, Materials informatics. *Mater. Today* **8**(10), 38–45 (2005)
3. R.F. Service, Science materials scientists look to a data-intensive future. *Science* **335**(6075), 1434–1435 (2012)
4. P. Jiang, X.S. Liu, Big data mining yields novel insights on cancer. *Nat. Genet.* **47**(2), 103–104 (2015)
5. P.V. Balachandran, S.R. Broderick, K. Rajan, Identifying the ‘inorganic gene’ for high-temperature piezoelectric perovskites through statistical learning. *Proc. R. Soc. A* **467**, 2271–2290 (2011)
6. M. Nishijima et al., Accelerated discovery of cathode materials with prolonged cycle life for lithium-ion battery. *Nat. Commun.* **5** (2014)
7. J. Carrete et al., Finding unprecedentedly low-thermal-conductivity half-Heusler semiconductors via high-throughput materials modeling. *Phys. Rev. X* **4**(1), 011019 (2014)
8. Y. Saad et al., Data mining for materials: computational experiments with AB compounds. *Phys. Rev. B* **85**(10), 104104 (2012)
9. A.W. Bosse, E.K. Lin, Polymer physics and the materials genome initiative. *J. Polym. Sci. Part B: Polym. Phys.* **53**(2), 89 (2015)
10. S. Broderick et al., An informatics based analysis of the impact of isotope substitution on phonon modes in graphene. *Appl. Phys. Lett.* **104**(24), 243110 (2014)
11. S. Aryal et al., A genomic approach to the stability, elastic, and electronic properties of the MAX phases. *Phys. Status Solidi (b)* **251**(8), 1480–1497 (2014)
12. M.W. Barsoum, *MAX Phases: Properties of Machinable Ternary Carbides and Nitrides* (Wiley, New York, 2013)
13. S.F. Pugh, XCH. Relations between the elastic moduli and the plastic properties of polycrystalline pure metals. *Lond. Edinb. Dublin Philos. Mag. J. Sci.* **45**(367), 823–843 (1954)
14. Y. Mo, P. Rulis, W.Y. Ching, Electronic structure and optical conductivities of 20 MAX-phase compounds. *Phys. Rev. B* **86**(16), 165122 (2012)
15. L. Wang, P. Rulis, W.Y. Ching, Calculation of core-level excitation in some MAX-phase compounds. *J. Appl. Phys.* **114**, 023708 (2013)
16. J. Hafner, J. Furthmüller, G. Kresse, *Vienna Ab-initio Simulation Package (VASP)* (1993), <http://www.vasp.at/>
17. M. Born, K. Huang, *Dynamical Theory of Crystal Lattices* (Clarendon Press, Oxford, 1956)
18. W.Y. Ching, P. Rulis, *Electronic Structure Methods for Complex Materials: The Orthogonalized Linear Combination of Atomic Orbitals*. (Oxford University Press, Oxford, 2012) p. 360
19. R. Ahuja et al., Structural, elastic, and high-pressure properties of cubic TiC, TiN, and TiO. *Phys. Rev. B* **53**(6), 3072–3079 (1996)
20. S.R. Nagel, J. Tauc, Nearly-free-electron approach to the theory of metallic glass alloys. *Phys. Rev. Lett.* **35**(6), 380–383 (1975)
21. M.W. Barsoum, *MAX Phases: Properties of Machinable Ternary Carbides and Nitrides* (Wiley-VCH, Weinheim, 2013)

22. M. Hall et al., The WEKA data mining software: an update. *ACM SIGKDD Explor. Newsl.* **11**(1), 10–18 (2009)
23. C. Dhakal, R. Sakidja, S. Aryal, W.Y. Ching, Calculation of lattice thermal conductivity of MAX phases. *J. Eur. Ceram. Soc.* **35**(12), 3203–3212 (2015)
24. D.T. Morelli, G.A. Slack, High lattice thermal conductivity solids, *High Thermal Conductivity Materials* (Springer, Berlin, 2006), pp. 37–68
25. C.L. Julian, Theory of heat conduction in rare-gas crystals. *Phys. Rev.* **137**(1A), A128 (1965)
26. D.R. Clarke, Materials selection guidelines for low thermal conductivity thermal barrier coatings. *Surf. Coat. Technol.* **163**, 67–74 (2003)
27. S.I. Ranganathan, M. Ostoja-Starzewski, Universal elastic anisotropy index. *Phys. Rev. Lett.* **101**(5), 055504 (2008)
28. C. Zener, *Elasticity and Anelasticity of Metals* (University of Chicago press, Chicago, 1948)
29. W. Voigt, *Lehrbuch Der Kristallphysik (mit Ausschluss Der Kristallogoptik)*. 1928: B.G. Teubner
30. A. Reuss, Berechnung der Fließgrenze von Mischkristallen auf Grund der Plastizitätsbedingung für Einkristalle. *ZAMM—J. Appl. Math. Mech./Zeitschrift für Angewandte Mathematik und Mechanik* **9**(1), 49–58 (1929)
31. R. Hill, The elastic behaviour of a crystalline aggregate. *Proc. Phys. Soc. Sect. A* **65**(5), 349 (1952)
32. C.C. Dharamawardhana, W.Y. Ching, *Universal Elastic Anisotropy in MAX Phases* (unpublished)
33. M.R. Lukatskaya et al., Cation intercalation and high volumetric capacitance of two-dimensional titanium carbide. *Science* **341**(6153), 1502–1505 (2013)
34. M. Naguib et al., 25th Anniversary article: MXenes: a new family of two-dimensional materials. *Adv. Mater.* **26**(7), 992–1005 (2014)
35. S. Aryal, R. Sakidja, L. Ouyang, W.-Y. Ching, Elastic and electronic properties of $\text{Ti}_2\text{Al}(\text{C}_{1-x}\text{N}_x)$ solid solutions. *J. Eur. Ceram. Soc.* **35**(12), 3219–3227 (2015)
36. Y. Mo, S. Aryal, P. Rulis, W.Y. Ching, Crystal structure and elastic properties of hypothesized MAX phase-like compound $(\text{Cr}_2\text{Hf})_2\text{Al}_3\text{C}_3$. *J. Am. Ceram. Soc.* **97**(8), 2646–2653 (2014)
37. C.C. Dharmawardhana, A. Misra, W.Y. Ching, Quantum mechanical metric for internal cohesion in cement crystals. *Sci. Rep.* **4**, 7332 (2014)
38. H. Strunz, *Mineralogische Tabellen* (Akad. Verl.-Ges. Geest u. Portig, Leipzig, 1982)
39. C.C. Dharmawardhana et al., Role of interatomic bonding in the mechanical anisotropy and interlayer cohesion of CSH crystals. *Cem. Concr. Res.* **52**, 123–130 (2013)
40. I.G. Richardson, The calcium silicate hydrates. *Cem. Concr. Res.* **38**, 137–158 (2008)
41. J. Schroers, Bulk metallic glasses. *Phys. Today* **66**(2), 32–37 (2013)
42. J.W. Yeh et al., Nanostructured high-entropy alloys with multiple principal elements: novel alloy design concepts and outcomes. *Adv. Eng. Mater.* **6**(5), 299–303 (2004)
43. A. Peker, W.L. Johnson, A highly processable metallic glass: $\text{Zr}_{41}\text{Ti}_{13}\text{Cu}_{12}\text{Ni}_{10}\text{OBe}_{22}$. *J. Appl. Phys. Lett.* **63**(17), 2342–2344 (1993)

Chapter 11

Symmetry-Adapted Distortion Modes as Descriptors for Materials Informatics

Prasanna V. Balachandran, Nicole A. Benedek and James M. Rondinelli

Abstract In this paper, we explore the application of symmetry-mode analysis for establishing structure-property relationships. The approach involves describing a distorted (low-symmetry) structure as arising from a (high-symmetry) parent structure with one or more static symmetry-breaking structural distortions. The analysis utilizes crystal structure data of parent and distorted phase as input and decomposes the distorted structure in terms of symmetry-adapted distortion-modes. These distortion-modes serves as the descriptors for materials informatics. We illustrate the potential impact of these descriptors using perovskite nickelates as an example and show that it provides a useful construct beyond the traditional tolerance factor paradigm found in perovskites to understand the atomic scale origin of physical properties, specifically how unit cell level modifications correlate with macroscopic functionality.

11.1 Introduction

One of the common objectives in the paradigm of materials informatics is the robust formulation of structure-property relationships. In materials informatics, normally, the “properties” of interest (e.g. Curie temperature, melting point, tensile strength, ductility, conductivity, polarization, hysteresis etc.) that we intend to optimize are well defined. However, what constitutes a “structure” is often not clear *a priori* and remains an outstanding issue. Note that in this paper, we restrict the scope of the

P.V. Balachandran (✉)

Theoretical Division, Los Alamos National Laboratory, Los Alamos 87545, USA
e-mail: pbalachandran@lanl.gov

N.A. Benedek

Department of Materials Science and Engineering, Cornell University, Ithaca 14853, USA
e-mail: nab83@cornell.edu

J.M. Rondinelli

Department of Materials Science and Engineering, Northwestern University,
Evanston 60628, USA
e-mail: jrondinelli@northwestern.edu

definition of “structure” to crystal structures, i.e. spatial arrangement of atoms in one, two or three-dimensions. Generally, the phenomenologically-derived descriptors (also referred to as features), e.g. Shannon’s ionic radius, Pauling’s electronegativity, pseudopotential radii of atomic orbitals and Pettifor’s chemical scale (to name a few), are utilized at a coarse-grained level to represent local or crystal structures and crystal chemistries. There are a number of reports in the literature, where structure-property relationships have been formulated for *bulk* materials using these descriptors that have even led to the discovery of new materials [1–6].

Although successful, these descriptors lack some desired characteristics. For example, when using these descriptors it is difficult to separate two materials that have the same chemical formula but different crystal symmetries (unless one appends the crystal symmetry as a separate feature). Similarly, owing to the remarkable progress in achieving high-quality coherent thin films, heterostructures and superlattices, heteroepitaxial synthesis has evolved into a reliable strategy to engineer new materials. In such ultra-thin films, strain fields at the thin film-substrate interface directly tune the local electronic states from which novel functionalities and phases prohibited or are absent in bulk materials are stabilized. Once again, the aforementioned descriptors fail under these contexts. Clearly, there is a need to develop more refined descriptors that carry physically relevant information, so that the constructed structure-property relationships not only merely reflect statistical correlations, but also provide avenues to probe mechanistic insights for better understanding. And this is not a trivial task.

Recently, computational codes based on *ab initio* [7–9] and classical methods [10] have also been explored for descriptor development. This approach is desirable, because these approaches contain the essential physics and enable rigorous materials modeling, which are absent in the phenomenological descriptors. Having said that, the cost of running expensive computer simulations on large systems could prove prohibitive and it is important to be wary of this shortcoming.

11.2 Distortion Modes as Descriptors

In this paper, we focus on developing descriptors based on distortion-mode decomposition analysis (or symmetry-mode analysis) that provide a rigorous basis set for studying crystal structures. Particularly, these descriptors are best-suited for problems (e.g. ferroelectricity, piezoelectricity, shape memory effect, ion transport etc..) in condensed matter systems that rely on structure-based materials design. In such materials, the ability to deterministically control local atomic structure would enable tuning many important electronic and structural functionalities, in turn critical for technological applications. In fact, one of the common themes is that these materials show some form of symmetry-breaking structural phase transitions and/or local structural distortions [e.g. cooperative atomic displacements (also known as “shuffles”), lattice strain, coupling between shuffles and strain].

Symmetry-mode analysis involves describing a distorted (low-symmetry) structure as arising from a (high-symmetry) parent structure with one or more static symmetry-breaking structural distortions. In the undistorted parent structure, symmetry-breaking distortion-modes have zero amplitude. The low-symmetry phase, however, will have finite amplitudes for each mode described by an irreducible representation (irrep) of the high-symmetry structure compatible with the symmetry breaking that are defined relative to specific k -points [11]. Additional details regarding distortion-mode decomposition analysis may be found in the literature [12–15].

The distortion-mode analysis is powerful, because it provides a complete and systematic basis to isolate multiple and complex distortions in crystals. By comparison of the amplitude of various modes, it is possible to directly assess each modes contribution to the mechanism underlying a structural and electronic phase transition. Furthermore, the distortion-mode analysis relies solely on crystal structure data, which enables both bulk and thin film stabilized structures with identical compositions to be evaluated on equal footing [16]. What is of particular utility in formulating quantitative structure-property relationships using distortion modes is that each irrep carries a physical representation of the displacive distortions—the unique atomic coordinates describing various symmetry-adapted structural modes. The relative importance of these modes on properties may then be mapped by means of *ab initio* computational methods or via detailed and systematic experimentation. Accessibility to computational methods make the distortion-mode analysis powerful, because it is possible to independently study various distortions and directly assess their role in structural and electronic phase transition mechanisms and macroscopic properties. Note that such direct comparison is not possible through aggregate parameters widely followed in the literature such as the tolerance factor, ionic radius, or electronegativity, i.e., when the composition is fixed. The physical basis that supports the usage of distortion-modes for materials informatics is grounded in Landau theory [17], where the free energy of a crystalline solid undergoing a phase transition from a high-symmetry parent phase to a low-symmetry distorted phase can be expressed in terms of one or more order parameters.

In this paper, we discuss the implications of symmetry-mode analysis as descriptors for materials informatics based on the perovskite structure class of materials. One of the motivations for choosing perovskites and oxides is based on the works of Benedek and Fennie [18] and Cammarata and Rondinelli [19], who used a combination of symmetry arguments and first-principles calculations to explore the connection between structural distortions and materials functionality. We have extended these guidelines to the family of perovskite nickelates (originally not considered by Benedek and Fennie [18]), where we uncover the meaning of the metric “tolerance factor” ($t = \frac{r_A+r_O}{\sqrt{2}(r_B+r_O)}$, where r_A , r_B and r_O are the Shannon’s ionic radii [20] of A , B and Oxygen elements in the ABO_3 chemical formula) and show that it encodes information that pertain to a set of the key distortion modes. It is this intriguing connection between t and distortion modes that makes t such an informative descriptor for capturing key structural and chemical trends in nickelates. Furthermore, we also show that t does not account for all distortion modes present in the ground state structure.

11.3 Perovskite Nickelates

The structure of perovskite oxides (see Fig. 11.1) are characterized by a three-dimensional network of corner-connected metal–oxygen octahedra, with alkali, alkaline-earth or lanthanide elements filling holes in the body centers of the octahedral network.

These nickelate oxides exhibit non-trivial changes in structure and physical properties, including sharp first-order temperature-driven metal to insulator transitions, unusual antiferromagnetic order in the ground state, and site- or bond-centered charge disproportionation owing to the valence and spin state flexibility of the Ni^{3+} cation [21]. Furthermore, it has been shown that both the electronic and magnetic transition temperatures can be modified by applying epitaxial strain when these materials are grown as thin films [22, 23].

In Fig. 11.2, we show the rare earth (R) cation-temperature phase diagram of bulk RNiO_3 nickelate perovskites. In our earlier work [15], we focused on two important characteristics in the phase diagram: the (i) metal to insulator transition temperature (T_{MI}) and the (ii) paramagnetic to antiferromagnetic phase transition temperatures (T_{N} , Néel Temperature). We uncovered key distortion modes and statistical correlations that govern the temperatures of the two phase transitions. One of the important findings is that R_3^+ and M_5^+ irreps capture the T_{N} trends that were previously unknown and the implications are that these distortions are not encoded in the widely recog-




Fig. 11.1 Crystal structure of an ideal cubic perovskite showing three-dimensional octahedral BO_6 connectivity with A-site cations filling the holes of the octahedral network




Fig. 11.2 Rare earth cation–temperature phase diagram of RNiO_3 perovskite nickelates [21]


nized t metric. Note that in RNiO_3 , the r_{Ni} is fixed; therefore, t for RNiO_3 is equivalent to r_R (i.e. the ionic size of the trivalent rare earth ion).

Our objective here is defined as follows, *can we use informatics to uncover the physical meaning of t in terms of the distortion modes?* We address this question by building a data set of distortion modes of known RNiO_3 perovskites (see the work of Balachandran and Rondinelli [15] for additional details on symmetry-mode analysis). We used a total of 10 RNiO_3 compounds for our analysis, where $R = \text{La}$, Nd , Pr , Tm , Lu , Dy , Er , Y , Ho and Yb . Except LaNiO_3 , all other nickelates were considered in the experimental ground state monoclinic $P2_1/c$ structure; we used the rhombohedral $R\bar{3}c$ ground state structure for LaNiO_3 .

11.3.1 Statistical Correlation Analysis

In Fig. 11.3, we show that the irreps are statistically correlated with one another, indicating that the distortions occur cooperatively in bulk nickelates. A strong positive correlation is found between irreps that describe distortions to the NiO_6 octahedra: M_2^+ , M_3^+ , X_5^+ , R_4^+ , and R_5^+ and T_{MI} . These five irreps fully describe the $Pnma$ (space group # 62) crystal structure relative to the cubic phase found in the metallic nickelates at high temperature, reinforcing the concept that the orthorhombic distortions are largely responsible for the electronic bandwidth-controlled transport behavior in nickelates. Our analyses also reveal the existence of a strong linear relationship between T_{N} and two irreps, R_3^+ and M_5^+ . The linear relationship is valid for both $T_{\text{MI}} = T_{\text{N}}$ and $T_{\text{MI}} > T_{\text{N}}$ nickelates, indicating that R_3^+ and M_5^+ contain additional

Fig. 11.3 Statistical correlation plot showing the positive (blue) and negative (red) pairwise correlation between distortion modes (M_2^+ , M_3^+ , X_5^+ , R_4^+ , R_5^+ , R_1^+ , R_3^+ and M_5^+), T_{MI} and T_N . © 2014 Reproduced with permission of the American Physical Society from [15]



information not captured by either the conventional $\langle \text{Ni-O-Ni} \rangle$ angle or tolerance factor descriptors. Although there are eight algebraically independent irreps necessary to decompose the monoclinic $P2_1/c$ phase, the presence of statistical correlation suggests redundancy—meaning, we can further reduce the complexity of the dataset and transform the statistically correlated irreps as linear combinations of one another. We used principal component analysis (PCA) to accomplish this objective.

11.3.2 Principal Component Analysis (PCA)

PCA is one of the well known linear data-dimensionality reduction methods [24]. PCA assumes that the dataset consists of a large number of intercorrelated descriptors that lie on a linear manifold. The purpose here is to reduce the dimensionality of a data set, while retaining maximum variability. This is achieved by transforming the original set of variables to a new set of derived variables, called the principal components (PCs), which are ordered so that the first few retain the most of the variation present in all of the original variables. The first PC accounts for the maximum variance (highest eigenvalue) in the dataset; the second PC is orthogonal to the first and accounts for most of the remaining variance. Thus, the m th PC is orthogonal to all others and has the m th largest variance in the set of PCs. Once all the PCs have been calculated, only those with eigenvalues above a critical level (a rule of thumb is to retain only those PCs whose eigenvalue is greater than or equal to 1) are retained.

Each PC is a linear combination of the weighted contribution of all attributes and the magnitude of the weight determines the relative impact of each descriptor in




Fig. 11.4 Scree plot showing the relative variance of each principal component (PC) from the RNiO_3 data set. The first three PC's together capture more than 95 % variance in the data set. After the third PC (circled), the relative variance captured by the subsequent PC's are small and can be ignored

affecting the PC. From the knowledge of the calculated PCs, one can determine the relative importance of each descriptor and the correlation between any two descriptors. Information pertaining to the relative importance of the descriptors will be helpful in identifying the dominant descriptors, whereas the correlation information will be helpful in screening the dominant descriptors to avoid choosing redundant descriptors. Here, we find that the first three PCs together capture 95 % of variation in the dataset. The scree plot is shown in Fig. 11.4. As a result, we have reduced the dimensionality of the dataset from 8 to 3. Therefore, we retain only the first three PCs for further consideration. In fact, the first two PC's alone capture 91 % of variation in the data.

In 11.1 and 11.2, we show the weighted contribution from the linear combination of irreps captured by PC1 and PC2, respectively.

$$\text{PC1} = -0.38R_1^+ - 0.12R_3^+ - 0.41R_4^+ - 0.42R_5^+ - 0.38X_5^+ - 0.41M_2^+ - 0.42M_3^+ - 0.06M_5^+ \quad (11.1)$$

$$\text{PC2} = 0.17R_1^+ + 0.66R_3^+ - 0.17R_4^+ - 0.10R_5^+ - 0.12X_5^+ - 0.01M_2^+ - 0.06M_3^+ - 0.69M_5^+ \quad (11.2)$$

Note that PC1 and PC2 captured 68 and 23 % of variation, respectively, in the dataset. PC1 captures descriptors associated with octahedral distortions that describe the orthorhombic crystal symmetry ($Pnma$) and also note that there is a significant contribution coming from the R_1^+ irrep that describes the octahedral breathing distortion, which is the primary order parameter for the phase transition from the paramagnetic metallic $Pnma$ structure to the paramagnetic insulating monoclinic $P2_1/c$ structure.




Fig. 11.5 Scatter plot between tolerance factor (t , y-axis) and *left* principal component 1 and *right* principal component 2. We find that principal component 1 correlates strongly with t ($R^2=0.90$), relative to principal component 2 ($R^2=0.74$)

On the other hand, in PC2 R_3^+ (Jahn-Teller distortion) and M_5^+ (out-of-phase tilting) distortions have the dominant contribution, but are orthogonal to PC1. One of the active areas of research in perovskite nickelates is to identify the mechanism responsible for the metal-to-insulator and the paramagnetic to antiferromagnetic phase transitions. Clearly, the insights hidden in the PC1 and PC2 should be rigorously explored using additional experimentation and theoretical simulations to elucidate the physical origin behind these correlations.

In Fig. 11.5, we show how PC1 and PC2 relate to the t . Clearly PC1 (Fig. 11.5a) correlates strongly with t with a correlation coefficient (R^2) of 0.9, relative to PC1 (Fig. 11.5b) whose R^2 is only modest at 0.74. Note that Fig. 11.5 also includes LaNiO_3 , whose ground state structure is $R\bar{3}c$; in sharp contrast with other nickelates whose ground state is $P2_1/c$. The key implications are that the octahedral distortions (in terms of M_2^+ , M_3^+ , X_5^+ , R_4^+ , and R_5^+) that describe the $Pnma$ symmetry and the breathing distortion (R_1^+) together collectively correlate strongly with the t metric. The full description of PC1 is given in 11.1. On the other hand, octahedral distortions described by irreps R_3^+ and M_5^+ (see 11.2) do not correlate strongly with t , indicating that the geometric t metric is much less sensitive to electronic-based structural effects such as Jahn-Teller distortions.

11.4 Summary

In summary, descriptor development is a critical component in the materials informatics research paradigm. The choice of the descriptions must be such that, in addition to helping accomplish statistical correlations between structure and property, it must provide mechanistic insights to address causal relationships. Distortion modes based on symmetry-mode analyses satisfy these requirements, which makes it very attractive for developing quantitative structure-property relationships in materials informatics.

Acknowledgments P.V.B. acknowledges funding support from the Los Alamos National Laboratory (LANL) Laboratory Directed Research and Development (LDRD) DR (#20140013DR) on Materials Informatics. J.M.R. acknowledges funding support from the NSF (DMR-1454688).

References

1. E.S. Machlin, T.P. Chow, J.C. Phillips, Structural stability of suboctet simple binary compounds. *Phys. Rev. Lett.* **38**, 1292–1295 (1977)
2. J.R. Chelikowsky, J.C. Phillips, Quantum-defect theory of heats of formation and structural transition energies of liquid and solid simple metal alloys and compounds. *Phys. Rev. B* **17**, 2453–2477 (1978)
3. P.B. Littlewood, Structure and bonding in narrow gap semiconductors. *Crit. Rev. Solid State Mater. Sci.* **11**(3), 229–285 (1983)
4. A. Zunger, Systematization of the stable crystal structure of all AB-type binary compounds: a pseudopotential orbital-radii approach. *Phys. Rev. B* **22**, 5839–5872 (1980)
5. T.R. Paudel, A. Zakutayev, S. Lany, M. d’Avezac, A. Zunger, Doping rules and doping prototypes in A_2BO_4 spinel oxides. *Adv. Funct. Mater.* **21**(23), 4493–4501 (2011)
6. P.V. Balachandran, S.R. Broderick, K. Rajan, Identifying the inorganic gene for high-temperature piezoelectric perovskites through statistical learning. *Proc. R. Soc. A: Math. Phys. Eng. Sci.* **467**(2132), 2271–2290 (2011)
7. J. Yan, P. Gorai, B. Ortiz, S. Miller, S.A. Barnett, T. Mason, V. Stevanovic, E.S. Toberer, Material descriptors for predicting thermoelectric performance. *Energy Environ. Sci.* **8**, 983–994 (2015)
8. B. Meredig, C. Wolverton, Dissolving the periodic table in cubic zirconia: data mining to discover chemical trends. *Chem. Mater.* **26**(6), 1985–1991 (2014)
9. L.M. Ghiringhelli, J. Vybiral, S.V. Levchenko, C. Draxl, M. Scheffler, Big data of materials science: critical role of the descriptor. *Phys. Rev. Lett.* **114**, 105503 (2015)
10. T. Das, T. Lookman, M.M. Bandi, A minimal description of morphological hierarchy in two-dimensional aggregates. *Soft Matter* **11**, 6740–6746 (2015)
11. B.J. Campbell, H.T. Stokes, D.E. Tanner, D.M. Hatch, ISODISPLACE: a web-based tool for exploring structural distortions. *J. Appl. Crystallogr.* **39**(4), 607–614 (2006)
12. C.J. Howard, H.T. Stokes, Group-theoretical analysis of octahedral tilting in perovskites. *Acta Crystallogr. Sect. B* **54**(6), 782–789 (1998)
13. J.M. Perez-Mato, D. Orobengoa, M.I. Aroyo, Mode crystallography of distorted structures. *Acta Crystallogr. Sect. A* **66**(5), 558–590, (2010). <http://dx.doi.org/10.1107/S0108767310016247>
14. D. Orobengoa, C. Capillas, M.I. Aroyo, J.M. Perez-Mato, AMPLIMODES: symmetry-mode analysis on the bilbao crystallographic server. *J. Appl. Crystallogr.* **42**(5), 820–833 (2009)
15. P.V. Balachandran, J.M. Rondinelli, Interplay of octahedral rotations and breathing distortions in charge-ordering perovskite oxides. *Phys. Rev. B* **88**, 054101 (2013)
16. I.C. Tung, P.V. Balachandran, J. Liu, B.A. Gray, E.A. Karapetrova, J.H. Lee, J. Chakhalian, M.J. Bedzyk, J.M. Rondinelli, J.W. Freeland, Connecting bulk symmetry and orbital polarization in strained $RNiO_3$ ultrathin films. *Phys. Rev. B* **88**, 205112 (2013)
17. J. Tolédano, P. Tolédano, *The Landau Theory of Phase Transitions*. (World Scientific, Singapore, 1987)
18. N.A. Benedek, C.J. Fennie, Why are there so few perovskite ferroelectrics?. *J. Phys. Chem. C* **117**(26), 13339–13349 (2013)
19. A. Cammarata, J.M. Rondinelli, Contributions of correlated acentric atomic displacements to the nonlinear second harmonic generation and response. *ACS Photonics* **1**(2), 96–100 (2014)
20. R.D. Shannon, Revised effective ionic radii and systematic studies of interatomic distances in halides and chalcogenides. *Acta Crystallogr. Sect. A* **32**, 751–767 (1976)
21. G. Catalan, Progress in perovskite nickelate research. *Phase Trans.* **81**, 729–749 (2008)

22. J. Chakhalian, J.M. Rondinelli, J. Liu, B.A. Gray, M. Kareev, E.J. Moon, N. Prasai, J.L. Cohn, M. Varela, I.C. Tung, M.J. Bedzyk, S.G. Altendorf, F. Strigari, B. Dabrowski, L.H. Tjeng, P.J. Ryan, J.W. Freeland, Asymmetric orbital-lattice interactions in ultrathin correlated oxide films. *Phys. Rev. Lett.* **107**, 116805 (2011)
23. J.M. Rondinelli, S.J. May, J.W. Freeland, Control of octahedral connectivity in perovskite oxide heterostructures: an emerging route to multifunctional materials discovery. *MRS Bull.* **37**, 261–270 (2012)
24. M. Ringnér, What is principal component analysis? *Nat. Biotech.* **26**(3), 303–304 (2008)

Chapter 12

Discovering Electronic Signatures for Phase Stability of Intermetallics via Machine Learning

Scott R. Broderick and Krishna Rajan

Abstract In this paper, we identify the signatures of the density of states (DOS) spectra which control the bulk modulus via a hybrid informatics driven analysis. The signatures of the DOS spectra then constitute the electronic structure fingerprint of the material. This provides an important step in the “inverse design” process because if we are able to compute bulk modulus from the DOS, then we can also compute the DOS from the bulk modulus, and in this way create a “virtual” DOS based on optimized properties. In this paper, we identify the signatures for bulk modulus, and associate the signatures with specific chemistry and crystal structure. Further, we identify the details in the electronic structure that result in Ni_3Al and Co_3Al having such different stabilities in L1_2 structure although they are seemingly isoelectronic. This paper lays out the methodology for extracting these features and has significant implications, such as in the identification of critical element substitutions, by developing a framework for accelerated and targeted materials design.

12.1 Introduction

This paper develops a template for “inverse design” of alloy chemistries, which we demonstrate for the density of states and bulk modulus of a material. The questions we are asking here is: (i) if we know the target property for a material we want, can we from that compute the chemistry and structure of the material and (ii) what signatures of the density states spectra dictate very different structural stabilities of seemingly isoelectronic systems? These issues represent an inverse logic to traditional materials design, such as through density functional theory (DFT),

S.R. Broderick · K. Rajan (✉)
Department of Materials Design and Innovation,
University at Buffalo—The State University of New York, Buffalo, NY, USA
e-mail: krajan3@buffalo.edu

© Springer International Publishing Switzerland 2016
T. Lookman et al. (eds.), *Information Science for Materials
Discovery and Design*, Springer Series in Materials Science 225,
DOI 10.1007/978-3-319-23871-5_12

where the input is the chemistry and structure, and the output is property and stability [1–4]. This provides an alternate approach to other linkages of DFT and informatics which seek to calculate properties for as many materials as possible, and provide a database that can be searched for the material closest to having the target properties [5]. We instead start with a relatively small database and require few, but clearly defined, DFT calculations. This logic further differs from the traditional definition of inverse design in materials science by going from property to condition, as opposed to defining inverse design as going from calculation to experiment [6, 7].

We have previously employed informatics for modeling the density of states (DOS) spectra as a function of the properties of constituent elements [8]. This work represented a new approach for rapidly modeling DOS spectra with an accuracy nearly equivalent to DFT calculations. Further, our prior work also demonstrated the capability of informatics for extracting signatures of the DOS spectra correlating to chemistry, crystal structure and stoichiometry [9, 10]. These prior works therefore introduced an approach for modeling DOS spectra based on modifications in the material chemistry and structure. This paper develops the next stage in the inverse design problem of connecting bulk modulus and density of states. That is, if we can extract modulus from the DOS spectra, then we can design a “virtual” DOS spectra which is optimized based on our property requirements. When connected with our prior works in connecting chemistry and DOS [8], and connecting crystal structure to DOS [9, 10], the framework is completed for going from target property to “virtual” DOS to crystal structure and chemistry, and therefore the computation of a “virtual” material with the target properties.

As the DOS represents all electronic interactions of a system, it theoretically contains information on all electronic properties [11–15]. However, the understanding of how these properties are captured by the DOS is not well understood. Therefore, another objective of this paper is in understanding the connection between DOS and property. That is, to identify what signatures related to the intensities of the DOS are controlling the material property. One example of a property which is known to be at least qualitatively represented within the DOS spectra for single elements is bulk modulus [16]. The Fermi energy (E_F) indicates the maximum occupancy by electrons at ground state conditions, with DOS values at energies greater than E_F representing unoccupied available states, while the transition from bonding to antibonding states is represented as a well-defined and extended valley [17]. Occupancy of a bonding state corresponds with an increase in strength, while additional occupancy of an antibonding state results in a decrease in strength. The bulk modulus can then be found to be related with the distance between the bonding-antibonding transition and E_F . We expand that logic here but for alloy systems.

12.2 Informatics Background and Data Processing

The ability to predict properties of new alloy systems from an input of elemental DOS requires the integration of principal component analysis (PCA) and partial least

squares (PLS). This work represents a hybrid approach because we are predicting the property as a summation of the PLS coefficients and the PCA weightings (i.e. $\text{Property} = f[\sum(\text{Component of PLS result} * \text{Component of PCA result})]$) as opposed to considering the two components independently (i.e. $\text{Property} = f[\sum(\text{Component of PLS result}) * \sum(\text{Component of PCA result})]$). That is, in the final development of an equation, the PLS and PCA components of the analysis cannot be separately extracted. This hybrid capacity of the approach is demonstrated in this paper.

The PCA serves to extract the unique patterns within the DOS spectra most correlated to the information discriminating the materials. A dimensionally reduced map can then be used to correlate the conditions of the materials to the signatures of the DOS spectra. The primary application however in this paper is the parameterization of the DOS spectra, where the parameterization is not based on a curve fitting, but rather by correlating the conditions of the material with the features. PLS is then used to develop a predictive model between these PCA derived parameters and material property, in the form of a quantitative structure-property relationship (QSPR).

The input DOS curves were calculated using the full-potential linearized augmented plane wave (FP-LAPW) method [18] within the density functional theory (DFT) [19] approach and implemented in the WIEN2K code [20]. The exchange-correlation term was determined within the generalized gradient approximation (GGA) using the scheme of Perdew and Wang [21]. DFT is based on the discovery that a relationship exists between the potential of a system and the electronic density and is able to model the electronic structure based on the relationships between these factors. The input into a DFT calculation is the chemistry and relative atom positions, and using quantum mechanical approximations DFT is able to model the electronic structure. Although the calculation is based on a k-space representation and structure is not directly involved in the calculation, we have previously shown using statistical learning that crystal structure is clearly represented within the DOS spectra [9, 10].

PCA classifies the data based on a set of orthogonalized axes (principal components) comprised of a combination of descriptors which maximize the variance in the data captured [22–25]. By applying PCA to the DOS spectra, the strongest patterns in the data can be identified in a limited number of dimensions. PCA operates by performing an eigenvector decomposition of the data. As such, the principal components (PCs) capturing the most information are associated with the largest eigenvalues of the covariance matrix and their corresponding eigenvectors. The original data is decomposed into two matrices: the scores (T) and loadings (P). The scores matrix classifies the samples, in this case different alloy chemistries and structures, as defined by their differences in the DOS. The loadings matrix contains information on how the different descriptors (here DOS at specific energy values) differentiate the samples. The PCA equation is summarized by (12.1), where E is the residual matrix, and X is the input data matrix.

$$X = T \cdot P^T + E \quad (12.1)$$

The loadings and scores matrices contain the principle patterns within the DOS curves and the scaling of those patterns to create the final DOS curve, respectively.

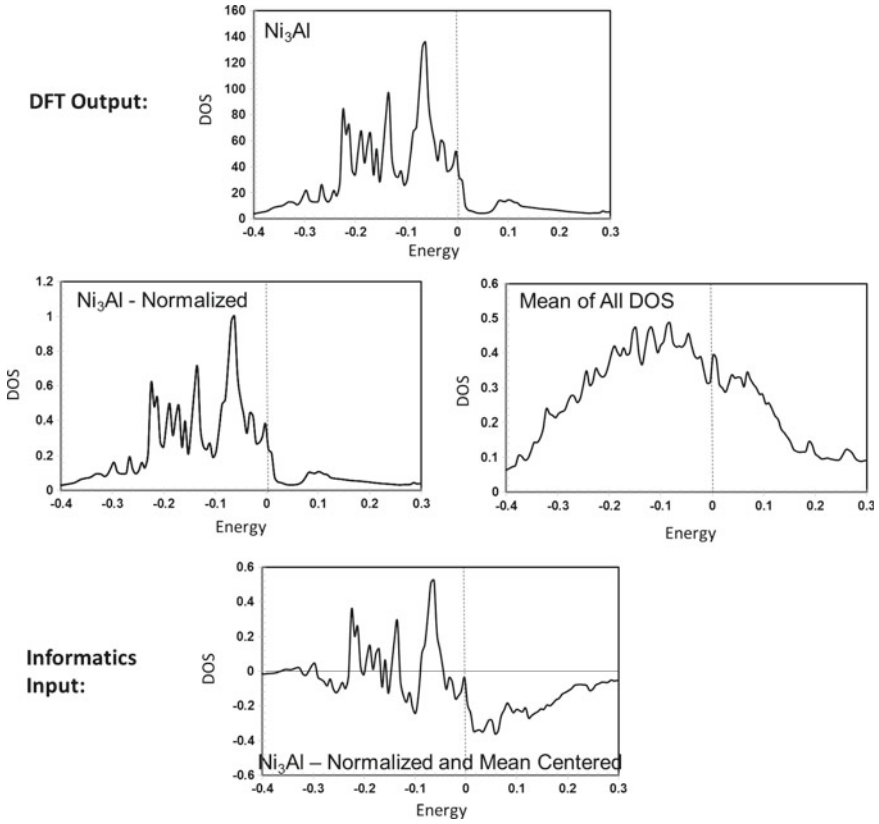


Fig. 12.1 The development of each row of the input matrix (X in (12.1)), demonstrated for Ni_3Al . Each row of X contains a unique alloy chemistry and structure. The columns of the input matrix contain every data point in the DOS curve, the rows contain different alloy systems, and the value in the matrix is the DOS, or intensity, at the specified energy. The DOS is first normalized by dividing every value by the maximum DOS value for the alloy, and then the mean of all DOS spectra at each respective energy is calculated and subtracted from the normalized spectra for each alloy. This processed (normalized and mean centered) DOS for each alloy is added as a separate row in the input matrix

The number of dimensions of T and P may be equal to the number of data points within the entire DOS curve and is on the order of hundreds of PCs in this case, although typically a significantly reduced number of PCs is sufficient for capturing the information of interest. The treatment of the input data is demonstrated in Fig. 12.1, and defines how the data from the DFT calculations is processed prior to being included in X . As an example, we show the DOS of Ni_3Al in the $L1_2$ structure. As our primary objective here is the extraction of patterns in the DOS spectra, we first normalize the DOS spectra by dividing all points by the maximum DOS value, and thereby the largest DOS value becomes unity. Then the mean for each energy value, across all DOS spectra included in the analysis, is calculated, with the DOS aligned

so that the Fermi energy (E_F) is equal to zero. This mean spectra is then subtracted from the DOS spectra. This processed DOS spectra is then used in the PCA analysis. The DOS spectra shown as the informatics input in Fig. 12.1 represents an entire row in matrix X from (12.1). This process is repeated on every system included in the analysis.

In PLS the training data is converted to a data matrix with orthogonalized axes, which are based on capturing the maximum amount of information in fewer dimensions [26–30]. The relationships discovered in the training data can be applied to a test dataset based on a projection of the data onto a high-dimensional hyperplane within the orthogonalized axis-system. Typical linear regression models do not properly account for the co-linearity between the descriptors, and as a result the isolated impact of each descriptor on the property cannot be accurately known. However, by projecting the data onto a high-dimensional space defined by orthogonal axes which are comprised of a linear combination of the spectral parameters defining the DOS curves, the impact of the descriptor on the property can be identified independent of all other descriptors.

PLS is used here to predict the relationship between spectral features and bulk modulus for different alloy chemistries. The prediction serves to create a connection between chemistry, electronic structure, and property. The PLS prediction requires two input matrices: a matrix which contains descriptors related to the input conditions (scores matrix) and a matrix containing the values which are to be predicted (bulk modulus), building a model between the input descriptors and the descriptor to be predicted. To ensure accuracy of the QSPR modeling and to verify that we are not over-fitting the data, we employ a cross validation to the predicted results. To this end, we compute both the root mean square error of calibration (RMSEC) and the root mean square error of cross validation (RMSECV). We perform a leave-one-out (LOO) cross validation and measure the accuracy of the model with and without the variable left out in the LOO approach. This step is repeated for removing each sample from the training data. That is, a model is built removing each sample, thereby ensuring that the physics captured in the model development is sufficiently robust that it can be used on new materials. The RMSEC and RMSECV values are then used to define the final predictive model. To select the number of latent variables with a suitable combination of accuracy and robustness, we define a criteria for selection of latent variables based on the ratio of $RMSECV(m)/RMSECV(m+1)$ where m is equal to the number of latent variables. From our criteria, m is selected such that it is the maximum number with the ratio below the threshold value of unity. To ensure that we are not over-fitting the data, a minimal number of parameters (PC scores values) are included so that the number of alloy chemistries is sufficiently larger than the number of parameters used as terms in the QSPR.

The systems that were modeled via DFT and are used in the analysis are listed in Table 12.1, with the bulk modulus (B) values also provided. The crystal structure type for each alloy is shown in parentheses.

Table 12.1 List of alloys modeled via DFT (crystal structure in parentheses) and used as input systems in this work

Alloy	B (GPa)
CuZn (B2)	113.8
CoTi (B2)	177.5
Ni ₃ Al (L12)	177.2
NiAl (B2)	189.0
CoAl (B2)	177.5
Co ₃ Al (L12)	123.9
NiTi (B2)	159.0
Be ₃ Co (D03)	156.5
TiAl (L10)	112.5
Fe ₃ Ni (L12)	141.0
FeNi ₃ (L12)	189.1
FeNi (L10)	183.5
FePd ₃ (L12)	192.3
Fe ₃ Pd (L12)	155.9
FePd (L10)	179.1

The modulus values are also listed, as calculated via DFT

12.3 Informatics-Based Parameterization of the DOS Spectra

The analysis described here represents a general methodology, and can therefore adapt with additional systems. As additional systems are added, the spectral patterns and subsequently the parameters will change. This flexibility in the approach is one of the primary benefits because it is robust enough to represent changes in systems or more range of possible systems. The output from the PCA is then spectral parameters for systems which we know the bulk modulus, and spectral parameters for systems which we do not know the bulk modulus. The parameters for the systems with the bulk modulus known are input into the PLS approach, and a model linking the spectral patterns and bulk modulus is then developed. This model can then be used to predict the bulk modulus of the new systems as a function of their spectral parameters. This logic is demonstrated here for bulk modulus, but should be applicable to any electronic property.

The results of the PCA on the DOS spectra is provided in Figs. 12.2 (scores plot) and 12.3 (loadings spectra). In the scores plot, we are able to extract some trends correlating to the new axis system. The first is that PC2 captures the crystal structures of the alloys, with those having L1₂ structure having positive PC2 and those having negative PC2 value being a structure besides L1₂. PC3 captures subtleties in the DOS spectra correlated with chemistry, which we observe through Co and Ni containing alloys trending towards positive PC3 and those with Fe and Pd trending towards





Fig. 12.2 PCA scores mappings of the alloys based solely on a DOS input. These values correspond with matrix T in (12.1). From these mappings of the first three PCs, we identify trends corresponding with crystal structure, chemistry, and electron valency. Our information based parameterization therefore captures variances related to these various factors which will then be represented in the development of the QSPR. The axes of the plots are defined by the loadings plots of Fig. 12.3

negative PC3. The physics captured by PC1 is harder to define, although it largely captures the relationship with d-electron valency, as those alloys with elements without d-electrons (Al and Be) generally having lower PC1 values. Therefore, the PCA analysis is able to capture subtle variations in the DOS associated with crystal structure, chemistry, and valency.

The loadings plot mathematically define the axes of the scores plots, with the axes a sum of the values at each energy but weighted according to the loadings values. Therefore, those features with larger loadings values more prominently define the axes. For example, in loadings spectra 1, the DOS at lowest energy have a negative correlation with PC1 scores value, while increasing DOS near E_F increases the PC1 value, which is also relatively insensitive to the changes at higher energy values, as seen by the loadings values closer to zero. This loadings pattern fits with our interpretation of the PC1 axis, as those elements without d-orbitals will increase the DOS at lower energies, and therefore decrease the PC1 scores value. Similar interpretation can be applied to the other PCs, as is shown for PC2 in Fig. 12.4. In this case, we identify the features in the DOS which promote the L₁₂ structure.

Fig. 12.3 The three most significant spectral patterns for differentiating the DOS spectra. These spectra represent the first three rows of matrix P in (12.1). These spectra define the axes of Fig. 12.2, and also define the physics associated with our parameterization which will be connected to bulk modulus. Further, the features of the DOS can be correlated with material conditions, such as is demonstrated in Fig. 12.4 for crystal structure



An issue when developing QSPRs is that we need the number of conditions (in this case unique chemistries and structures) to be well larger than the number of predictor variables (in this case the parameterization of the DOS spectra as represented through the PC scores values). Otherwise, the risk of over-fitting the model becomes high. To address this challenge, beyond employing the cross validation approach discussed in Sect. 12.2, we reduce the number of parameters included. The importance of each PC, as measured through variance, is listed in Table 12.2. We select six PCs as the number to include, as this represents a dimensionality lower than the number of conditions (fifteen), while capturing greater than 90 % of the total variance in the DOS spectra.




Fig. 12.4 Correlating signatures in the DOS spectra with the $L1_2$ crystal structure. In Fig. 12.2, we identified PC2 as separating $L1_2$ from other structures, and specifically with $L1_2$ structures having positive PC2 value. Increasing the DOS at the regions with positive loadings values increases the PC2 scores value, while increasing DOS at the regions with negative PC2 loadings decreases the scores value. Therefore, we find that those compounds with larger DOS values below and at the Fermi energy are more like $L1_2$, while those with higher DOS above E_F are less correlated with the $L1_2$ structure

Table 12.2 Variance captured by each PC

PC	% Variance	% Total variance
1	62.6	62.6
2	11.2	73.8
3	6.4	80.1
4	4.6	84.8
5	4.3	89.0
6	2.8	91.8
7	2.1	93.9
8	1.9	95.8
9	1.3	97.1
10	0.8	98.0
11	0.8	98.8
12	0.5	99.3
13	0.4	99.7
14	0.2	99.9
15	0.1	100.0
16	0.0	100.0

In order to reduce the number of predictor variables used in developing our QSPR, we select only the first six PCs, which capture over 90% of the information contained by the DOS spectra. We therefore include six parameters derived from the DOS in modeling bulk modulus. The information represented by the other PCs is moved to the residual matrix (E in (12.1))




Fig. 12.5 Residual spectra for Ni_3Al . This spectra represents the information captured in PC7 through PC16, and which is not included in modeling bulk modulus. By removing this residual signal, we have reduced noise or other information not contributing relevant information for modeling material properties for Ni_3Al . A comparable residual spectra is calculated for every alloy system in the analysis

In (12.1), matrices T and P therefore have six dimensions each. The loadings spectra from the other PCs after the first six are put into the residual matrix (E in (12.1)). As an example of the information in the residual matrix, the row corresponding to Ni_3Al is shown in Fig. 12.5. This information is therefore further removed from the input spectra shown in Fig. 12.1 for Ni_3Al .

We have decomposed the DOS spectra for each alloy into seven components (PC1 through PC6 and the residual spectra). The components corresponding to relevant signal for Ni_3Al are shown in Fig. 12.6, with each spectra corresponding to a different PC. The sum of these components then is the row corresponding to Ni_3Al in the matrix TP^T in (12.1). This sum also is equal to the informatics input spectra from Fig. 12.1 minus the residual matrix from Fig. 12.5. The deconvolution to these six spectra is based on the processing of the DFT output, the PCA analysis and the removal of the residual signal. These spectra are then used to determine the parameters used in extracting the modulus. The parameter is calculated as ratio of these patterns to the respective loadings spectra. For example, $T_1P_1^T$ for Ni_3Al divided by the PC1 spectra of Fig. 12.3 results in a value of -3.94 . Similar logic is used for calculating the other five parameters for Ni_3Al , and is also repeated for every alloy system.





Fig. 12.6 Ni_3Al spectra from the first six PCs. The parameters used in the modeling are then these spectra divided by the respective loadings spectra for each PC. The scaling value is then the parameter for that PC. For example, the six parameters resulting from Ni_3Al are -3.94 , 3.43 , 1.13 , -3.16 , -1.94 and 1.68 . Parameters are computed in the same way for each alloy. The collection of these parameters are then the predictor matrix for extracting the bulk modulus

12.4 Identifying the Bulk Modulus Fingerprint

Based on the predictor matrix we developed as described in Sect. 12.3, we develop a QSPR for the bulk modulus using PLS. The PLS prediction is based on correlating the features of the DOS (as represented through \mathbf{T}) with the bulk modulus. The output of the PLS model is then a coefficient matrix β and a constant C , such that the bulk modulus is defined as in (12.2), where the bulk modulus (B) of material i is a function of the product of the PLS coefficient and the scores value corresponding to each PC j . The correlation between the input mean-centered DOS curve, with the DOS input for every energy k , is then defined by (12.3).

$$B_i = \sum_{j=1}^6 \beta_j * T_{i,j} + C \quad (12.2)$$

Fig. 12.7 The result of the hybrid approach for predicting bulk modulus from an input of alloy DOS spectra. The accuracy of the results demonstrates that bulk modulus is clearly represented within the DOS spectra, and that it can be quantitatively extracted via statistical learning



$$B_i = \sum_{j=1}^6 \sum_{k=1}^{1000} \beta_j X_{i,k} P_{j,k} + C \tag{12.3}$$

Based on (12.2) and (12.3), the terms of β were calculated as 1.09, -1.65 , 1.27, -3.19 , -2.77 and 2.48, in order of j from one to six. To test and ensure the robustness of the model, the cross-validation approach described in Sect. 12.2 was utilized. The result of the model comparing the informatics modeled B with that calculated via DFT is shown in Fig. 12.7. The high accuracy of this approach shows that we are capturing the features in the DOS which control the bulk modulus.

Based on a comparison of the magnitudes of the coefficients in β , PC1 which has the fewest features impacts the bulk modulus the least. The other PCs, such as PC4, PC5 and PC6, which have higher weighting, also contain more features as seen in Fig. 12.6. For Co_3Al , the fourth and fifth patterns were most important (as determined by the value for parameter times coefficient) for determining the bulk modulus. To extract the features which most impact the modulus, we utilize (12.4), which also highlights the hybrid approach, as the PLS and PCA derived components are summed together based on the component number, and not using the two approaches individually. This therefore converts the single modulus value to a spectra with dimensionality equivalent to the number of values at unique energy intervals. The spectral values ($B_{i,k}$) are a measure of the contribution of the DOS at each energy to the bulk modulus. We therefore develop spectra which correlate the features of the DOS with the modulus, as the features with largest intensity represent the signatures of the DOS which impact the modulus.

$$B_{i,k} = \underbrace{\sum_{j=1}^6}_{\text{Hybridization}} \underbrace{\beta_j T_{i,j} K_{j,k}}_{\text{PCA Derived}} + C \tag{12.4}$$

$\begin{matrix} \text{PLS} \\ \text{Derived} \\ \downarrow \end{matrix}$

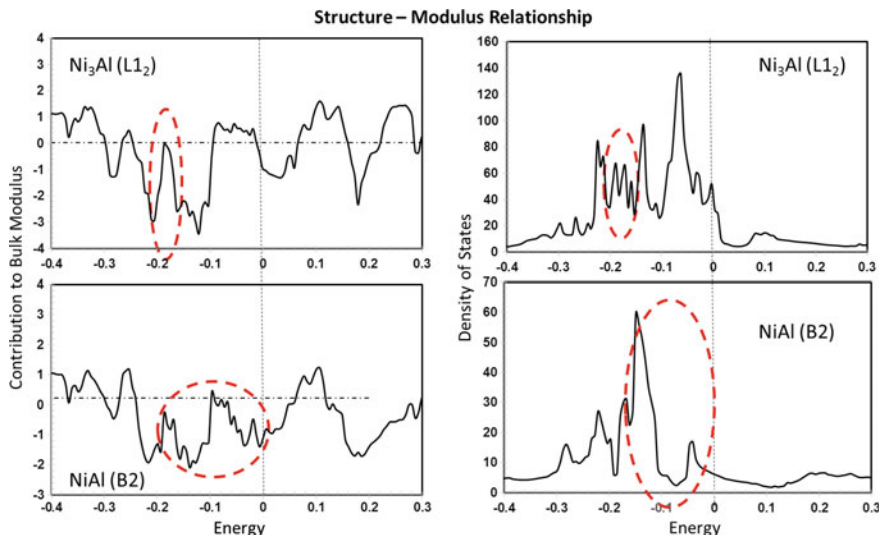


Fig. 12.8 Identification of signatures of the DOS spectra for Ni-Al alloys in $L1_2$ and B2 structure. The circled regions define the energies where the largest magnitude features are. The corresponding features of the DOS spectra at those energies are then extracted. We notice for both structures, the bonding-anti-bonding transition is identified as a feature of the DOS corresponding to bulk modulus, and therefore we identify this transition as a signature of the Ni-Al alloys

The comparison between the weighting on the bulk modulus and the DOS spectra is possible as they have the same energy values. Therefore, we can trace the energy corresponding to the signature back to the original input spectra. This is represented in Figs. 12.8 and 12.9, where we compared, respectively, the signature of Ni-Al alloys in different structures and $L1_2$ structures with different chemistries. In this way, we identify the signatures common to crystal structure and the features common to chemistry. The circled regions in these figures correspond to the highest intensity features in terms of contribution to the bulk modulus. The circled regions within the input DOS spectra are at the same energies and therefore identify the most important features of the DOS in terms of contribution to bulk modulus. In the case of changing chemistry, we identify similarities in terms of Ni-Al chemistries (structure-modulus relationship). We find that for both $L1_2$ and B2 structure, the bonding-anti-bonding transition is a prominent feature. Conversely, when comparing $L1_2$ structures but with different chemistries (chemistry-modulus relationships), we find that a doublet peak between larger peaks is identified in each case, with only the bonding-anti-bonding transition identified as a signature for Ni_3Al .

We therefore have through this work identified two signatures of the DOS for Ni_3Al and used them to explain the differences in bulk moduli of alloys, as well as the differences in stability for alloys with similar electronic structures. We have also for Ni_3Al correlated one signature to structure and two signatures to chemistry. This result is summarized in Fig. 12.10. In the case of bonding-anti-bonding transition as

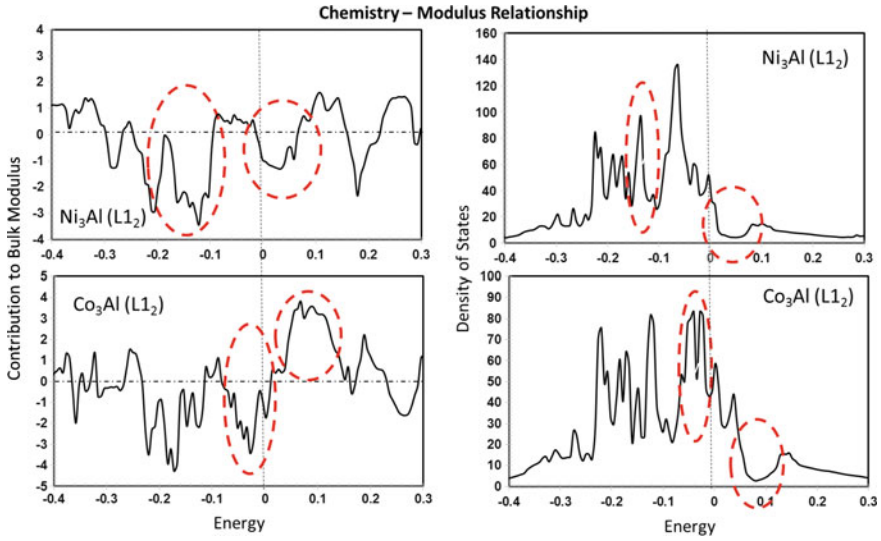


Fig. 12.9 Identification of signatures of the DOS spectra of $L1_2$ structures for Ni_3Al and Co_3Al . The interpretation of this figure is the same as in Fig. 12.8. In this case, we identify the circled doublet peak in both, and therefore identify it as a signature of the $L1_2$ structure. This difference in features define the difference in stability for Co_3Al and Ni_3Al , which have similar electronic structure but very different stabilities in $L1_2$ structure




Fig. 12.10 We identified two signatures of the DOS for Ni_3Al in the $L1_2$ structure which determine the bulk modulus. The first signature is the circled doublet peak which is due to the $L1_2$ structure, as described in Fig. 12.9. Further, the signature corresponding with the bonding-anti-bonding transition is due to the chemistry. The correlation with modulus, electronic signatures, crystal structure and chemistry provide a pathway for inverse design for chemical substitutions

a signature, this is not surprising as that is a factor in determining single element transition metal bulk modulus, as discussed in the introduction. However, the signature of the peaks correlating to $L1_2$ structure would not be identified otherwise. As discussed, engineering the intensity of these peaks leads to controlling the bulk modulus of the material. When combined with our prior work in connecting structure and chemistry to the DOS, we now have a template for multi-scale “inverse design” of new alloys.

12.5 Summary

In this paper, we developed a hybrid informatics approach for extracting bulk modulus from the DOS spectra and identifying subtle features in the DOS spectra which dictate differences in stability of electronically similar alloys. By connecting property and DOS spectra, we can now develop “virtual” DOS which correspond to target property, thereby representing an inverse design approach where we start with the property and calculate the material, as opposed to the traditional approach. The approach developed here first extracted parameters based on the comparison of the DOS spectra with the signals corresponding to material characteristics. The modulus was then modeled based on the quantitative relationship between the spectral weightings and the property, thus developing electronic structure-crystal structure-property relationships. The natural extension of this work is predicting the influence of alloying additions on DOS and the use of our approach as a means for searching for stability of multicomponent systems without doing large numbers of first principles calculations, as well as rapidly exploring the role of rare earth additions compared to non-rare earth additions in terms of electronic structure.

Acknowledgments We acknowledge support from NSF grant no. DMR-13-07811 and Air Force Office of Scientific Research grant no. FA9550-12-1-0456. KR acknowledges support from the Erich Bloch Endowed Chair-University at Buffalo: The State University of New York.

References

1. C.C. Fischer et al., *Nat. Mater.* **5**, 641 (2006)
2. S.V. Dudiy, A. Zunger, *Phys. Rev. Lett.* **97**, 046401 (2006)
3. G.H. Jóhannesson et al., *Phys. Rev. Lett.* **88**, 255506 (2002)
4. C.E. Mohn, W. Kob, *Comput. Mater. Sci.* **45**, 111 (2009)
5. A. Jain, S.P. Ong, G. Hautier, W. Chen, W.D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, K.A. Persson, *Appl. Phys. Lett. Mater.* **1**(1), 011002 (2013)
6. A. Zakutayev et al., *J. Am. Cer. Soc.* **135**, 10048–11054 (2013)
7. L. Yu et al., *Adv. Energy Mater.* **3**, 43–48 (2013)
8. S.R. Broderick, K. Rajan, *Europhys. Lett. (EPL)* **95**, 57005 (2011)
9. S.R. Broderick, H. Aourag, K. Rajan, *Stat. Anal. Data Min.* **1**, 353 (2009)
10. S.R. Broderick, H. Aourag, K. Rajan, *J. Am. Ceram. Soc.* **94**, 2974 (2011)

11. M. Finnis, *Interat. Forces Condens. Matter* (Oxford University Press, Oxford, 2003)
12. N.W. Ashcroft, N.D. Mermin, *Solid State Physics* (Brooks/Cole, Boston, 1976)
13. J.R. Alvarez, P. Rez, *Acta Mater.* **49**, 795 (2001)
14. W. Zhou, H. Wu, T. Yildirim, *Phys. Rev. B (Condens. Matter Mater. Phys.)* **76**, 184113 (2007)
15. S.F. Matar, M.A. Subramanian, *Mater. Lett.* **58**, 746 (2004)
16. L. Cheng-Bin et al., *Chin. Phys.* **14**, 2287 (2005)
17. J. Hauglund et al., *Phys. Rev. B* **48**, 11685 (1993)
18. D.J. Singh, L. Nordstrom, *Planewaves, Pseudopotentials, and the LAPW Method* (Springer, Berlin, 2006)
19. P. Hohenberg, W. Kohn, *Phys. Rev.* **136**, B864 (1964)
20. P. Blaha, K. Schwarz, (Vienna University of Technology, Austria, 2002)
21. J.P. Perdew et al., *Phys. Rev. B* **46**, 6671 (1992)
22. C. Suh et al., *Data Sci. J.* **1**, 19 (2002)
23. A. Daffertshofer et al., *Clin. Biomech.* **19**, 415 (2004)
24. L. Ericksson et al., *Multi- and Megavariate Data Analysis: Principles, Applications* (Umetrics Ab, Umea, 2001)
25. H. Berthiaux et al., *Chem. Eng. Process.* **45**, 397 (2006)
26. S. Wold, M. Sjostrom, L. Eriksson, *Chemom. Intell. Lab. Syst.* **58**, 109 (2001)
27. D.V. Nguyen, D.M. Roche (2002), p. 39
28. R. Rosipal, N. Kramer, in *Subspace, Latent Structure and Feature Selection Techniques*, ed. by C. Saunders, et al. (Springer, Berlin/Heidelberg, 2006), p. 34
29. P. Geladi, B.R. Kowalski, *Anal. Chimica Acta* **185**, 1 (1986)
30. S. de Jong, *Chemom. Intell. Lab. Syst.* **18**, 251 (1993)

Part III
Combinatorial Materials Science
with High-throughput
Measurements and Analysis

Chapter 13

Combinatorial Materials Science, and a Perspective on Challenges in Data Acquisition, Analysis and Presentation

Robert C. Pullar

Abstract Combinatorial Materials Science is the rapid synthesis and analysis of large numbers of compositions in parallel, created through many combinations of a relatively small number of starting materials. It is, therefore, essential that for a truly combinatorial approach both synthesis and measurement must be high-throughput, to handle the large number of samples required. Since the first serious attempts at combinatorial searches in Materials Science in the mid 1990s, the technique is still very much in its infancy, falling way behind the progress made in biomedical and organic combinatorial chemistry, despite attracting increasing interest from industry. The most investigated materials by combinatorial methods are catalysts and phosphors, and most work has been on libraries in deposited thin film form. This chapter will give a broad overview of the different synthetic strategies used, with a particular look at the difficulties of producing thick film or bulk ceramic/metal-oxide libraries. A vast number of characteristics can be quantified in combinatorial materials libraries, from compositional, crystal phase, structural and microstructural information, to functional properties including catalytic/photocatalytic, optical/luminescent, electrical/dielectric, piezoelectric/ferroelectric, magnetic, oxygen-conducting, water-splitting, mechanical, thermal/thermoelectric, magnetoelectric/optoelectric/magneto-optic/multiferroic, bioactive/biocompatible, etc. This chapter will cover the range of high-throughput measurements open in combinatorial Materials Science, and especially the challenges in presenting and displaying the large and complex amount of data obtained for functional materials libraries. To this end, the use of glyphs is looked at, glyphs being data points that also contain extra levels of information/data in graphic form.

R.C. Pullar (✉)

Departamento de Engenharia de Materiais e Cerâmica/CICECO - Aveiro Institute of Materials, Universidade de Aveiro, Campus Universitário de Santiago, 3810-193 Aveiro, Portugal
e-mail: rpullar@ua.pt

13.1 Combinatorial Materials Science—20 Years of Progress?

Combinatorial materials science is the rapid synthesis and analysis of large numbers of compositions in parallel, created through many combinations of a small number of starting materials. It is, therefore, essential for a truly combinatorial approach that both synthesis and measurement must be high-throughput, to handle the large number of samples required. Combinatorial searching was initiated in the 1960s for the solid-phase synthesis of peptides by Merrifield [1] (Fig. 13.1), who later won a Nobel prize for this, but it took until the 1990s for industry to adopt this technique, which is now deemed essential in the pharmaceutical industry, where both sample preparation and analysis are carried out by robots.

The first compositional gradients studied were those naturally occurring in co-deposited thin films to construct alloy phase diagrams [2]. In 1970, Hanak proposed his ‘multiple sample concept’ in the *Journal of Materials Science* as a way around the traditional, slow, manual, laboratory preparation procedures used to make samples for testing [3]. Robotic search methods for cuprate superconductor ceramics was explored by the GEC Hirst Laboratories (Wembley UK) in the early 1990s [4], and a series of combinatorial searches in Materials Science were carried out in 1995 by Xiang, Schultz et al. [5], on a 128 sample combinatorial library of luminescent materials obtained by co-deposition of elements on a silicon substrate. This milestone paper was published in *Science*, and had a picture of the combinatorial library on the cover (Fig. 13.2).

Since then, the interest in combinatorial materials science searches has increased greatly over the last 20 years, to the extent that there are now regular conferences on this specific field [6, 7]. The existence of a whole journal dedicated to *Combinatorial Chemistry* since 1999 (now renamed *ACS Combinatorial Science*), and special issues of *Measurement Science and Technology* on combinatorial materials science (e.g. 2005, vol. 16, issue 1), show how the field is growing (Fig. 13.3). There have also been several high-profile review papers on combinatorial methods [8] and high-throughput analysis [9, 10]. Using data from Scopus, it can be seen that the number of




Fig. 13.1 The first true combinatorial synthesis, created by R.B. Merrifield, for the high-throughput parallel synthesis of peptides [1]



Fig. 13.2 The cover of the *Science* issue in 1995 containing the paper by Xiang and Schultz, with a photograph of a section of the luminescent thin film combinatorial library [5]

combinatorial and high-throughput papers is steadily increasing each year, but that clearly progress in Materials Science is lagging behind severely (Fig. 13.4). Indeed, if all the publications on combinatorial and high-throughput are broken down into their Scopus subject areas (Fig. 13.4), it can be seen that a quarter are in Biochemistry, and another quarter in Medicine or Engineering, indicating the dominance of the biomedical sector in this field. Materials Science accounts for only 4% of all combinatorial and high-throughput articles over this period, and the situation is not rapidly improving, as looking at 2014 only, Materials Science is still in last place of all these categories, with only 6%.

Breaking down research into general combinatorial and high-throughput by country, it can be seen that the USA dominates hugely producing over 1/3 of all papers, but that a rapidly industrialising China is now in second place, ahead of the UK, Germany, Japan and France (Fig. 13.5). Looking at the institutions that have produced the most articles, all are in North America except for the University of Cambridge (UK) in 5th place, and the University of Tokyo, in 14th Place (Fig. 13.5).




Fig. 13.3 A selection of journals and special issues devoted to combinatorial synthesis

However, if this data is analysed only for articles related to Materials Science (Fig. 13.6) it paints a different picture, with Japan now in a clear second place to the USA, which dominates even more, and four Japanese institutions are in the top ten, including first (Tokyo Institute of Technology), third (National Institute for Materials Science Tsukuba) and fourth (Japan Science and Technology Agency) places. In second position is the National Institute of Standards and Technology (NIST, USA), which has initiated a very large research programme into Combinatorial Materials Science.

Currently, industry is already heavily involved in the development of synthesis techniques, and the development and automation of measurements, suitable for combinatorial searches—indeed, it should be noted in Fig. 13.5 that the biomedical company Pfizer is in 13th place. Major companies investing in such research also include Hitachi, General Electric, Kodak, Ciba, Hoechst, Bosch, Bayer, BF Goodrich, Siemens, Dow, Englehard, Dupont, L’Oreal, ICI, PPG, Unilever, Procter & Gamble, Intel, Heraeus, Alcoa, Celanese, Rhodia, Shell, Exxon-Mobil, Volkswagen, Honeywell, Degussa, Azko Nobel, Lucent Technologies (Bell Labs) and BASF [11], and Kurt J Lesker Co. have developed commercial combinatorial PLD (Pulsed Laser Deposition) systems. However, combinatorial and high-throughput methods for materials science are in still their infancy. The main activity is in the USA and Japan, with the leading countries in the EU being Germany and the UK, reflecting the output of academic papers shown above. A search on Scopus revealed a total of 17 800 patents on combinatorial and high throughput synthesis, but only 1000 of these were related to materials science.

Yearly Combinatorial & High Throughput Publications from Scopus



Combinatorial Publications per Subject Area




Fig. 13.4 Yearly publications on combinatorial and high-throughput topics since 1995, and below, publications on combinatorial or high-throughput per Subject Area over this period, data from Scopus. Searches were for (TITLE-ABS-KEY(combinatorial OR “high throughput”)), (TITLE-ABS-KEY(“high throughput”)) and (TITLE-ABS-KEY(combinatorial OR “high throughput”) AND TITLE-ABS-KEY(“materials science” OR ceramic OR composite OR film OR sol-gel), respectively, with results classified under (SUBJAREA, “MATH”) excluded




Fig. 13.5 The countries and institutions that have published the most articles on general combinatorial and high-throughput topics between 1995–2014 (data from Scopus)




Fig. 13.6 The countries and institutions that have published the most articles on combinatorial and high-throughput Materials Science between 1995–2014 (data from Scopus)

13.2 Combinatorial Materials Synthesis

Much current high-throughput combinatorial research is focused on biotechnology and biological systems [12]. However, here I shall only look at the state of the art in Materials Science of metals, oxides and ceramics. To date, most such combinatorial high-throughput methods use thin films, deposited on the nanoscale by various methods. If we break down the combinatorial and high throughput Materials Science papers by type of material investigated, we can see that the vast majority are on thin films and/or nanoparticles and nanosynthesis, usually by deposition




Fig. 13.7 Yearly publications on combinatorial and high-throughput Materials Science since 1995, for the named topics (data from Scopus)

(Fig. 13.7), and that very few are on bulk or thick film ceramics despite the fact that such materials can have completely different properties and applications than their thin film/nanoscale analogues. In fact, it can be seen that research into such materials peaked in 2008, while the other categories have continued to increase, demonstrating both the difficulty and the need in developing combinatorial techniques for bulk, sintered ceramics.

Several methods of high-throughput thin-film synthesis techniques have been developed for exploring new compositions, as well as for optimising process parameters of materials. Methods to prepare different types of combinatorial thin-film libraries include discrete sequentially masked depositions [13] or composition spread co-deposition [14] by molecular beam epitaxy (MBE), pulsed laser deposition (PLD), liquid source misted chemical deposition (LSMCD) [15], composition-gradient molecular layer epitaxy [16], ion beam sputtering deposition (IBSD) and chemical vapour deposition (CVD) [17]. All of these techniques tend to result in libraries with uneven thickness and stoichiometry, but they allow for easy mapping of the structural changes or functional properties.

MBE, or combinatorial laser MBE (CLMBE), uses a mask pattern which can be designed on a computer to use a masked carousel, evaporating several targets with a laser to deposit epitaxial layers on the substrate, with variations in relative stoichiometry across the library (Fig. 13.8). MBE can be more accurate than PLD, as it involves monolayer epitaxial growth, but this naturally makes it very slow, which is not ideal for high-throughput synthesis. Other deposition methods such as PLD,

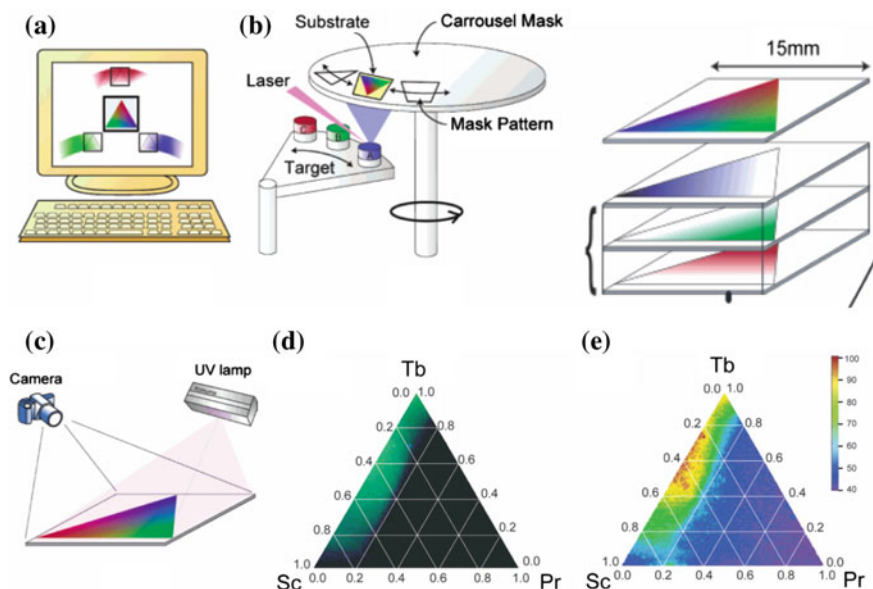


Fig. 13.8 Combinatorial laser molecular beam epitaxy (CLMBE) process: **a** Mask designed on computer; **b** compositional spread created on substrate one layer at a time using masks; **c** rapid high-throughput analysis possible, such as luminescence under UV light; **d** photograph of $\text{Tb}_{1-x-y}\text{Sc}_x\text{Pr}_y\text{Ca}_4\text{O}(\text{BO}_3)_3$ library under 254 nm UV excitation; **e** emission intensity map same ternary thin film library of $\text{Tb}_{1-x-y}\text{Sc}_x\text{Pr}_y\text{Ca}_4\text{O}(\text{BO}_3)_3$ [18]




Fig. 13.9 The pulsed laser deposition (PLD) combinatorial film deposition process

CVD and IBSD are quicker, and use similar masking effects or shutters (Fig. 13.9), but result in more variable compositions.

Combinatorial methods which can be applied to bulk or thick film ceramics usually use either high-throughput synthesis of powders [19] (Fig. 13.10) or ink-jet printing methods [20] (Fig. 13.11). Much of the work on combinatorial powder synthesis involves only the production of a library of powders with a compositional

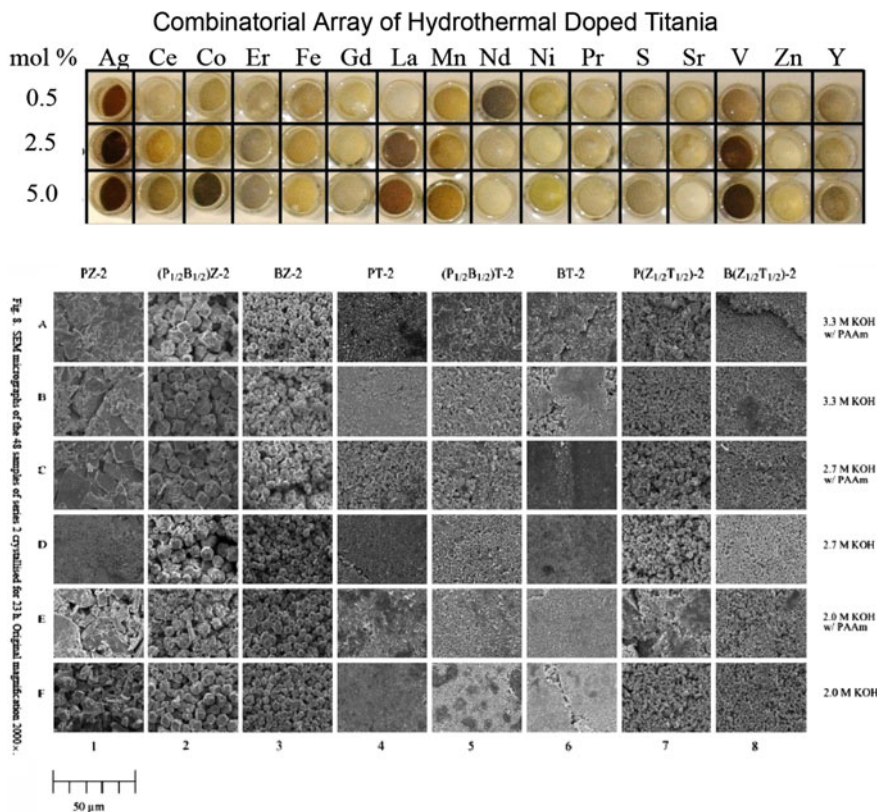


Fig. 13.10 Combinatorial powder library of doped TiO₂ made continuous hydrothermal synthesis [24], and SEM images of a 48 sample library of perovskite powders also produced via a hydrothermal batch process [19]

spread, with no subsequent high-throughput processing, e.g. autopipetting of sols to produce small amounts of combinatorial powders [21], powder metallurgy using acoustic vibration valves to dispense powders [22], solution combustion synthesis of combinatorial libraries of photocatalytic perovskites in microwells [23], or continuous hydrothermal synthesis of combinatorial libraries of oxide powders from nanoparticle suspensions [24, 25] (Fig. 13.10). Some workers also incorporate a high-throughput processing method with synthesis, e.g. a combinatorial robot system for measuring, mixing and moulding liquid samples by automatic micropipette to produce a library for ceramics on a pallet [26], or robotic dosing and planetary ball milling of 40 different samples with a parallel pressing of 5 samples at a time [27].

The ink-jet printing process creates a thick film library already laid out on a substrate ready for processing [28, 29], and this method has been particularly successful in the discovery of new phosphors [30, 31], such as the 121 sample library of the

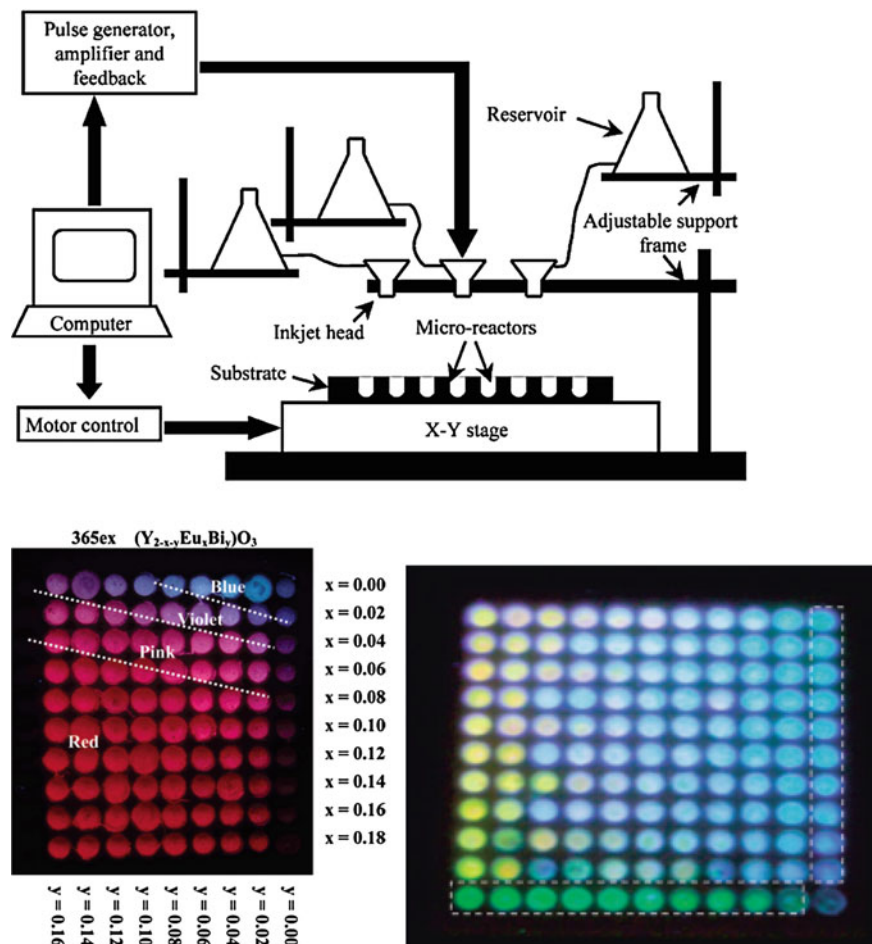


Fig. 13.11 Schematic diagram of the combinatorial inkjet printing process, and libraries thus created of red $Y_{2-x-y}Eu_xBi_yO_3$ and blue $KrSr_{1-x-y}PO_4:Tb^{3+}_xEu^{2+y}$ phosphors under UV light [31, 32]

$KrSr_{1-x-y}PO_4:Tb^{3+}_xEu^{2+y}$ UV phosphor system or the red Y_2O_3 -based phosphor libraries created by Chan et al. [31, 32] (Fig. 13.11).

The author, R. C. Pullar, was part of the *Functional Oxides Discovery using Combinatorial Methods* (FOX) project, using the London University Search Instrument (LUSI) robot to make sintered combinatorial libraries of ceramic compositions. LUSI automatically created sintered bulk ceramic libraries by ink-jet printing multicomponent mixtures on substrates, then robotically loading the libraries into a flatbed 4-zone furnace for firing, with up to 100 °C difference between each zone, unloaded the samples, and could also place them on a test bed for measurement [33]. The ceramics under investigation in the FOX project were dielectrics, ferroelectrics [34–36] and ionic conductors [37], and libraries were made and characterised (Fig. 13.12).







Fig. 13.12 The LUSI robot [33], a printed and sintered Ba_{1-x}Sr_xTiO₃ library on a single 50 mm long substrate (each dot 1–2 mm wide), SEM images of the sintered library, EDS measurements showing the variation in composition cross the library, and dielectric measurements showing the functional gradient in Curie temperature across the library [34, 35]

Other workers have also made and characterised libraries of thin film dielectric ceramics, such as ternary oxide $ZrO_x-SnO_y-TiO_z$ [14], microwave dielectrics [38], high ϵ_r (50–80) $HfO_2-TiO_2-Y_2O_3$ dielectrics [39], sol-gel piezoelectrics [40], and a 64 sample LSMCD ferroelectric $Bi_{3.75}La_xCe_{0.25-x}Ti_3O_{12}$ library [15]. Bulk ceramic piezoelectrics have also been studied via combinatorial methods [41]. Many other ceramics have been investigated such as high-throughput analysis of semi-conductors [42], combinatorial physical vapour deposition of gold nanoparticles [43], sol-gel metal oxide nanoparticles [26], SOFC ceramics [24], catalysts [44], pigments [45], gas sensor materials [46], electrochemical electrodes [47] and hydrogen energy storage materials [48]. Maybe the best known new material discovered via combinatorial methods is the Co-doped TiO_2 dilute magnetic semiconductor [49], discovered by chance in a combinatorial search of 162 thin film photocatalyst candidate materials [16], and leading to an explosion of interest in such materials. In his review [10], Zhao lists 23 new materials successfully discovered by combinatorial high-throughput searches, and examples include $Zr_{0.2}Sn_{0.2}Ti_{0.6}O_2$ dielectrics from libraries of 30 multicomponent systems [50], high ϵ_r microwave dielectrics [14, 39, 51], cobalt oxide magnetoresistance materials [52], hydrogen storage candidates [48], novel photocatalysts [23], and improved catalysts from a library of thousands of samples [44]. In their analysis of new generation capacitance materials for random access memory devices, replacing amorphous silica with optimised materials based on $ZrO_2-SnO_2-TiO_2$, Koinuma and Takeuchi [8] suggest that 900 one-by-one sputtering preparations would have been required to fully explore the combinations of the ternary Zr-Sn-Ti oxide system prepared in the compositional spread by van Dover et al. [50].

Most combinatorial materials searches involve thin film techniques, and for some applications where the end product will be exclusively in thin film form, this does make sense. However, many materials are also required in bulk form, and the bulk properties can be quite different to those of thin films, where surface diffusion, strain effects from substrate-lattice mismatch, and surface and skin electrical effects dominate. For example, ferroelectric functions are highly dependent upon strain effects in thin films. Also, most thin films are epitaxial or single crystal, and hence have no grain boundaries, which can have a large effect on electrical, magnetic, dielectric, mechanical and transport properties. From the point of view of constructing large materials properties data bases for data mining and prediction of novel compositions, it could be argued that bulk properties are much more relevant than those of thin films. Furthermore, for many applications bulk or thick film ceramics are required, e.g. multilayer chip capacitors, low temperature co-fired ceramics (LTCC), structural and engineering ceramics, refractories, clays, glazes and household ceramics, SOFC and ionic conductors, electromagnetic and radar absorbing materials (RAM), catalyst supports, substrates, etc.

As discussed above all current bulk ceramic combinatorial projects either just make a combinatorial library of powders through a high-throughput synthesis process (e.g. hydrothermal) [24], or they use a solution based process to deposit or print a library on a substrate (e.g. ink-jet printing) [29]. In the first case, there is no high-throughput processing, and each sample in the library must be individually

prepared (e.g. pressed in a die) from the powder, usually by hand. In the second case, the solution chemistry, stability during a printing run, drying in a regular shape, and reaction with, or lack of adherence to, the substrate, becomes a serious issue, especially in complex multi-component systems. The issue of reactions with, or lack of adherence to, substrates is also an issue for all ink-jet based combinatorial processes, as they need a substrate that can be both printed on and heated during sintering [28].

It was Dr. Pullar's experiences with ink-jet printing that led him to consider an alternative solution, which gave the benefits of a bulk or thick film ceramic library, but in a much simplified, novel adapted tape casting process. Using a minimum of solvent, the combinatorial components could be mixed with a commercial mixer tip, designed for mixing adhesives, polymers and dental cements. Unlike a solution based process, much less volume is lost on drying, leading to a denser green body that should produce dense ceramics, and no segregation or precipitation effects should occur. The libraries can be made either on a substrate, or a release tape which can be removed before firing, avoiding substrate problems if necessary. As the sintering step is often a rate limiting step in combinatorial ceramics, a multiple zone furnace was used to simultaneously fire five libraries at different temperatures. This technique has been used to create sintered libraries of magnetoelectric $\text{SrFe}_{12}\text{O}_{19}/\text{BaTiO}_3$ composite ceramics, in compositional steps of 10%, in which the two phases, one magnetic and the other dielectric, did not react, maintaining their respective characters (Fig. 13.13) [53].

13.3 High-Throughput Measurement and Analysis

The importance of combinatorial high-throughput materials science has been clearly shown above, and although still in its infancy, the fact that so many industries are investing in developing such techniques demonstrates their belief in its future significance. Once established, its impact on Materials Science will be enormous, like on that of the pharmaceutical industry, and increasingly biomedicine and biochemistry. However, to be successful, combinatorial materials synthesis also requires high-throughput measurement. The diverse spectrum of functionalities in materials represents a significant challenge in high-throughput characterisation, and often involves development of novel measurement methods [9]. Zhao's review paper [10] is good overview of the techniques available for combinatorial high-throughput analysis, although it does concentrate on the low-micro and nanoscale, which is by no means all that is of interest. It must be understood that the aim of characterisation in combinatorial science is a broad brush mapping or analysis of the sample to show trends and unexpected or complimentary properties, not a precise measurement—that can come later on materials of interest. The properties of combinatorial libraries can be measured as a function of composition to give a functional gradient, which can also vary with processing conditions between identical libraries processed differently. Properties that can be investigated include:

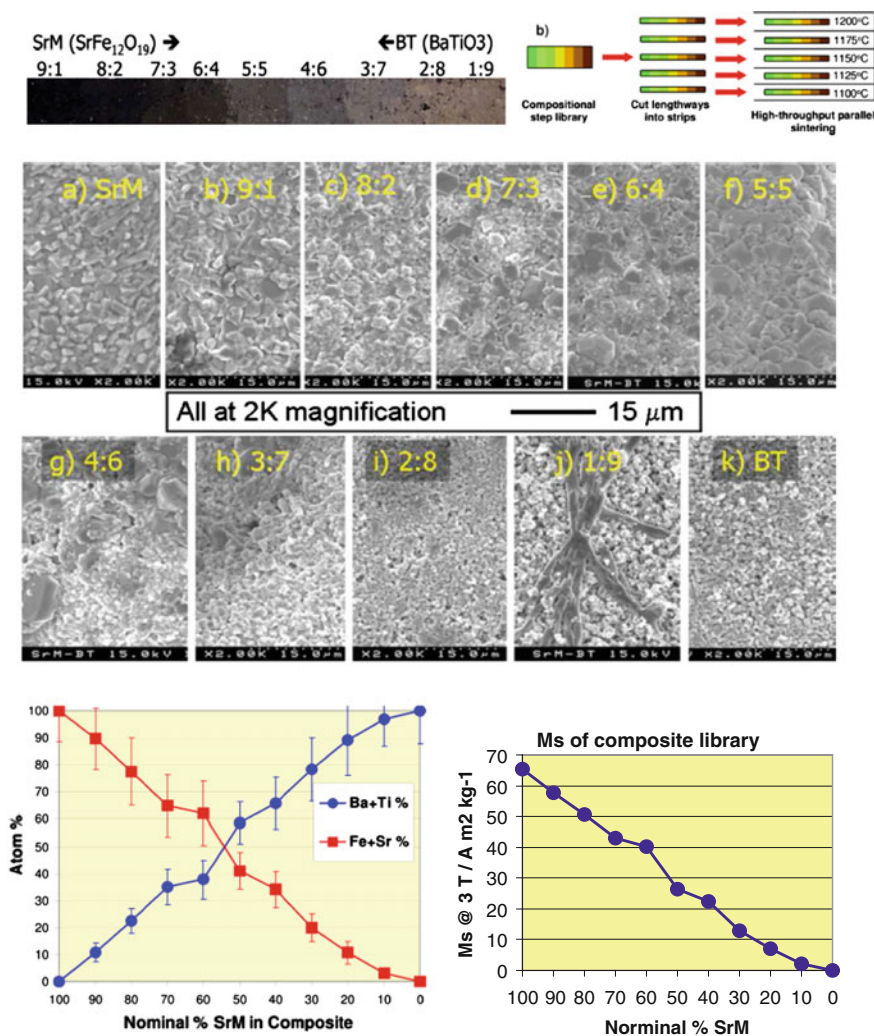


Fig. 13.13 A photograph of sintered, bulk $\text{SrFe}_{12}\text{O}_{19}/\text{BaTiO}_3$ library (with compositional ratios along the library from 9:1 to 1:9 for SrM:BT); a diagram of the parallel high-throughput firing process; SEM images of the microstructure of the library; EDS spectra of compositional variation along the library; magnetic measurements showing the functional variation in magnetisation (M_s) along the library [53]

- **Composition** and phase purity/solid solutions/lattice parameters/crystal structure
- **Microstructure**/density/porosity/grain boundaries and segregation
- **Mechanical properties**: hardness, elastic modulus, stress/strain, etc.
- **Electrical properties**: Conductivity, superconductivity, ionic conduction, oxygen vacancies

- **Dielectric properties:** permittivity, ferroelectricity, piezoelectricity, capacitance, Curie points
- **Magnetic properties:** domains, magnetisation, hysteresis loops, Curie points,
- **Optical properties:** electro-optics, magneto-optics, luminescence/fluorescence
- **Thermal properties:** Thermal conductivity, thermoelectrics, thermal creep, thermal expansion
- **Multiferroics:** multiferroic and magnetoelectric coupling, responses to direct and indirect stimuli
- **Chemical reactions:** catalysis, selectivity, redox reactions, fuel cells, water splitting/H₂ production
- **Band gaps:** photocatalysis, solar energy, semiconductivity, smart materials, etc.
- **Biomaterials:** antibacterials, human compatibility, biological markers, etc.

Many of these parameters can also vary with changes in measurement temperature, pH, wavelength of light, or applied electrical or magnetic fields, adding yet another layer of complexity, and creating yet more data.

The most basic tools for characterising or mapping composition and phases present are XRD (x-ray diffraction) and EDS (energy dispersive spectroscopy, also known as EDX—energy dispersive x-ray analysis). Use of Real Time Multiple Strip (RTMS) XRD detectors, such as the PANalytical PIXcel range or Shimadzu One-Sight, has become essential for the rapid high-throughput structural characterisation of combinatorial libraries, by greatly speeding up the measurement time with little or no loss of resolution, meaning that scans that would normally take hours can be carried out instead in a few minutes (Fig. 13.14). As well as identifying the phases present, XRD also gives structural information, lattice parameters, etc., and the coefficient of thermal expansion (CTE) can be evaluated from changes in lattice parameters with temperature [10].

EDS is another very rapid technique, in which measurements take a few minutes, which can be used to identify elements present in a scanning electron microscope (SEM) image (see Fig. 13.13), and can also map the distribution of those elements. While doing EDS, SEM images can also be rapidly taken, to study microstructure, porosity, sintering, liquid phases, phase/grain boundaries, etc., and Electron Back-Scatter Diffraction (EBSD) is also a useful tool for identifying crystal structures in reference to a library of known structures. Many points can be measured in a minute on polished samples, and it can also measure changes in orientation in anisotropic samples.

Scanning probe techniques such as Atomic Force Microscopy (AFM), Piezoresponse Force Microscopy (PFM) and Magnetic Force Microscopy (MFM) are often collectively called Scanning Probe Microscopy (SPM, Fig. 13.15), and are ideal tools for high-throughput mapping of the functional gradients of combinatorial libraries, and can map a library in minutes. PFM can show piezoelectric grains and domains, and with measurements before and after poling, hysteresis loops and piezo coefficients (d_{33}) can be measured. MFM can show magnetic domain structure, and with an external magnetic field it can show hysteresis loops and coercivity (H_c) values for each point measured, but it cannot be used to measure magnetic moments or

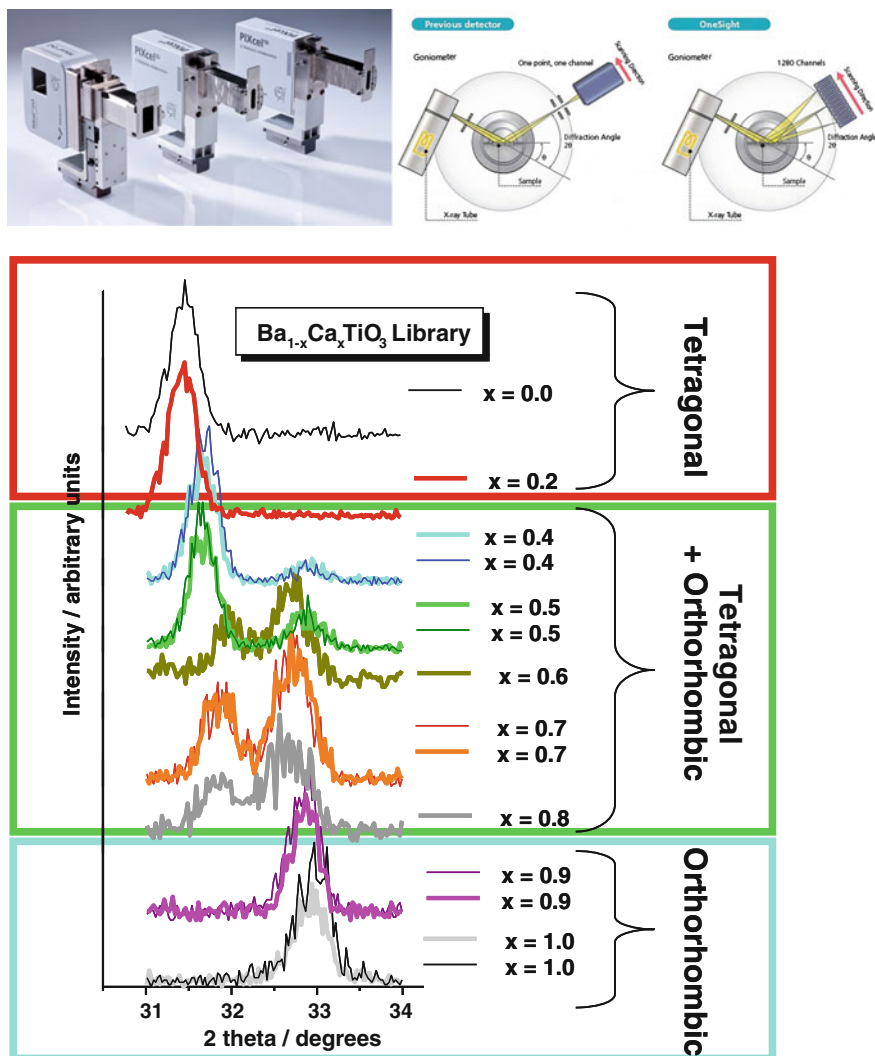


Fig. 13.14 RTMS XRD detectors, which work by simultaneously measuring over a range of angles, greatly speed up XRD measurements of combinatorial libraries. The measurements below, taken in only 2 min each by the author over a range of 20–70°, clearly show the change in structure from tetragonal BaTiO_3 to orthorhombic CaTiO_3 across a bulk ceramic $\text{Ba}_{1-x}\text{Ca}_x\text{TiO}_3$ library [36]

saturation magnetisation (M_s) values, as it is measuring the remnant magnetisation. Nanoindentation can be carried out by the AFM tip to give mechanical properties such as hardness and elastic modulus. A related technique is Atomic Force Acoustic/ultrasonic Microscopy (AFAM), which vibrates the AFM cantilever in contact mode, the change in resonant frequencies giving information about stiffness

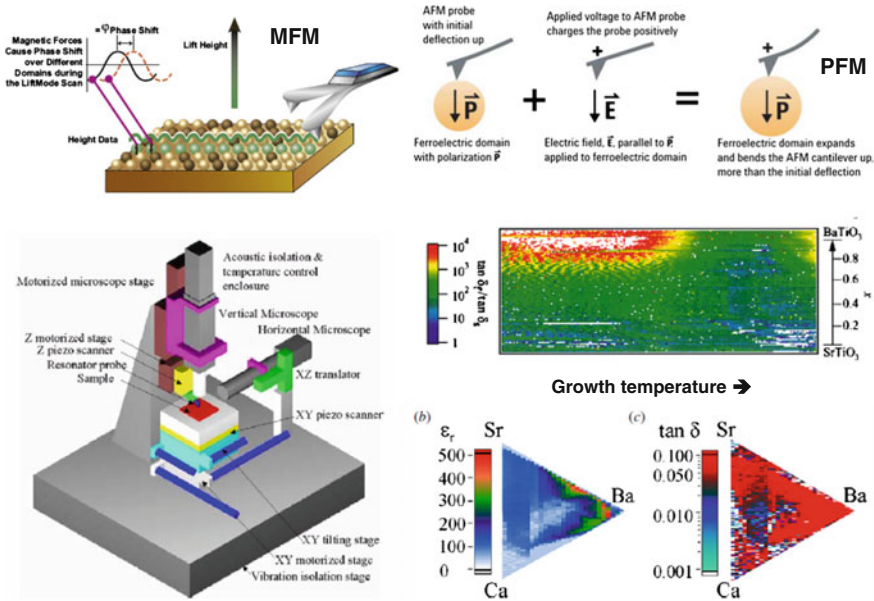


Fig. 13.15 Scanning Probe Microscopy (SPM) techniques useful for high-throughput combinatorial libraries: *Top*, Magnetic Force Microscopy (MFM) can measure magnetic properties and map magnetic domains, and Piezoresponse Force Microscopy (PFM) can measure piezoelectric hysteresis loops at different spots on a library, and map piezoelectric domains. *Bottom*, the Evanescent Microwave Probe (EMP), or Scanning Evanescent Microwave Microscope (SEMM), can map dielectric properties over a library such as permittivity (ϵ_r) or dielectric loss ($\tan \delta$), as shown by the $\tan \delta$ maps of a $\text{Ba}_{1-x}\text{Sr}_x\text{TiO}_3$ thin film library with varying growth temperature, and the ϵ_r and $\tan \delta$ maps of a ternary $\text{Ba}_{1-x-y}\text{Sr}_x\text{Ca}_y\text{TiO}_3$ thin film library [54]

and local elastic constants. The Evanescent Microwave Probe (EMP), or Scanning Evanescent Microwave Microscope (SEMM), is a kind of SPM that measures the change in dielectric properties of a metal tip embedded in a microwave resonator just above, or in contact with, the surface [54]. Interaction with the sample changes the resonant frequency and dielectric loss of the resonator, and from this electrical conductivity, permittivity (ϵ_r) and quality factor (Q) of the sample can be calculated and mapped in minutes, and although accurate quantitative measurements are problematic, EMP has been used on combinatorial libraries [54] (Fig. 13.15).

Electrical conductivity measurements are very important in combinatorial searches for insulators, superconductors, semiconductors and thermoelectrics, and also in dielectrics and ferroelectrics along with permittivity and dielectric loss ($\tan \delta$, $Q \approx 1/\tan \delta$). Bulk samples and thick films can be analysed by a simple capacitance method, if top and bottom electrodes are applied, to measure all of these values quickly, and over a range of temperatures with longer runs to give Curie points (T_c) and ferroelectric/relaxor behaviour. The author has used such a method to simultaneously measure multiple points in bulk dielectric ceramic combinatorial libraries

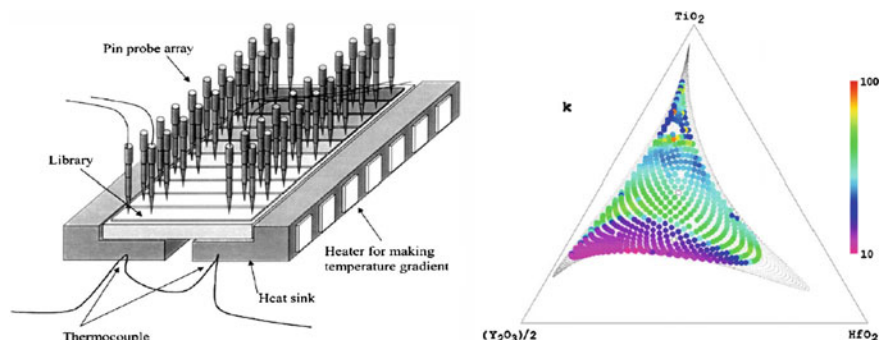


Fig. 13.16 Example of a multiple electrode array used for high-throughput electrical measurements of libraries, and the map of permittivity over a ternary Ti-Hf-Y oxide thin film library [39]




Fig. 13.17 General schematic diagram for the high-throughput optical measurement of combinatorial libraries, where light could be of various wavelengths (UV, visible, laser) and many kinds of spectrometer could be applied, and optical measurements of transmittance and band gap at UV wavelengths for a $\text{Zn}_{1-x}\text{Mg}_x\text{O}$ thin film library

between 150–450 K, clearly showing T_c and phase transitions [35, 36]. An 8 mm², 64 multielectrode array for high-throughput impedance spectroscopy, suitable for dielectrics, semiconductors and electrocatalysts, has also been developed [55] and used to test for gas sensing abilities across a library, and ϵ_r and leakage current have also been mapped from capacitance-voltage (C-V) and current-voltage (I-V) measurements of ferroelectric PLD thin film libraries [39] (Fig. 13.16).

Optical techniques such as FTIR, Raman and UV-Vis spectroscopy, colourimetry, cathodoluminescence (CL) and photoluminescence (PL) are clearly suitable for high-throughput analysis and mapping (Fig. 13.17), needing short measurement times and measuring only the spot in the beam at one time. The first two techniques can give information on chemical bonding, phonon modes (and dielectric loss), polarisation, 2D spectra with changes over time/environment/temperature, and fingerprints of molecules and structures. The others have been used in combinatorial searches for pigments, phosphors, diodes and display materials. CCD cameras have also been used to measure the output/absorption of combinatorial libraries [10, 56], and spatially resolved infrared imaging used as a high throughput hydrogen storage candidate screening technique [48].




Fig. 13.18 *Left* SMOKE map of Co-Fe-Ni ternary system, showing the magnetic hysteresis loop extracted from just one of the pixels/data points [10]. *Right* scanning SQUID images of $\text{La}_{1-x}\text{Ca}_x\text{MnO}_3$ taken at 7 K, showing the magnetic domains and the transition from strong to weak magnetization [57]

Magnetic techniques other than MFM of great interest, and needing more development, are Scanning Magneto-Optical Kerr Effect (SMOKE) probes and scanning SQUID microscopy (Fig. 13.18). SMOKE can measure M_s if an external magnetic field is applied, and a hysteresis loop and coercivity data can be extracted from every pixel on the combinatorial map [10], but it is difficult to measure low electrically conductive oxides. Scanning SQUIDS have also been used to map combinatorial thin films [57], and a recent development is one that can operate at room temperature, although they cannot use an external field, and therefore only measure remnant magnetisation and are not currently capable of quantitative measurements.

Thermal conductivity mapping can be measured by the change in thermoreflectance of sample, heated by a femto-second pulsed laser, which has been coated with an Al film to absorb the 770 nm Ti:sapphire laser [10]. Mass spectroscopy (MS) has been used a lot in high-throughput catalytic analysis [44, 58], and a robot system measuring with an electrode array to form 16 electrochemical cells has been used in combinatorial searches for new electrode materials [47]. Catalysis is probably the field where the most progress has been made in combinatorial Materials Science, driven by industry and the relative ease of high-throughput measurement, and such synthesis and measurement systems are quite well established now, with automatic data extraction and analysis [58] (Fig. 13.19).

IR-sensography has been developed where an IR-camera acts as an external optical detector system for sensor libraries, detecting small temperature changes due to physisorption or chemisorption [59], and high-throughput impedance screening has been used on gas-sensing materials in variable atmospheres and temps [55?]. There has been a lot of development on the high-throughput search for electrochemical and gas sensor materials, including a 16 sample SnO_2 -based gas sensor library for testing as an electronic nose, and a complete high throughput assembly consisting of a 64 sample reactor for the sensor libraries (Fig. 13.20), with IR-cameras, switching multimeter for dc-resistance measurements and impedance measurements, a test gas supply array for different test gases, and software for control of experimental flow, data recording, data evaluation, data mining and a database [46, 60]. Ambitious projects like this are where the eventual future of combinatorial materials science lies, in fully integrated, automated high-throughput synthesis and measurement systems, with artificial intelligence (AI) driven control, analysis and data mining [61].

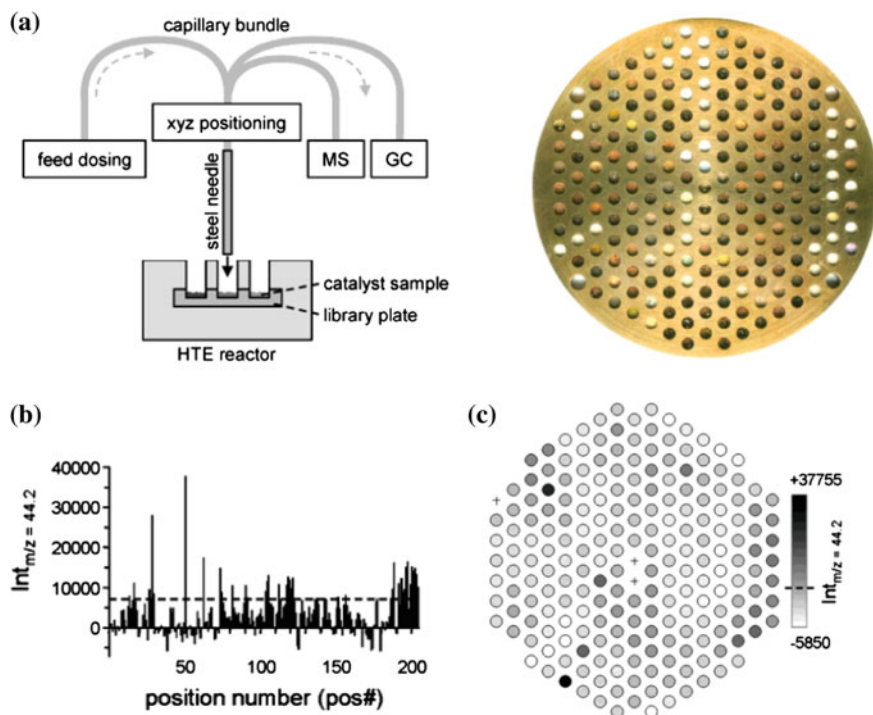


Fig. 13.19 High-throughput set up for discovery of catalysts, where the same probe is used to both deposit the samples and measure the library after processing. **a** A schematic diagram of the setup and a photograph of the 207 sample library produced, **b** measured screening results for catalysis across the library and **c** visualisation of the results on the layout of the sample [58]

13.4 Data Analysis and Presentation

If we are going to generate large amounts of data, we also need to be able to analyse it, understand it, and interpret it in a way which is comprehensible. In an ideal combinatorial system, the flow would be as shown in Fig. 13.21, with synthesis, processing and measurement of libraries all carried out by a single robot, which would then feed the data to a data base, which could be data mined and used to predict the next likely candidate systems to be investigated based on all results so far, feeding back to the next step in the synthesis process. In reality, combinatorial Materials Science is still a long way from such an automated feedback process, although much progress is being made on data mining using various forms of statistical analysis, AI and evolutionary software [62–65].

Other authors will deal with this topic, but also of great importance is how we can interpret and present such a multitude of results and data to a human audience, in such a way that it can be easily comprehended. The number of degrees of freedom in a simple binary component system are staggering—in $A_{1-x}B_x$ we have x

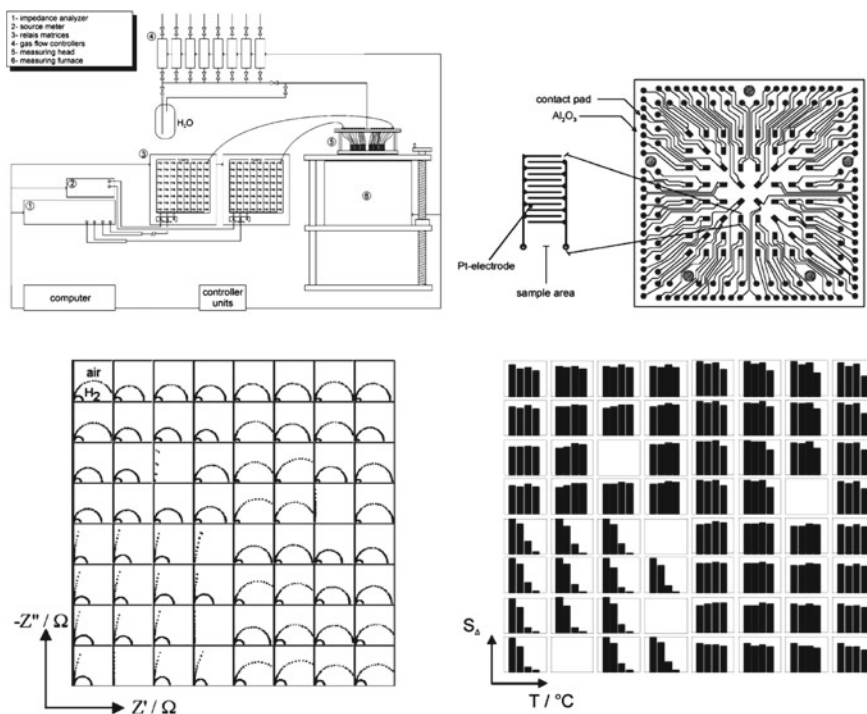


Fig. 13.20 Schematic diagram of the set up for the combinatorial synthesis and high-throughput measurement of 64 sample gas sensor libraries, the multi-electrode array used for the library, and measured results showing the Argand plots of impedance, and sensitivity to 25 ppm of H_2 (S_a , height of bar) with measuring temperature (each progressive bar, at 250, 300, 350 and 400°C) for all 54 samples in the library [60]

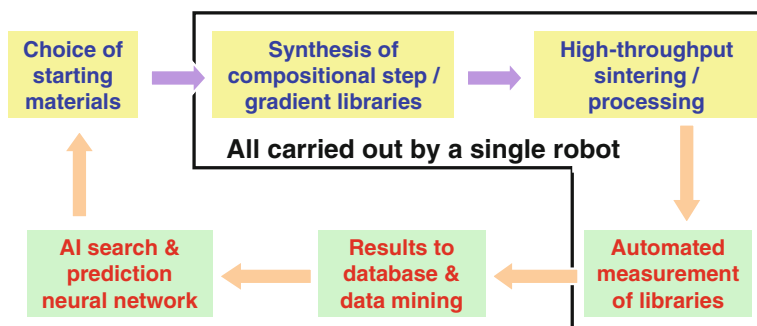


Fig. 13.21 The ideal combinatorial synthesis, processing, measurement and analysis set-up




Fig. 13.22 Triangular phase diagram map of a ternary library where change in colour/contrast depicts variation in a property

compositions, which could be processed for various temperatures, periods or pressures/atmospheres, and we want to show evolution/existence of crystalline phases, perhaps details on microstructure, and functional properties, ideally all in a single image or graphic. Clearly this is quite a challenge, and with ternary $A_{1-x-y}B_xC_y$ systems it becomes even more so. Triangular phase diagram maps can be used to show variation in composition with position and variation in a property with change in colour or contrast (Fig. 13.22). However, problems arise when we want to see the effects on various properties, or a plot of data, for each data point. A very interesting overview of the various possibilities for high-throughput analysis, and examples of unusual ways to display those results, is detailed in the review by Potyrailo et al. [66].

A solution to this is the use of glyphs—a glyph being a single data point that contains extra data in graphic form. A simple way to achieve this is to have each point on a binary or ternary plot as a different colour, size, shape or transparency, where the variation in these characteristics represents variations in functional or structural properties, for example the plot by the author showing the ternary $Ba_xSr_yCa_zTiO_3$ system in Fig. 13.23. In this plot the position represents the composition, the colour shows the position of the main XRD peak, telling us whether we have orthorhombic or tetragonal phases and solid solutions, and the size of the point shows the measured permittivity, from 147 for the smallest to 3573 for the largest. Further changes in feature such as point shape or transparency/fill could be used to indicate other properties in the same plot. Another example is shown in Fig. 13.24, from the author's paper [53], of a binary magnetoelectric composite SrM-BT combinatorial library. The black line shows the real change in composition, with the red area representing SrM and green BT, the plotted spheres show the magnetisation (the centre of each sphere being the data point), and the relative volume of the 3D spheres represents the permittivity (the larger permittivity values are too big to be represented as areas of 2D circles). It can easily be seen that the evolution in composition is not completely




Fig. 13.23 Plot by the author of the ternary Ba_xSr_yCa_zTiO₃ system, using glyphs to show composition (position), position of the main XRD peak (colour) and permittivity (size of the point)

linear along the library (or else the red/green diving line would a straight diagonal), that these non-linear variations in composition are reflected in the magnetisation values, and as magnetisation of the composite decreases permittivity increases. The second image in Fig. 13.24 shows seven degrees of data in a single plot: The quaternary Co-Te-Mn-Cr oxide catalyst system composition is depicted by the 3D pyramid, and the activity of each sample in the library towards one of three different molecules is shown by changes in colour, size and transparency of the data point [66, 67].

Another approach is to use pie chart glyphs for each point, as also shown in Fig. 13.24. One example is from the author's paper [53], showing the same magneto-electric composite library, but this time with the points representing the magnetisation on the y axis and the supposed % of SrM on the x axis, and a pie chart showing the actual measured relative proportions of the two ceramic phases in that sample of the composite library, in an easy-to-comprehend manner. It can again be seen that non-linear behaviour in the magnetisation through the library matches discontinuities in the quantity of the SrM magnetic phase. The final image in Fig. 13.24 shows a structural phase diagram produced using the weights various x-ray diffraction patterns, with the position of each point showing the composition in a ternary Fe-Pd-Ga material, and the glyph pie charts showing the relative proportions in that composition of seven possible phases found throughout the library, indicated by different coloured sections in the pie charts [68].

Data point glyphs can also be used to contain actual images, such as plots, SEM images or photographs. Two examples are shown in Fig. 13.25. The first example shows a compositional triangle in a ternary substituted ferroelectric BiFeO₃ library, but at each data point is a small image glyph of the measured ferroelectric polarisation hysteresis loop [69]. Not only can the general shape of each loop also be

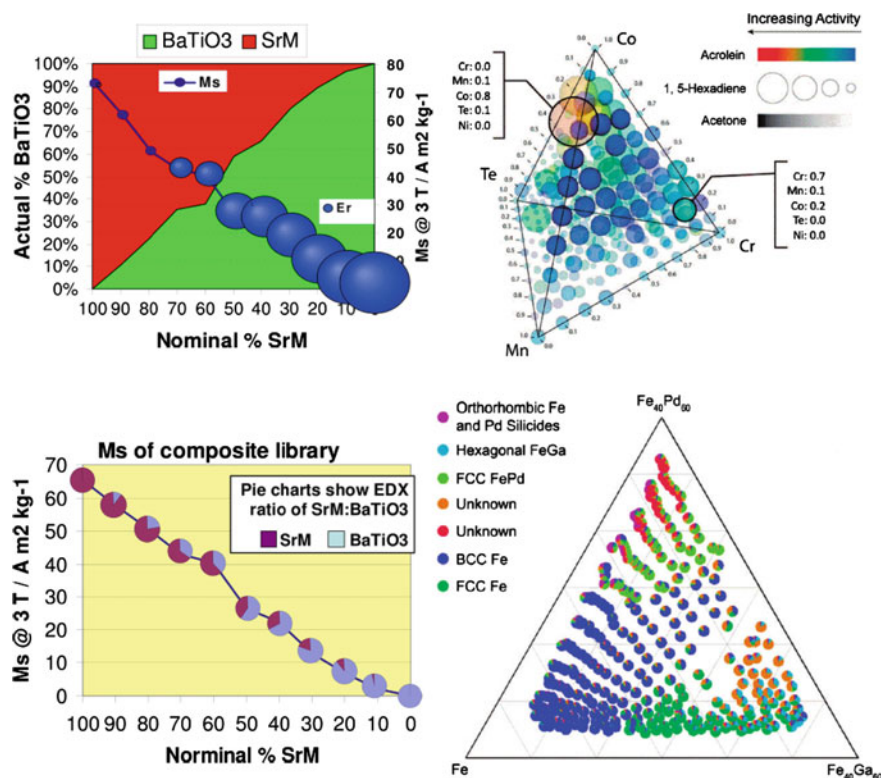


Fig. 13.24 Various uses of glyphs in displaying complex series of combinatorial data. From *top left* a chart showing relative proportions of two phases (*red* and *green*, *black line*), magnetisation (position on *left axis*) and permittivity (volume of 3D sphere) [53]; A four component catalyst library, with composition shown by position in the 3D pyramid, and catalytic activity with three molecules shown by colour, size and transparency of the data point [66, 67]; A chart showing magnetisation (on *left axis*), supposed % of SrM in the composite (*bottom axis*) and actual compositional proportion by a pie chart glyph for the data points [53]; Triangular map of composition, with pie charts and the use of 7 different colours for the data point glyphs to show the composition of phases at each point [68]

seen qualitatively, and those compositions with a good hysteresis loop easily identified, the individual loops can then be enlarged to give fully quantitative information on that loop. The second shows the compositional map of a Fe-Co-Mo alloy, with a magnetic hysteresis loop glyph at each data point, which again can be enlarged to give fully quantitative magnetic data [63]. In this case, the plot contains both in-plane (IP) and out-of-plane (OP) measurements. This last paper was part of a combinatorial search for replacements for the increasingly expensive rare earth magnets.

Clearly, many other kinds of glyph could be used as a data point. Furthermore, if plots are in online or electronic form, they can be interactive, with enlarged plots, more details, or even several different properties plots and images given when the

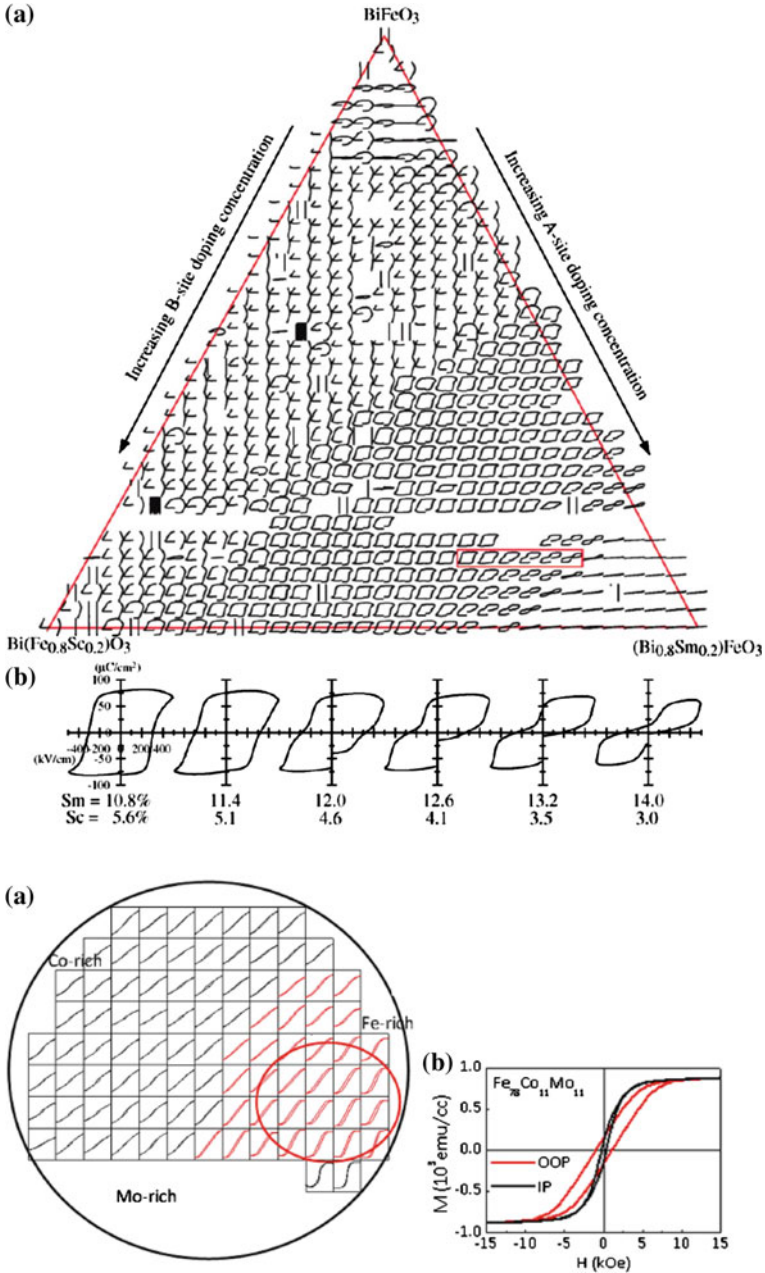


Fig. 13.25 Top, pseudoternary compositional map of ferroelectric BiFeO_3 , co-doped with (Bi, Sm) and (Fe, Sc), with the ferroelectric hysteresis loop measured at each point shown as a glyph. The six loops highlighted in the red rectangle as shown in enlarged form below, demonstrating the fully quantitative nature of these data [69]. Bottom, compositional map of a Fe-Co-Mo alloy, with magnetic hysteresis loop glyphs at each data point (a), which can be enlarged to give fully quantitative magnetic data (b) [63]

relevant glyph is clicked upon, touched or activated. This opens up a whole new area of interactive combinatorial data display and analysis, and exciting new way to handle and explore the large amount of data generated in high-throughput searches.

Acknowledgments The author would firstly like to thank the FCT (Fundação para a Ciência e a Tecnologia in Portugal), and the FCT Ciência 2008 program and grant SFRH/BPD/97115/2013 are acknowledged for funding the author during the writing and publication of this chapter. The author would also like to thank the publishers and copy write holders of all figures from previous sources used in this chapter, which have been referenced in the relevant figure caption.

References

1. R.B. Merrifield, Solid phase peptide synthesis. I. The synthesis of a tetrapeptide. *J. Am. Chem. Soc.* **85**, 2149–2153 (1963)
2. K. Kenedy, T. Stefansky, G. Davy, V.F. Zacky, E.R. Parker, Rapid mapping for determining ternary-alloy phase diagrams. *J. Appl. Phys.* **36**, 10–3808 (1965)
3. J.J. Hanak, The multiple sample concept in materials research; synthesis, compositional analysis and testing of entire multi-component systems. *J. Mater. Sci.* **5**, 964–971 (1970)
4. S.R. Hall, M.T.R. Harrison, The search for new superconductors. *Chem. Br.* **30**, 739–742 (1994)
5. X.-D. Xiang, X. Sun, G. Briceno, Y. Lou, K.-A. Wang, H. Chang, W.G. Wallace-Freedman, S.-W. Chen, P.G. Schultz, A combinatorial approach to materials discovery. *Science* **268**, 1738–1740 (1995)
6. Proceedings of the first Japan-US Workshop on Combinatorial Materials Science and Technology. *Appl. Surf. Sci.* **189**, 175–371 (2002)
7. Proceedings of the Second Japan-US Workshop on Combinatorial Materials Science and Technology. *Appl. Surf. Sci.* **223**, 1–267 (2004)
8. H. Koinuma, I. Tekeuchi, Combinatorial solid-state chemistry of inorganic materials. *Nat. Mater.* **3**, 429–438 (2004)
9. R.A. Potyrailo, I. Takeuchi, Role of high throughput characterization tools in combinatorial materials science. *Meas. Sci. Tech.* **16**, 1–4 (2005)
10. J.-C. Zhao, Combinatorial approaches as effective tools in the study of phase diagrams and composition-structure relationships. *Prog. Mater. Sci.* **51**, 557–631 (2006)
11. J. Ouellette, Combinatorial materials synthesis. *Ind. Phys.* **4**, 24–27 (1998)
12. E.W. McFarland, W.H. Weinberg, Combinatorial approaches to materials discovery. *Trends Biotechnol.* **17**, 107–115 (1999)
13. Y. Matsumoto, M. Murakami, Z. Jin, A. Ohtomo, M. Lippmaa, M. Kawasaki, H. Koinuma, Combinatorial laser molecular beam epitaxy (MBE) growth of Mg–Zn–O alloy for band gap engineering. *Jpn. J. Appl. Phys.* **38**, L603–L606 (1999)
14. R.B. van Dover, L.F. Schneemeyer, R.M. Fleming, Discovery of a useful thin film dielectric using a compositional-spread approach. *Nature* **392**, 24–27 (1998)
15. K.W. Kim, M.K. Jeon, K.S. Oh, T.S. Kim, Y.S. Kim, S.I. Woo, Combinatorial approach for ferroelectric material libraries prepared by liquid source misted chemical deposition method. *Proc. Natl. Acad. Sci. USA* **104**, 1134–9 (2007)
16. T. Fukumura, M. Ohtani, M. Kawasaki, Y. Okimoto, T. Kageyama, T. Koida, T. Hasegawa, Rapid construction of a phase diagram of doped Mott insulators with a composition-spread approach. *Appl. Phys. Lett.* **77**, 3426–3428 (2000); T. Fukumura, M. Kawasaki, Z. Jin, H. Kimura, Y. Yamada, M. Haemori, Y. Matsumoto, K. Inaba, M. Murakami, R. Takahashi, T. Hasegawa, H. Koinuma, Combinatorial search for transparent oxide diluted magnetic semiconductors, in Proceedings of the Materials Research Society, vol. 700 (2001) S2.6

17. A. Kafizas, G. Hyett, I.P. Parkin, Combinatorial atmospheric pressure chemical vapour deposition (cAPCVD) of a mixed vanadium oxide and vanadium oxynitride thin film. *J. Mater. Chem.* **19**, 1399–1408 (2009)
18. R. Takahashi, H. Kubota, M. Murakami, Y. Yamamoto, Y. Matsumoto, H. Koinuma, Design of combinatorial shadow masks for complete ternary-phase diagramming of solid state materials. *J. Comb. Chem.* **6**, 50–53 (2004)
19. R. Wendelbo, D.E. Akporiaye, A. Karlsson, M. Plassen, A. Olafsen, Combinatorial hydrothermal synthesis and characterisation of perovskites. *J. Eur. Ceram. Soc.* **26**, 849–859 (2006)
20. J.R.G. Evans, M.J. Edirisinghe, P.V. Coveney, J. Eames, Combinatorial searches of inorganic materials using the ink-jet printer; science, philosophy and technology. *J. Eur. Ceram. Soc.* **21**, 2291–2299 (2001)
21. C.J. Vess, J. Gilmore, N. Kohrt, P.J. McGinn, Combinatorial synthesis of oxide powders with an autopipetting system. *J. Comb. Chem.* **6**, 86–90 (2004)
22. S. Yang, J.R.G. Evans, Device for preparing combinatorial libraries in powder metallurgy. *J. Comb. Chem.* **6**, 549–555 (2004)
23. J. Ding, J. Bao, S. Sun, Z. Luo, C. Gao, Combinatorial discovery of visible-light driven photocatalysts based on the ABO₃-type (A) Y, La, Nd, Sm, Eu, Gd, Dy, Yb, B) Al and In) binary oxides. *J. Comb. Chem.* **11**, 523–526 (2009)
24. A. Cabañas, J.A. Darr, E. Lester, M. Poliakoff, Continuous hydrothermal synthesis of inorganic materials in a near-critical water flow reactor; the one-step synthesis of nano-particulate Ce_{1-x}Zr_xO₂ (x=0-1) solid solutions. *J. Mater. Chem.* **11**, 561–568 (2001)
25. R. Wendelbo, D.E. Akporiaye, A. Karlsson, M. Plassen, A. Olafsen, Combinatorial hydrothermal synthesis and characterisation of perovskites. *J. Eur. Ceram. Soc.* **26**, 849–859 (2006)
26. I. Yanase, T. Ohtaki, M. Watanabe, Combinatorial study on nano-particle mixture prepared by robot system. *Appl. Surf. Sci.* **189**, 292–299 (2002)
27. T.A. Stegk, R. Janssen, G.A. Schneider, High-throughput synthesis and characterization of bulk ceramics from dry powders. *J. Comb. Chem.* **10**, 274–279 (2008)
28. Y. Zhan, L. Chen, S. Yang, J.R.G. Evans, Thick film ceramic combinatorial libraries: the substrate problem. *QSAR Comb. Sci.* **26**, 1036–1045 (2007)
29. M.M. Mohebi, J.R.G. Evans, A drop-on-demand ink-jet printer for combinatorial libraries and functionally graded ceramics. *J. Comb. Chem.* **4**, 267–274 (2002)
30. Z.-L. Luo, B. Geng, J. Bao, C. Gao, Parallel solution combustion synthesis for combinatorial materials studies. *J. Comb. Chem.* **7**, 942–946 (2005)
31. T.-S. Chan, C.-C. Kang, R.-S. Liu, L. Chen, X.-N. Liu, J.-J. Ding, J. Bao, C. Gao, Combinatorial study of the optimization of Y₂O₃:Bi. *Eu Red Phosphors.* *J. Comb. Chem.* **9**, 343–346 (2007)
32. T.-S. Chan, Y.-M. Liu, R.-S. Liu, Combinatorial search for green and blue phosphors of high thermal stabilities under UV excitation based on the K(Sr_{1-x-y})PO₄:Tb³⁺_xEu²⁺_y system. *J. Comb. Chem.* **10**, 847–850 (2008)
33. J. Wang, J.R.G. Evans, London University Search Instrument: a combinatorial robot for high-throughput methods in ceramic science. *J. Comb. Chem.* **7**, 665–672 (2005)
34. R.C. Pullar, Y. Zhang, L. Chen, S. Yang, J.R.G. Evans, N. McN, Alford, manufacture and measurement of combinatorial libraries of dielectric ceramics, part I: physical characterisation of Ba_{1-x}Sr_xTiO₃ libraries. *J. Eur. Ceram. Soc.* **27**, 3861–3865 (2007)
35. R.C. Pullar, Y. Zhang, L. Chen, S. Yang, J.R.G. Evans, P.Kr. Petrov, A.N. Salak, D.A. Kiselev, A.L. Kholkin, V.M. Ferreira, N.McN. Alford, Manufacture and measurement of combinatorial libraries of dielectric ceramics, part II: dielectric measurements of Ba_{1-x} libraries. *J. Eur. Ceram. Soc.* **27**, 4437–4443 (2007)
36. R.C. Pullar, Y. Zhang, L. Chen, S. Yang, J.R.G. Evans, A.N. Salak, D.A. Kiselev, A.L. Kholkin, V.M. Ferreira, N. McN, Alford, dielectric measurements on a novel Ba_{1-x} (BCT) bulk ceramic combinatorial library. *J. Electroceram.* **22**, 245–251 (2009)
37. J.C.H. Rossiny, S. Fearn, J.A. Kilner, Y. Zhang, L. Chen, Combinatorial searching for novel mixed conductors. *Solid State Ion.* **177**, 1789–1794 (2006)
38. B. Wessler, V. Jehanno, W. Rossner, W.F. Maier, Combinatorial synthesis of thin film libraries for microwave dielectrics. *Appl. Surf. Sci.* **223**, 30–34 (2004)

39. M.L. Green, P.K. Schenck, K.-S. Chang, J. Ruglovsky, M. Vaudin, Higher- κ dielectrics for advanced silicon microelectronic devices: a combinatorial research study. *Microelectron. Eng.* **86**, 1662–1664 (2009)
40. R.-P. Herber, C. Schröter, B. Wessler, G.A. Schneider, High throughput screening of piezoelectric response of ferroelectric thin films with automated scanning probe microscopy. *Thin Solid Films* **516**, 8609–8682 (2008)
41. J.L. Jones, A. Pramanick, J.E. Daniels, High-throughput evaluation of domain switching in piezoelectric ceramics and application to $\text{PbZr}_{0.6}$ doped with La and Fe. *Appl. Phys. Lett.* **93**, (152904) (2008)
42. T. Chikyow, P. Ahmet, K. Nakajima, T. Koida, M. Takakura, M. Yoshimoto, H. Koinuma, A combinatorial approach in oxide/semiconductor interface research for future electronic devices. *Appl. Surf. Sci.* **189**, 284–291 (2002)
43. S. Guerin, B.E. Hayden, D. Pletcher, M.E. Rendall, J.-P. Suchsland, L.J. Williams, Combinatorial approach to the study of particle size effects in electrocatalysis: synthesis of supported Gold nanoparticles. *J. Comb. Chem.* **8**, 791–798 (2006)
44. P. Cong, A. Dehestani, R. Doolen, D.M. Giaquinta, S. Guan, V. Markov, D. Poojary, K. Self, H. Turner, W.H. Weinberg, Combinatorial discovery of oxidative dehydrogenation catalysts within the Mo-V-Nb-O system. *Proc. Natl. Acad. Sci. USA* **96**, 11077–11080 (1999)
45. S.J. Henderson, A.L. Hector, M.T. Weller, High throughput synthesis of pigments by solution deposition. *Mater. Res. Soc. Symp. Proc.* **848**(FF3.17), 151–156 (2005)
46. J. Scheidtmann, A. Frantzen, G. Frenzer, W.F. Maier, A combinatorial technique for the search of solid state gas sensor materials. *Meas. Sci. Tech.* **16**, 119–127 (2005)
47. K. Takada, K. Fujimoto, T. Sasaki, M. Watanabe, Combinatorial electrode array for high-throughput evaluation of combinatorial library for electrode materials. *Appl. Surf. Sci.* **223**, 210–213 (2004)
48. C.H. Olk, Infrared screening of combinatorially prepared hydrogen sorbing metal alloys. *Mater. Res. Soc. Symp. Proc.* **801**, 75–88 (2003)
49. Y. Matsumoto, M. Murakami, T. Shono, T. Hasegawa, T. Fukumura, M. Kawasaki, P. Ahmet, T. Chikyow, S. Koshihara, H. Koinuma, Room-temperature ferromagnetism in transparent transition metal-doped titanium dioxide. *Science* **291**, 854–6 (2001)
50. R.B. van Dover, L.F. Schneemeyer, R.M. Fleming, Discovery of a useful thin film dielectric using a combinatorial-spread approach. *Nature* **392**, 162–164 (1998)
51. H. Chang, I. Takeuchi, X.-D. Xiang, A low loss composition region identified from a thin film composition spread of $(\text{Ba}_{1-x-y}\text{Sr}_x\text{Ca}_y)\text{TiO}_3$. *Appl. Phys. Lett.* **74**, 1165–1167 (1999)
52. G. Briceño, H. Chang, X. Sun, P.G. Schultz, X.-D. Xiang, A class of Cobalt Oxide magnetoresistance materials discovered with combinatorial synthesis. *Science* **270**, 273–275 (1995)
53. R.C. Pullar, Combinatorial bulk ceramic magnetoelectric composite libraries of strontium hexaferrite and barium titanate. *ACS Comb. Sci.* **14**, 425–433 (2012)
54. C. Gao, B. Hu, I. Takeuchi, K.-S. Chang, X.-D. Xiang, G. Wang, Quantitative scanning evanescent microwave microscopy and its applications in characterization of functional materials libraries. *Meas. Sci. Technol.* **16**, 248–260 (2005)
55. U. Simon, D. Sanders, J. Jockel, C. Hepel, T. Brinz, Design strategies for multielectrode arrays applicable for high-throughput impedance spectroscopy on Novel gas sensor materials. *J. Comb. Chem.* **4**, 511–515 (2002)
56. I. Takeuchi, W. Yang, K.-S. Chang, M. Aronova, R.D. Vispute, T. Venkatesan, L.A. Bendersky, Monolithic multi-channel UV detector arrays and continuous phase evolution in $\text{Mg}_x\text{Zn}_{1-x}\text{O}$ composition spreads. *J. Appl. Phys.* **94**, 7336–7340 (2003)
57. Y.K. Yoo, F. Duewer, H. Yang, D. Yi, J.-W. Li, X.-D. Xiang, Room-temperature electronic phase transitions in the continuous phase diagrams of perovskite manganites. *Nature* **406**, 704–708 (2000)
58. P.-A.W. Weiss, C. Thome, W.F. Maier, MS-express: data-extracting and -processing software for high-throughput experimentation with mass spectrometry. *J. Comb. Chem.* **6**, 520–529 (2004)

59. J. Klein, S.A. Schunk, IR-Sensography™—expanding the scope of contact-free sensing methods. *Meas. Sci. Tech.* **16**, 221–228 (2005)
60. U. Simon, D. Sanders, J. Jockel, T. Brinz, Setup for high-throughput impedance screening of gas-sensing materials. *J. Comb. Chem.* **7**, 682–687 (2005)
61. Combinatorial and artificial intelligence methods in materials science, in *MRS Proceedings Volume 700* (2001), <http://www.mrs.org/publications/epubs/proceedings/fall2001/s/>
62. M.Z. Pesenson, S.K. Suram, J.M. Gregoire, Statistical analysis and interpolation of compositional data in materials science. *ACS Comb. Sci.* **17**, 130–136 (2015)
63. A.G. Kusne, T. Gao, A. Mehta, L. Ke, M.C. Nguyen, K.-M. Ho, V. Antropov, C.-Z. Wang, M.J. Kramer, C. Long, I. Takeuchi, On-the-fly machine-learning for high-throughput experiments: search for rare-earth-free permanent magnets. *Sci. Rep.* **4**, 6367 (2014)
64. G. Pilania, C. Wang, X. Jiang, S. Rajasekaran, R. Ramprasad, Accelerating materials property predictions using machine learning. *Sci. Rep.* **3**, 2810 (2013)
65. A. Yosipof, O.E. Nahum, A.Y. Anderson, H.-N. Barad, A. Zaban, H. Senderowitz, Data Mining and Machine Learning Tools for Combinatorial Material Science of All-Oxide Photovoltaic Cells. *Mol. Inf.* (2015). doi:10.1002/minf.201400174
66. R. Potyrailo, K. Rajan, K. Stoewe, I. Takeuchi, B. Chisholm, H. Lam, Combinatorial and high-throughput screening of materials libraries: review of state of the art. *ACS Comb. Sci.* **13**, 579–633 (2011)
67. K. Rajan, Materials informatics. *Mater. Today* **8**(10), 38–45 (2005)
68. C.J. Long, D. Bunker, X. Li, V.L. Karen, I. Takeuchi, Rapid identification of structural phases in combinatorial thin-film libraries using x-ray diffraction and non-negative matrix factorization. *Rev. Sci. Instrum.* **80**, 103902 (2009)
69. D. Kan, R. Suchoski, S. Fujino, I. Takeuchi, Combinatorial investigation of structural and ferroelectric properties of A- and B-site Co-doped BiFeO₃ thin films. *Integr. Ferroelectr.* **111**, 116–124 (2009)

Chapter 14

High Throughput Combinatorial Experimentation + Informatics = Combinatorial Science

Santosh K. Suram, Meyer Z. Pesenson and John M. Gregoire

Abstract Many present, emerging and future technologies rely upon the development high performance functional materials. For a given application, the performance of materials containing 1 or 2 elements from the periodic table have been evaluated using traditional techniques, and additional materials complexity is required to continue the development of advanced materials, for example through the incorporation of several elements into a single material. The combinatorial aspect of combining several elements yields vast composition spaces that can be effectively explored with high throughput techniques. State of the art high throughput experiments produce data which are multivariate, high-dimensional, and consist of wide ranges of spatial and temporal scales. We present an example of such data in the area of water splitting electrocatalysis and describe recent progress on 2 areas of interpreting such vast, complex datasets. We discuss a genetic programming technique for automated identification of composition-property trends, which is important for understanding the data and crucial in identifying representative compositions for further investigation. By incorporating such an algorithm in a high throughput experimental pipeline, the automated down-selection of samples can empower a highly efficient tiered screening platform. We also discuss some fundamental mathematics of composition spaces, where compositional variables are non-Euclidean due to the constant-sum constraint. We describe the native simplex space spanned by composition variables and provide illustrative examples of statistics and interpolation within this space. Through further development of machine learning algorithms and their prudent implementation in the simplex space, the data informatics community will establish methods that derive the most knowledge from high throughput materials science data.

S.K. Suram · M.Z. Pesenson · J.M. Gregoire (✉)
Joint Center for Artificial Photosynthesis, California Institute of Technology, Pasadena,
CA 91125, USA
e-mail: gregoire@caltech.edu

© Springer International Publishing Switzerland 2016
T. Lookman et al. (eds.), *Information Science for Materials
Discovery and Design*, Springer Series in Materials Science 225,
DOI 10.1007/978-3-319-23871-5_14

14.1 Tailoring Material Function Through Material Complexity: The Utility of High Throughput and Combinatorial Methods

Many technological industries ranging from manufacturing to renewable energy rely on the discovery of new high-performance solid state materials. A common approach to discovery of advanced materials is through increasing chemical complexity, for example through the incorporation of several elements into a single material. This long-standing approach in materials research traditionally involves the synthesis and evaluation of one composition at a time. Most of the single-element and binary-composition spaces were effectively investigated in the 20th century by this low-throughput method, and the frontier has thus been pushed to higher order ternary, quaternary, etc. composition spaces. Due to the vast number of possible sets of elements and compositions in a given composition space, systematic experimental investigation of these high-order composition spaces require sophisticated tools for high throughput synthesis and evaluation of new compositions. Recent advancements in experimental methods for the rapid synthesis of material libraries and rapid measurement of material properties are yielding vast ensembles of complex data [14, 17, 19, 46, 49, 60, 66]. A tenet of materials science is the development of composition-property relationships, and the automated identification of relationships within high throughput datasets requires the development of new informatics tools.

In this chapter we discuss a high throughput experimental pipeline which motivates the development of specific informatics tools. In particular, we note the importance of tiered screening wherein a high throughput pipeline contains a series of experimental measurements that operate at disparate sample throughput. To avoid bottlenecks, a sample down-selection method must be implemented. The informatics challenge arises in the automated identification of a subset of samples for lower throughput measurements such that the selected subset retains maximal “information content,” or maximal ability to establish composition-property relationships with the incomplete dataset. With high throughput datasets in hand, analysis of compositional trends requires prudent practices for the statistical analysis of compositional data. We review unique attributes of compositional data, and through illustrative examples, show that informatics and statistical algorithms for compositional data must account for the non-Euclidean nature of compositional variables.

14.2 Materials Datasets as an Instance of Big Data

High throughput materials science requires handling enormous amounts of complex data produced by modern high-throughput experimental technologies. Many modern techniques produce data beyond what can be readily processed, and even fields with well-established data archives and methodologies, such as genomics, are facing new and mounting challenges in data management and exploration. Besides the size in

Table 14.1 The number of unique compositions in a discrete composition library is shown for several values of the number of components n and composition steps δ

num. steps	10	20	30	40	50
δ	10 %	5 %	3.33 %	2.5 %	2 %
$n = 2$	11	21	31	41	51
$n = 3$	66	231	496	861	1,326
$n = 4$	286	1,771	5,456	12,341	23,426
$n = 5$	1,001	10,626	46,376	135,751	316,251
$n = 6$	3,003	53,130	324,632	1,221,759	3,478,761

The number of δ steps (“num. steps”) between 0 and 100 % is also listed

bytes, these modern data sets are multivariate, high-dimensional and consist of wide ranges of spatial and temporal scales. All of this severely restricts the capability of traditional approaches to modeling, analysis, and visualization of data. The data not only consist of vectors in Euclidean spaces, but may also include new types of data (e.g. tensor fields), or any of those data types defined not just on a Euclidean space, but on manifolds or graphs as we discuss for compositional variables.

To conceptualize the extent of data in explorations of high order composition spaces, one can count the number of discrete composition samples required to cover the composition space with a fixed composition interval. A composition space of n components contains $n - 1$ degrees of freedom due to *closure*, the requirement that the individual concentrations sum to 1. Table 14.1 provides the number of unique compositions in an n -component composition space with composition interval δ , for several illustrative values of these parameters. For unexplored composition spaces, using a fine composition step is desirable to mitigate the possibility of missing a new high performance material, and to explore high order composition spaces, high throughput techniques are required. The systematic exploration of the various combinations of the components is an example of combinatorial materials science, and given the vast number of possible combinations and the speed at which they are probed experimentally, we refer to these investigations as high throughput (HiTp).

For each composition, experimental investigations may include a scalar measurement of performance (so-called “FOM”, Figure of Merit), multi-dimensional data such as images and spectra, or any combination thereof. In particular, the extent of data may vary for each composition in a given library, and the disparate dimensionality of data increases the data complexity and the challenges for informatics algorithm development.

Data mining and machine learning embrace many sophisticated data structures. One of the main difficulties in data mining is caused by dependencies between multiple variables/parameters. Identifying a set of independent variables/parameters can be seen as a particular case of a general approach called data dimensionality reduction. When data points are close to a hyperplane in a Euclidean space, methods such as principal component analysis (PCA) and correspondence analysis (CA) are widely used for dimension reduction. These methods allow one to extract major dependen-

cies between physical variables. In case when data lie on a non-Euclidean space, more sophisticated methods of analysis are required. Complex data sets cannot be adequately understood without detecting various scales that might be present in the signal. However, traditional multiresolution analysis (MRA) tools based on wavelets are restricted mostly to one-dimensional or two-dimensional signals. Thus, in order to accurately extract information from modern data sets, the development of multiscale analysis applicable to functions defined on manifolds and graphs is of great importance. Extending multiresolution analysis from Euclidean to curved spaces and networks presents a significant challenge to applied mathematics. This is an emerging field, which is still being developed. Wavelet-type bases and frames consisting of nearly exponentially localized band-limited functions are imperative for computational harmonic analysis and its applications in statistics, approximation theory, and so on. For the two-dimensional sphere and group of its rotations, frames have already found a number of important applications in statistics and crystallography [6, 18, 58]. An adaptive multiscale approach to data analysis based on synchronization was suggested in [59]. The approach is nonlinear, data driven in the sense that it does not rely on *a priori* chosen basis, and can be extended to automatically determine the scale for complex signals defined on graphs/manifolds (regarding analysis on compact manifolds, see also remarks in section “Composition Spread and Distances” in this chapter). Overall, MRA is a necessary, indispensable approach to efficient representation/analysis of complex information (signals, images, etc.) produced in high throughput and combinatorial experiments.

Traditional statistical methods may lead to erroneous dependencies and incorrect inferences when applied to modern complex data, as we demonstrate in this chapter. But even if data consist of usual vectors in a Euclidean space, there are still many open issues. One of them is related to the so-called null-hypothesis testing (NHST). It has been recognized lately that it is necessary to perform NHST to more instructional effect sizes, confidence intervals (CIs) by applying Meta-analysis [7, 15, 30]. Although CIs are more informative than NHST approach since, in some form, they quantify the uncertainty, their meaning is often misunderstood. In fact, CIs are intimately connected with NHST, and both are superseded by Bayesian techniques [42, 43, 73].

The complexity of data calls for application and development of adequate techniques, which are more powerful than the conventional ones and tailored to specific types of experimental data. Statistical methods are often considered simply a toolbox and are consequently utilized superficially in data analysis. To make full use of the deluge of complex data, researchers must transcend the notion of toolbox statistics and engage in the independent applied science of statistics and informatics. Combinatorial Science is Data-Driven and its main premise is that discovery and optimization of materials can be made efficient if directed by statistical inference based on the experimental data. In other words, Combinatorial Materials Science cannot be truly realized without modern statistics and more generally, informatics. Informatics here refers to the management of complete data-lifecycle: the storage, integration, compression of data as well as quantification of uncertainty and mining/analysis of data via statistical learning and data mining techniques.

Moreover, experimental techniques for the generation of data are often developed independent from the development of analysis techniques for that data. Statistical analysis should not be subsequent, but should rather be a part of the experimental design [12]. In this chapter we describe the development of both experimental techniques and analytical methods, and while these developments cannot take place strictly simultaneously, we note the importance of iterative developments of both sides of the high throughput methods.

14.3 High Throughput Experimental Pipelines: The Example of Solar Fuels Materials Discovery

A high throughput experimental pipeline is comprised of a network of experimental methods that are interlinked in a process workflow to enable a complete cycle of high throughput experiments [26]. The high-level summary of a high throughput pipeline for the discovery of solar fuels materials is shown in Fig. 14.1. The pipeline contains 3 primary sectors: the synthesis of material libraries, the screening of materials via measurements of material performance, [29, 31, 32, 34, 39, 51, 74, 75] and the characterization of materials [28]. The materials screening portion is split according to two primary types of functional materials for solar fuels technology, and the screening of light absorber and of electrocatalyst materials each involves unique experiments.

Several data-related aspects of the pipeline that are not shown are data management, data analysis, and design of experiments. The informatics-based aspect that is shown relates to the active down-selection of samples. That is, to create a throughput-matched series of screening experiments, a higher throughput coarse screening method is coupled to a lower throughput fine screening method through the judicious selection of a subset of the samples. The 3 electrocatalyst screening experiments listed in Fig. 14.1 have been described in recent publications with the higher throughput method being the parallel imaging of O₂ bubbles produced by electrocatalysis of the oxygen evolution reaction (OER) [75]. The two other experiments are serial experiments performed by a scanning drop cell (SDC) device




Fig. 14.1 (top) Sectors of the accelerated discovery pipeline, with the screening sector split for the 2 general material functions of light absorption and electrocatalysis [26]. (bottom) Tiered screening experiments are shown for evaluating electrocatalyst libraries where sample down-selection occurs between subsequent screening experiments

to quantify OER electrocatalytic activity [29]. These experiments include the collection of cyclic voltammogram with rapid sweep rate and then a longer measurement of catalyst overpotential at a fixed current density, with the experiment duration being sufficiently long to demonstrate that any anodic current could not be dominated by a sample corrosion process. While the throughput of each technique depends on the choice of experimental parameters, for the screening of material libraries with approximately 1800 composition samples on a library plate, the throughput of each stage is approximately 180, 10 and 2 samples per minute, respectively. While some throughput matching can be achieved through duplication of instruments for performing the lower-throughput techniques, practical throughput matching is attained through sample down-selection at each juncture.

Transition from the screening portion of the pipeline to materials characterization often involves another substantial down-selection of samples. While development of HiTp materials characterization [27, 35, 40] and related analysis techniques [44, 45] is an active field of research, the characterization throughput is often lower than the final screening in a tiered screening pipeline. For a given composition region, a systematic variation in a materials characterization attribute may correspond to a variation in the performance metric. By partitioning a composition space into regions which exhibit systematic trends in performance, samples can be selected for detailed characterization to capture the attribute-property relationships both within and among the composition regions. Implementing this strategy into a down-selection algorithm is a primary goal of informatics for high throughput pipelines.

14.4 An Illustrative Dataset: Ni-Fe-Co-Ce Oxide Electrocatalysts for the Oxygen Evolution Reaction

In the following sections, we present the challenges and initial progress in two areas of informatics related to high throughput materials discovery: the automated down-selection of samples in a tiered screening pipeline, and the statistical analysis of compositional variables as a critical aspect of identifying composition-property relationships. Both of these discussions will use simple, synthetic datasets as illustrative examples. In addition, examples will be provided using an experimental dataset from the high throughput mapping of OER catalyst activity over a pseudo-quaternary composition space of metal oxides containing all possible combinations of Ni, Fe, Co and Ce with 3.33 at.% intervals. For details on materials synthesis and experimental methods, we refer the reader to previous reports. [31, 34] Here we provide a map of a primary figure of merit for OER electrocatalysts for solar fuels applications, the overpotential required to provide 10 mA cm⁻² geometric catalytic current density. The results are summarized in Fig. 14.2 with Fig. 14.2a showing an example map of the FOM for an array of samples on a library plate, which are mapped onto composition space in Fig. 14.2b. The composition mapping of the pseudo-quaternary spread is performed as a stack of pseudo-ternary triangles with increasing Ce concentration.

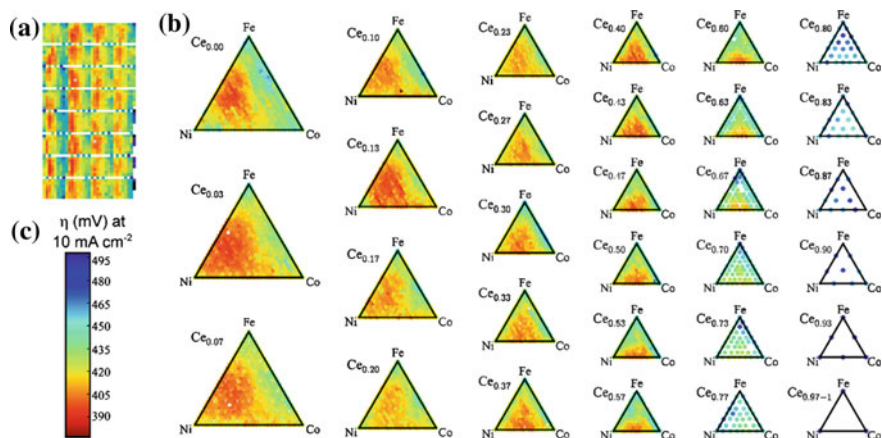


Fig. 14.2 A FOM for solar fuels applications, the overpotential for delivering OER geometric current density of 10 mA cm^{-2} , is measured on composition libraries **a** and mapped to composition space **b** using a false color scale **c**. The (Ni-Fe-Co-Ce) O_x composition space is shown as a stack of Ni-Co-Ce composition plots with increasing Ce concentration [26]

The common FOM color scale is shown in Fig. 14.2c. This dataset is most representative of the third tier of electrocatalyst screening described above, although it is a full dataset for which we can perform down-selection informatics to choose a sample subset for additional screening or characterization experiments. We can also analyze FOM trends with (sub-)compositional variables, as will be illustrated in the final section.

14.5 Automating Sample Down-Selection for Maximal Information Retention: Clustering by Composition-Property Relationships

As described above, HiTp experimentation typically involves the coarse, rapid measurement of a FOM or property of interest for each sample in a material library. Appropriate down-selection methods are essential to ensure generation of information-rich experimental data that lead to knowledge and discovery. While a combinatorial material library may include variation of a number of process parameters such as synthesis temperature or processing parameters, [11, 13] we continue this discussion in the context of composition libraries.

For demonstration purposes, a synthetic dataset with four distinct composition regions that are governed by different composition-property relationships is shown in Fig. 14.3a and the resulting down-selection using top ‘z’ percentile performing compositions is shown in Fig. 14.3b. It is evident that down-selection based on top ‘z’

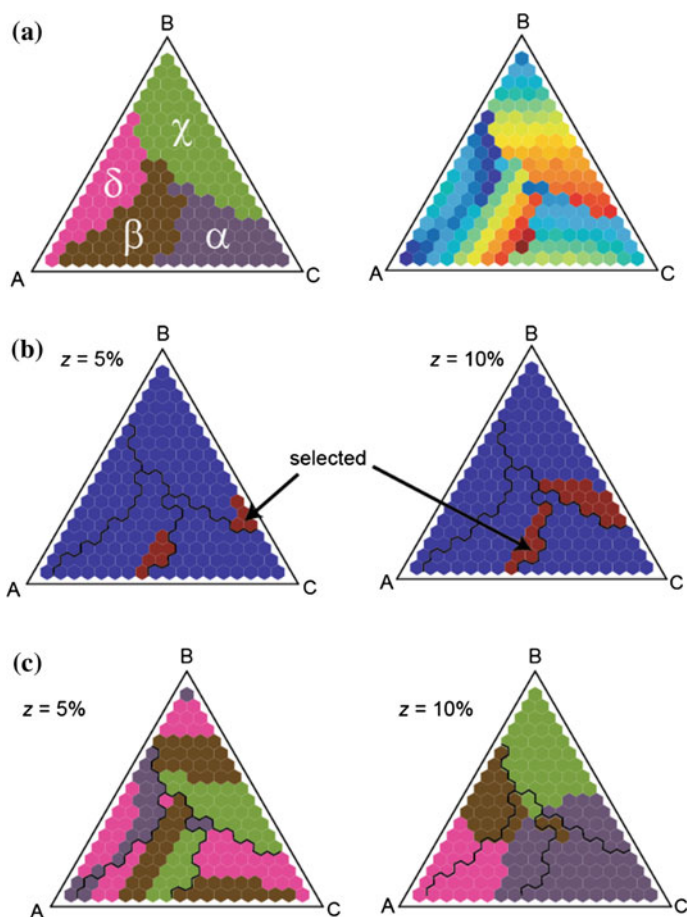


Fig. 14.3 **a** A ternary composition space (with 5 at.% step) is partitioned into 4 property fields (*left*), and a synthetic composition-property plot is obtained by applying distinct polynomial functions to the compositions of each property field (*right*), **b** down-selection of compositions by selecting top ‘z’ percentile of compositions based on their property value. The downselected compositions, colored red, are very sensitive to the choice of ‘z,’ which is usually fixed based on throughput matching of successive experiments. The property field boundaries are overlaid for comparison and **c** Clustering of the ternary composition library using a Euclidean distance metric on the property space (*left*) and composition-property space (*right*). Clustering using only the property yields clusters with compositions scattered over the library, while adding the compositions to the clustering metric yields clusters that are mostly connected in composition space but do not match the original property fields, whose boundaries are overlaid for comparison

performers is highly sensitive to the value of z, typically imposed by the throughput capabilities of the HTE workflow and more importantly are incapable of capturing the composition-property relationships thus necessitating the need for more sophisticated partitioning/clustering techniques.

Traditional clustering techniques such as k-means clustering on either composition-property space or property space (Fig. 14.3c) alone that depend on spatial statistics are incapable of capturing composition-property relationships. In this context, we discuss the role of evolutionary statistical methods and information theory concepts in identifying several composition-property relationships and generating information rich experimental databases.

14.5.1 Down-Selection for Maximal Information Content

The above discussion on tiered pipeline screening introduces the importance of down-selection for maximal information content in the context of a high-throughput workflow. From a materials discovery point of view, maximal information down-selection in a HiTp pipeline allows generation of information-rich experimental databases that allow us to extract knowledge pertinent to composition-(micro)structure-experimental parameters-property relationships. This knowledge is typically un-accessible via other main facets of materials design based discoveries namely, first principles computations and materials informatics as applied to existing databases. Thus, generation of information rich experimental databases provides a unique opportunity to exploit the capabilities of HiTp to simultaneously perform exploratory and knowledge based search for new materials. These experimental databases can be used as an input to data mining methods to extract empirical relationships, they also form an important resource to develop sophisticated first principles based models that are applicable to higher order compositions spaces and in-operando conditions.

Distance and density based clustering approaches that are typically ubiquitous in the field of clustering and have been successfully applied for materials science applications where spatial statistics are relevant are inapplicable for partitioning the composition space for maximizing information content. Alternately, information theory based metrics provide access to higher order statistics [22, 36, 37] necessary for clustering/classification in complex data structures. Specifically, Shannon entropy criterion has been successfully applied as a supervised classification algorithm for unravelling crystal chemistry design rules [41] and discovery of materials [4]. The selection, crossover and mutation based evolutionary operations of genetic programming enable complex data relationships to be captured as genetic trees, resulting in its application for supervised classification of complex data [52]. Other evolutionary techniques such as genetic algorithms [5] and particle-swarm optimization [72] have also been used for clustering data. However, they use cluster variance-based objective functions and thus are unable to capture non-hyperspherically shaped clusters which are typical of phase/property fields in materials science.

While several data mining algorithms have been applied to a) capture the function relating the input and output variables [9, 68] and/or b) cluster data based on input variables [63] and/or c) classify complex data structures in supervised classification; [47] these approaches are insufficient to cluster data based on the (dis)similarity

in the function relating input and output variables. For this purpose an approach that is capable of capturing and classifying several underlying composition-property relationships is required. Mathematically, this is achieved by identifying clusters that maximize divergence amongst composition-property relationships described by them. Genetic programming is a well-accepted and robust methodology for capturing functional relationships, whereas, divergence is measured using information theory based concepts. In the following sections, the concepts of multi-tree genetic programming as applied to a materials discovery problem using an information theory based objective function are introduced and refined. We utilize genetic programming trees to represent the functions that map compositions and HiTp property measurements to memberships in a fixed number of clusters. The clustering is defined over the composition space such that the optimized trees cluster the compositions based on the *functional relationships* between composition and measured property. This method of clustering allows selection of representative compositions from each cluster for further investigation and characterization, resulting in an information rich experimental materials genome with respect to composition-characterization attribute-property relationships.

14.5.2 Information-Theoretic Approach

Using information-theoretic approach, clustering composition space such that the similarity of composition-property relationships among different clusters is minimized while similarity of composition-property relationships within a given cluster is maximized can be represented as minimizing cross “between cluster” information potential while maximizing self “within cluster” information potential. An attractive metric to achieve this for a two class system is the Cauchy-Schwarz divergence [38, 64] and is expressed as

$$D_{cs}(p_1, p_2) = -\ln \frac{\int p_1(x)p_2(x)dx}{\sqrt{\int p_1^2(x)dx \int p_2^2(x)dx}}, \quad (14.1)$$

where $p_k(x)$ is the probability distribution of x in class C_k and x is the (multi-dimensional) composition coordinate.

In case of discrete data; probability distribution functions can be estimated using a Parzen window [38] with a Gaussian kernel:

$$p(x) = \frac{1}{n} \sum_{i=1}^n G(x - x_i, \sigma^2), \text{ where } G(x - x_i, \sigma^2) = \frac{1}{(2\pi\sigma^2)^{d/2}} \exp\left(-\frac{\|x - x_i\|^2}{2\sigma^2}\right) \quad (14.2)$$

The kernel width, σ , is an a priori specified parameter; n is number of observations; d is the dimension of the dataset.

Using (14.2), Jenssen et al. [37] show that the divergence function of (14.1) can be estimated as

$$D_{cs}(p_1, p_2) \approx -\ln \frac{\sum_{x_i \in C_1} \sum_{x_j \in C_2} G_{ij, 2\sigma^2}}{\sqrt{\sum_{x_i, x_{i'} \in C_1} G_{ii', 2\sigma^2} \sum_{x_j, x_{j'} \in C_2} G_{jj', 2\sigma^2}}} \quad (14.3)$$

where $G_{ij, \sigma^2} = G(x_i - x_j, \sigma^2)$

The fact that every composition in a composition library belongs to exactly one property field is imposed using a membership value (${}^i m_k$) for data point i in cluster k as

$${}^i m_{k'} = 1 \text{ for } k' = k \quad \& \quad {}^i m_{k'} = 0 \text{ for } k' \neq k. \quad (14.4)$$

And, ${}^i \mathbf{m}$ is defined as the vector of membership values for a data point i over the set of clusters. Using these membership notations, [8] extend the Cauchy-Schwarz divergence function to a c -cluster problem ($c \geq 2$), as:

$$D_{cs}(p_1, p_2, \dots, p_c) \approx -\ln \frac{\frac{1}{2} \sum_{i, j=1}^n (1 - {}^i \mathbf{m}^T {}^j \mathbf{m}) G_{ij, 2\sigma^2}}{\sqrt{\prod_{k=1}^c \sum_{i, j=1}^n {}^i m_k {}^j m_k G_{ij, 2\sigma^2}}} \quad (14.5)$$

In this objective function, the denominator scales as a power of the number of clusters (c) whereas the numerator varies comparatively very slowly with c , resulting in a denominator dominated objective function with increasing number of clusters. Therefore, we introduce a modified form of Cauchy-Schwarz divergence function such that the numerator and denominator remain invariant to the number of clusters and is described as:

$$D_{cs}(p_1, p_2, \dots, p_c) \approx -\ln \frac{\sqrt{\left(\sum_{i, j=1}^n (1 - {}^i \mathbf{m}^T {}^j \mathbf{m}) G_{ij, 2\sigma^2} \right) \left(\frac{c}{c-1} \right)}}{c \left(\prod_{k=1}^c \sum_{i, j=1}^n {}^i m_k {}^j m_k G_{ij, 2\sigma^2} \right)^{\frac{1}{2c}}} \quad (14.6)$$

To introduce the modified Cauchy-Schwarz divergence function (14.6) in an optimization algorithm, a continuous membership function is required, because the binary membership defined in (14.4) does not provide a continuous divergence function with respect to alterations in membership of a given data point in a given cluster. Further, to accurately cluster property fields, the membership values should be based on composition-property relationships. To facilitate this, continuous membership values in the range [0, 1] are introduced by defining a membership function $m_k(x_f)$ for each cluster such that ${}^i m_k = m_k(x_f^i)$. The domain for the probability distribution functions for Parzen window estimation is the composition space which enables compositional connectedness in the clusters, whereas the domain for mem-

bership functions is a combined composition and property space, with coordinate represented as xf which enables the membership functions to represent composition-property relationships. Additionally, by constraining the membership values to sum to one, they can be regarded as a set of posterior probabilities:

$$m_k(xf) = P(C_k|xf), \sum_{k=1}^c m_k(xf) = 1. \quad (14.7)$$

14.5.3 Genetic Programming Based Clustering

Genetic trees are computer programs capable of learning complex relationships present in the data. In a c -class dataset, there are c functional relationships between composition and property that need to be learnt or distinguished from each other. Thus, we utilize a multi-tree genetic programming (MT-GP) framework developed by Muni et al., [52] and [8] such that each tree learns the functional relationship between composition and property for one of the classes in the data. In this representation each tree (T_k) is defined on the composition-property space where the scalar is used $T_k(xf)$ to generate membership values, $m_k(xf)$ as described below. This reduces the functional relationship based clustering problem to optimal identification of composition-property relationships by MT-GP such that the resulting membership values maximize the Cauchy-Schwarz divergence function (14.6).

The algorithm is based upon the construct illustrated in Fig. 14.4, where each cluster is represented by a hierarchical tree of root, leaf and terminal nodes in the MT-GP chromosome. The leaf nodes and the root nodes are chosen from the set of operators $\{+, -, \times, \div\}$. The terminal nodes are numerical and the domain includes the composition, property parameter space and random integer constants in $[0, 10]$. For the tree representing a cluster k , the sequence of operators that terminate with numeric values comprise a nested algebraic function $T_k(xf)$.

Initialization, mutation, selection, crossover and termination proceed using standard genetic programming techniques and are discussed elsewhere, [69] although crossover in a MT-GP approach differs from crossover in traditional genetic programming. A crossover between any two selected parent chromosomes with ‘ c ’ trees can occur in ${}^c C_2$ ways because the k th tree in chromosome ‘ i ’ does not have to crossover with the k th tree in chromosome ‘ j ’ given that genetic tree-property field mapping is not necessarily the same for all chromosomes. For each pair of multi-tree chromosomes selected as parents (using a probability p_{cross} here, set to 1), pairs of trees are randomly selected with one tree from each of the parent chromosomes contributing to the pair such that every tree in the parent chromosomes is present in exactly one pair. To balance between exploratory and exploitative capabilities of genetic programming, we define a base probability ($p_{\text{treecross}}$) and probability multiplier (p_{cm}) such that the probability for crossover of the k th randomly selected pair of trees for a given pair of parent chromosomes is $p_{\text{cm}} \cdot p_{\text{treecross}}$. Values of $p_{\text{treecross}}$ in the range

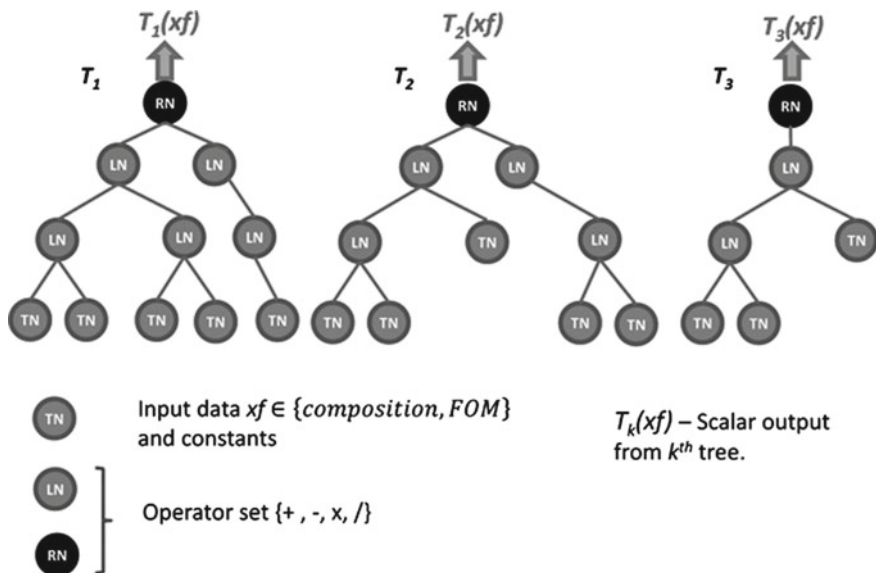


Fig. 14.4 A schematic of a multi-tree chromosome in an MT-GP approach for 3 clusters and maximum depth 3. Abbreviations used are TN: terminal node, LN: leaf node, RN: root node

0.6–0.8 and $p_{treecross} \times (p_{cm})^{k-1}$ in the range 0.8–1.0 are found to be reasonable estimates to ensure robust convergence. However, further research is required to identify optimal values of these parameters using various case studies.

14.5.4 Calculating Membership

Boric and Estévez [8] related the output of the trees $T_k(xf)$ to membership values $m_k(xf)$ using a Sigmoid transformation followed by normalization:

$$T'_k(xf) = \frac{1}{1 + e^{-T_k(xf)}} \quad \& \quad m_k(xf) = \frac{T'_k(xf)}{\sum_{k'=1}^c T'_{k'}(xf)} \tag{14.8}$$

The scalar outputs from different trees could be of varying magnitudes depending on the distinct composition-property function they map, and thus could result in membership values that are skewed towards a particular function. To avoid this, relative memberships within each class are obtained by first normalizing the output of the trees $T_k(xf)$ with respect to the minimum and maximum values of $T_k(xf)$ for a

given ‘ k ’ and then normalizing the relative memberships such that $m_k(xf)$ represent posterior probabilities (14.9).

$$T'_k(xf) = \frac{T_k(xf) - T_k^{\min}(xf)}{T_k^{\max}(xf) - T_k^{\min}(xf)} \quad \& \quad m_k(xf) = \frac{T'_k(xf)}{\sum_{k'=1}^c T'_{k'}(xf)} \quad (14.9)$$

The most representative class label set $\widehat{k}(x)$ is computed using

$$\widehat{k}(x) = \underset{k}{\operatorname{argmax}}(m_k(xf)) \quad (14.10)$$

To ensure that variations in each feature vector are given equal importance, the composition vectors and the property vectors need to be converted to unit standard deviation prior to the MT-GP analysis.

14.5.5 Application to a Synthetic Library

Figure 14.5 shows the optimal membership set obtained after clustering the dataset shown in Fig. 14.1 assuming the presence of four clusters, using a Gaussian kernel width, $\sigma = 0.17$ at. %. Figure 14.5 also shows the clustering of compositions based on their maximum membership class ($\widehat{k}(x)$). Given that the number of property fields in the synthetic dataset and the number of clusters used in the MT-GP algorithm are equal, the association of a synthetic property field and calculated cluster is easily made by evaluating the maximum intersection of the composition points. The clusters in Fig. 14.5 are colored corresponding to the association of property fields in Fig. 14.1, and comparison between these composition maps reveals 14 misclassified samples, approximately 8% of the data points. The misclassified samples lie on the boundaries




Fig. 14.5 (left) Maps of the membership of each composition in the four optimized MT-GP trees. (right) The four clusters obtained by taking the maximum membership for each composition with the property field boundaries from Fig. 14.1 overlaid for comparison. The 14 misclassified compositions are marked by red borders

between different property fields, where the continuous membership parameters show partial membership in each of the neighboring fields. That is, the MT-GP algorithm produces the correct property fields with the boundaries blurred by 1 or 2 composition intervals.

14.5.6 Experimental Dataset

To demonstrate structure-property relationship clustering on experimental data, we use the (Ni-Fe-Co-Ce) O_x catalyst performance dataset from Fig. 14.2. The 5429 FOM values and corresponding 4-component compositions are used as the input for the MT-GP algorithm with 4 trees, each with maximum depth 4, and $\sigma = 0.17$. We choose 4 clusters (4 trees) to demonstrate the capability of our algorithm to capture important composition-FOM relationships.

One of the essential genetic operators is division which allows capturing of complex composition-property relationships. However, this adds special constraints for treatment of compositions along ternary faces, binary lines, and unary end points which have at least one composition component as zero. To avoid division by zero, we shift all the compositions by $\Delta = 0.01$ at. %. Using maximum membership to define representative clusters, the stacked-ternary representation of the 4 optimal clusters obtained from MT-GP is shown in Fig. 14.6.

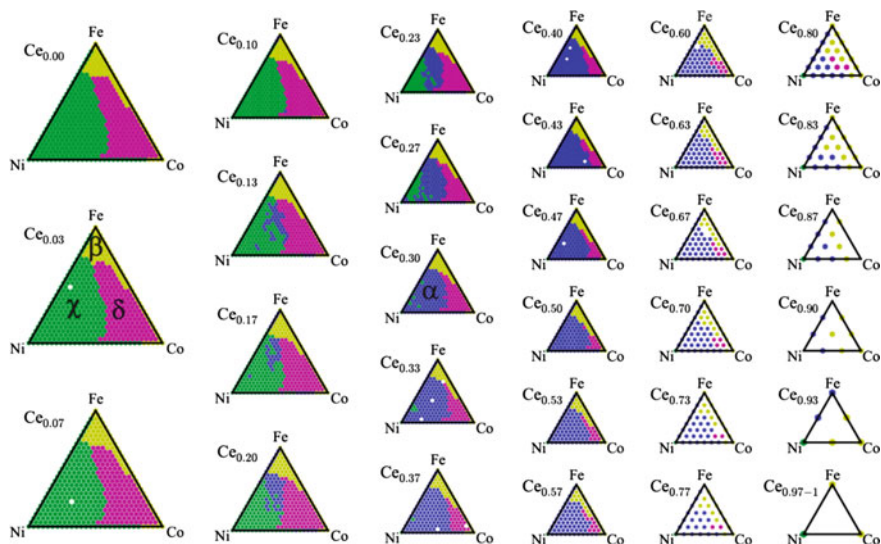


Fig. 14.6 Mapping of the most representative cluster onto quaternary compositions in a (Ni-Fe-Co-Ce) O_x library

In any experiment dataset of composition-property information, there is no known optimal solution for composition clusters. For the dataset of Fig. 14.2, two unique, highly-active catalyst composition regions have been identified and classified through additional electrochemical characterization [33]. The recently discovered catalyst composition region contains little to no Fe and approximately 50% Ce, which is identified as the α cluster. Traditional mixed-transition-metal oxides with at least approximately 50% Ni comprise the low-Ce region of highly active catalysts, which is identified as the χ cluster. The MT-GP algorithm provides information for two other clusters with lower catalytic activity. Given that the FOM explored is convoluted due to experimental noise and has limited dynamic range, the excellent clustering results suggest that the MT-GP algorithm can be successfully deployed for automated down-selection routines.

While further research is necessary to develop a non-parametric MT-GP based clustering algorithm, the approach presented establishes a protocol for identifying distinct, complex composition-property relationship fields from combinatorial materials science data and presents a significant step towards developing information rich experimental materials genomes. In addition, the compositional connectedness of clusters is encouraged by Euclidean metric based Gaussian kernels. While Gaussian kernels capture clusters effectively for the test cases demonstrated, compositional data are defined on the simplex, as described below, requiring additional development of clustering and down-selection methods.

14.6 The Simplex Sample Space and Statistical Analysis of Compositional Data

A primary objective of Combinatorial Materials Science is to unravel the composition dependence of materials properties. Probably no other field has so much of its data intrinsically expressed as percentages (compositional data) as do chemistry and combinatorial materials science. Since the percentages sum to a constant the composition sample space is not the usual Euclidean space. Indeed, the constraint of constant sum doesn't allow the components of a composition to vary from $-\infty$ to ∞ and a composition of N elements is confined to a restricted part of the Euclidean space called the simplex $S^N = \{\mathbf{x}; x_k \geq 0, \sum_{k=1}^N x_k = 1\}$ [1, 2, 53]. Conventional statistical analysis of such data does not incorporate the inherent relationships between the elements, even though they are crucial for the physics and chemistry of materials. Moreover, conventional processing of compositional data introduces artifacts such as spurious correlation, while compositional statistical methods enable more accurate extraction of composition-property relationships. The relationships between compositions and their properties are intrinsically multivariate and compositional data require special methods of processing.




Fig. 14.7 Demonstration of subcompositional incoherence using the dataset of Fig. 14.2. For each Fe concentration, the lowest overpotential value from the set of samples with that Fe concentration is shown under 2 subcompositional representations of the quinary oxides: (*black*) quantification of Fe, Co, Ce and Ni and (*red*) quantification of Fe, Co and Ce

In this section we present the importance of CDA for materials science. The concepts and importance of closure and sub-compositional incoherence are discussed, and to demonstrate the consequences of these concepts for interpretation of experimental data, we begin with an analysis of composition trends within the data of Fig. 14.2. An illustrative, practical compositional analysis is to evaluate the relationship between Fe concentration and electrocatalytic activity. To generate Fig. 14.7, we consider the discrete Fe concentration intervals of 3.33 at.% and for the set of samples with a given Fe concentration, extract the lowest overpotential value. The corresponding compositional trend indicates how good a catalyst can be with a given Fe concentration. An important realization about the discussion of this composition library is that the above figures have only considered the composition of the 4 cations, as the oxygen stoichiometry is unknown. That is, the samples have been treated as quaternary subcompositions of the quinary parent compositions. Figure 14.7 shows the composition trend calculated using the Fe-Co-Ce-Ni subcomposition space and the analogous trend using Fe-Co-Ce compositions, which may result from an experiment wherein the Ni concentrations are unknown. The striking differences between the overpotential trends using these two subcomposition spaces highlight an inherent complexity of compositional data.

Using illustrative synthetic datasets and mathematical descriptions of CDA, we demonstrate that Euclidean-based correlation structure should not be used to interpret associations among measured elemental concentrations. In particular, we demonstrate induced correlations and subcompositional incoherence of the Pearson correlation coefficient. These effects are caused by the constant sum constraint that restricts the sampling space to a simplex instead of the usual Euclidean space. Since statistical measures such as mean, standard deviation, etc., are defined for the Euclidean space, traditional correlation studies, multivariate analysis, and hypothesis testing may lead to erroneous dependencies and incorrect inferences when applied to compositional data. These issues demonstrate that prior to applying usual statistical methods data

should be transformed to remove the constant sum constraint. Logratio transforms remove the data-sum constraint by mapping the components of the compositions into a Euclidean space, thus enabling one to apply classical statistical methods. Moreover, a metric vector space structure can be introduced in the simplex (via the simplicial metric based on logratios), thus enabling meaningful statistical analysis of compositional data. We apply logratio analysis to interpolation of simulated composition data. Comparison of a consistent compositional interpolation based on balances with traditional linear approach reveal discrepancies between their results that are crucial for correct statistical analysis of composition-property relationships. Altogether these results demonstrate the importance of using physically/chemically adequate and mathematically consistent approaches to compositional data, particularly in high-order composition spaces.

14.6.1 The Closure Effects—Induced Correlation

The traditional way to describe the pattern of variability of data is through the estimates of the raw mean, covariance, and correlation matrices. Individual components of compositional data are not free to vary independently: if the proportion of one component decreases, the proportion of one or more other components must increase, thus leading to an artificial correlation that is, in fact, caused by the constant sum constraint. Indeed, the closure, or in other words the constant sum constraint, affects correlation between variables. Consider for example a set of N -part compositions that can be treated as a $M \times N$ matrix \mathbf{W} where N is the number of elements in the composition, and M is the number of measurements, or samples, with the component sum $\sum_{k=1}^N w_{ik} = 1$, where $i = 1, \dots, M$. Let us denote $Y_k = w_{ik}$, $i = 1, \dots, M$, to be the k th column of the matrix \mathbf{W} . Since

$$\text{cov}(Y_k, \sum_{j=1}^N Y_j) = 0$$

we have

$$\sum \text{cov}(Y_k, Y_j) = -\text{var}(Y_k), j \neq k \quad (14.11)$$

so the sum of the covariances of any variable is negative. Thus each variable must be negatively correlated with at least one other variable and, in general, there is a strong bias toward negative correlation between variables of (relatively) large variance. One of the critical consequences of closure for materials science is that usual correlation analysis can produce misleading associations between elemental concentrations. This is especially consequential since visualization of results by Composition-Structure-Property correlation maps is so important in materials science [10, 76]

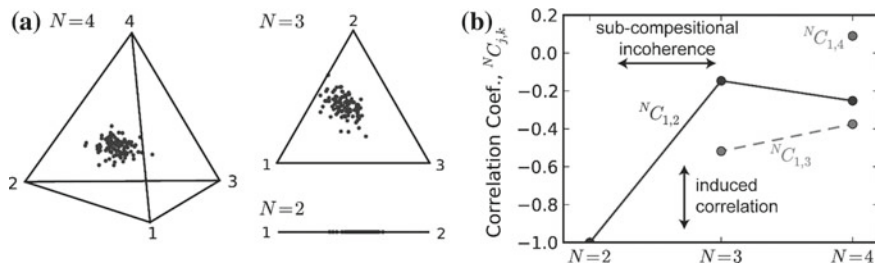


Fig. 14.8 **a** A set of 100 compositions generated from normal distributions of element quantities with normalization into the quaternary ($N = 4$), ternary ($N = 3$) and binary ($N = 2$) composition spaces. **b** The correlation of the concentration of element 1 with each other element is shown, with the magnitudes demonstrating induced correlation, and their variation with respect to N showing sub-compositional incoherence [61]

14.6.2 Illustrative Example

As an example, consider a set of M materials each containing N elements for which we would like to ascertain if there is correlation in the concentration of element 1 with respect to the other elements. As shown in Fig. 14.8a, a synthetic dataset is created by generating random quantities of the $N = 4$ elements from normal distributions. Due to the randomness, the element-pairwise correlation over the $M = 400$ materials is negligible when considering the quantities of the elements, which is non-normalized data. Measurements of the (normalized) composition of each material produce the $M \times N$ closed dataset w_{ik} . Using this simulated data, the Pearson correlation coefficient ${}^4C_{k,l}$ of the concentration vectors \mathbf{Y}_k and \mathbf{Y}_l (elements k and l) can be calculated, where the superscript 4 indicates the dimension of the composition space ($N = 4$). Consider an extension of this example in which the concentration of the 4th element cannot be measured so instead composition measurements are made in the $N = 3$ space and correlations ${}^3C_{k,l}$ are calculated, and a similar exercise can be performed for $N = 2$.

The values of ${}^N C_{k,l}$ plotted in Fig. 14.8b demonstrate some limitations of the usual statistics. Indeed, the correlation coefficients are skewed towards negative values due to the normalization-induced correlation, as indicated by (14.11). In fact, for the $N = 2$ case, the correlation coefficient is -1 because due to the normalization $x_{i,2} = 1 - x_{i,1}$. In other words, correlation structure of a composition cannot be used to interpret correlations among the measured elemental concentrations and vice versa. It should be mentioned that other distance-based statistics like means, variances and standard deviations, as well as tasks such as clustering and multidimensional scaling have similar limitations.

14.6.3 Sub-Compositional Coherence

An n -part composition (x_1, x_2, \dots, x_n) with $\sum_{i=1}^n x_i = 1$ is called a subcomposition of an m -part composition (x_1, x_2, \dots, x_m) with $\sum_{i=1}^m x_i = 1$, if $m > n$ and (x_1, \dots, x_n) is a subset of the elements (x_1, \dots, x_m) . A consequence of the constant sum constraint for compositional data is that sub-compositions may not reflect the variations present in the parent data, and as a result covariance of elements may change substantially between different subsets of the parent data set. Every composition is a sub- or a parent- composition depending on the objective of an experiment or the goal of data analysis. An experimentalist or a data analyst may not be able to take into account all elements (some elements may not be accessible), or may disregard some of the available elements if they are not pertinent to the objective. The following principle of sub-compositional coherence is an important concept of compositional analysis: any compositional data analysis should be done in a way that produces the same results in a sub-composition, regardless of whether we analyzed only that sub-composition or a parent composition. Subcompositional incoherence of Pearson correlation coefficient is demonstrated in Fig. 14.8, where for a given pair of elements, ${}^N C_{k,l}$ varies with the order N of the composition space.

These effects of closure on statistical analysis of compositional data, induced correlations and subcompositional incoherence, make traditional statistical methods invalid, and artificial correlation obtained by applying such techniques may lead to false scientific discoveries and incorrect predictions. Moreover, methods that are based on a correlation matrix of observations, such as: factor analysis, principal component analysis (PCA), cluster analysis, kriging interpolation to name just a few, would lead to inaccurate, warped results. Thus correlation analysis, and multivariate statistical analysis in general, of compositional data require special techniques in order to avoid producing false results.

14.6.4 Principled Analysis of Compositional Data

The fundamental building block of statistical analysis is the probabilistic model. A well-defined sample space is one of the basic elements in a probabilistic model, and as noted above, the composition sample space is a simplex [2]. All standard statistical methods assume that the sample space is the entire Euclidean space, while compositional data clearly do not satisfy this assumption. In order to deal with the closure effects described in the previous section, an approach based on a family of transformations, the so-called logratio transformations, has been introduced [1]. These transformations based on logarithm of ratios of compositions map the components of the compositions onto a Euclidean space, thus enabling one to apply classical statistical methods. In what follows, we briefly describe a few key concepts of such analysis [2]. The so-called *alr* transform is defined for a given N -element composition \mathbf{x} as an $(N - 1)$ -element vector \mathbf{z} with the following components:

$$\mathbf{z} = \text{alr}(\mathbf{x}) = (\ln(x_1/x_N), \dots, \ln(x_{N-1}/x_N)) \quad (14.12)$$

where one of the composition components is chosen as common divisor. This logratio transform is invertible since there is a one-to-one correspondence between any N -part composition \mathbf{x} and its logratio vector \mathbf{z} . This means that any statement about the components of a composition can be expressed in terms of logratios and vice versa.

By defining the sum

$$s_i = \sum_{j \neq i} \exp(z_j - z_i) = \sum_{j \neq i} x_j/x_i$$

the transformation from logratio to composition coordinates is given by

$$x_i = 1/(s_i + 1). \quad (14.13)$$

Because alr depends on the choice of x_N , this transform is not employed in our calculations and a more suitable transform is discussed below. Later we utilize (14.12) and (14.13) only to illustrate the results of spatial compositional interpolation. To build a vector space structure on the simplex the following operations were introduced by Aitchison. The *closure* operation C is defined as follows

$$\mathbf{x} = C[u_1, \dots, u_N] = (u_1/(u_1 + \dots + u_N), \dots, u_N/(u_1 + \dots + u_N)); \quad u_i \geq 0, \mathbf{x} \in S^N \subset R^{N-1}$$

where u_i represent the raw data such as element quantities. *Perturbation* \oplus is an equivalent of addition in the Euclidean space and defined as

$$\mathbf{w} = \mathbf{x} \oplus \mathbf{y} = C[x_1 y_1, x_2 y_2, \dots, x_N y_N]; \quad \mathbf{w}, \mathbf{x}, \mathbf{y} \in S^N \subset R^{N-1} \quad (14.14)$$

Powering \odot is an equivalent of multiplication a vector by a scalar and defined as

$$\mathbf{w} = a \odot \mathbf{x} = C[x_1^a, x_2^a, \dots, x_N^a]; \quad \mathbf{x} \in S^N, a \in R$$

Aitchison inner product replaces the Euclidean inner product and defined as

$$\langle \mathbf{x}, \mathbf{y} \rangle_A = 1/N \sum_{i=1}^N \sum_{j>i}^N \ln(x_i/x_j) \ln(y_i/y_j); \quad \mathbf{x}, \mathbf{y} \in S^N \subset R^{N-1} \quad (14.15)$$

Thus the norm of a vector, or its simplicial length, is $\|\mathbf{x}\| = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle_A}$. This enables one to compute distances between compositional vectors, projections of compositional vectors, etc.

The Aitchison distance is defined as

$$d_A(\mathbf{x}, \mathbf{y}) = \{1/N \sum_{i=1}^N \sum_{j>i}^N [\ln(x_i/x_j) - \ln(y_i/y_j)]^2\}^{1/2}; \quad \mathbf{x}, \mathbf{y} \in S^N \subset R^{N-1} \quad (14.16)$$

Establishing a metric vector space structure in the simplex and utilizing orthonormal bases facilitates application of complex statistical methods to analysis of composi-

tional data. The so-called isometric logratio (ilr) transform has important conceptual advantages and enables one to use balances, a particular form of ilr coordinates in an orthonormal basis. A balance is defined as

$$b_{pq} = ([pq/(p+q)]^{1/2})\ln(g(x_p)/g(x_q)) \quad (14.17)$$

where $g(\cdot)$ is the geometric mean of the argument, x_p is the group with p parts and x_q is the group of q parts which are obtained by sequential binary partition (see [53] and references there). However, there is no obvious ‘optimal’ basis, and the compositional biplot approach should be used to find one [2]. For an analysis to be subcompositionally coherent, it suffices to define variables using the ratios of the composition values. The quantities x_1/x_2 and $\ln(x_1/x_2)$ are invariant under changes of the composition order as they quantify the relative magnitudes of elemental concentration rather than their absolute values, though the interpretation of the results in terms of the original variables is not always trivial. To study correlation structure of compositions Aitchison introduced a variation matrix $T = \{\tau_{ij}\}$ of dimensions $N \times N$, with the elements

$$\tau_{ij} = \text{var}[\ln(Y_i/Y_j)] \quad (14.18)$$

When τ_{ij} are large, there is no proportionality between the corresponding elements. If, however, the elements i and j are exactly proportional then $\tau_{ij} = 0$. The scale of these variations can be determined by introducing total variance as a normalized sum of the variances of all logratios


$$V_{tot} = 1/(2N) \sum_{i=1}^N \sum_{j=1}^N \tau_{ij} \quad (14.18a)$$

The variation matrix T (14.18, 14.18a) is instrumental in the analysis of associations between elemental concentrations in compositions. Such analysis will be discussed in greater detail in our forthcoming paper dedicated to covariance structures of screening libraries [62]. In what follows we apply balances (14.17) to spatial interpolation of compositional data.

14.6.5 Composition Spread and Distances

The evaluation of a composition spread is crucial in elucidating the composition dependence of materials properties. As a global measure of spread, one can use the metric variance (also known as total variance or generalized variance) that is the average squared distance from the center to the dataset [70, 71]. There are various measures of spread for compositional data and they all are based on the distance defined in (14.16). that is very different from the Euclidean distance. As it is always the case with non-Euclidean geometries, there is more than meets the eye, so to get

(a)



(b)

Compositions	c1	c2	c3	c4	c5	c6	c7
c1	0.000	1.165	2.297	5.488	4.802	5.302	1.931
c2		0.000	1.132	4.323	3.727	5.457	1.837
c3			0.000	3.191	2.754	5.831	2.373
c4				0.000	1.895	7.727	5.078
c5					0.000	5.927	3.786
c6						0.000	3.620
c7							0.000

Fig. 14.9 Distances and straight lines in the composition space (see the text for details). **a** straight lines in a simplex. **b** Log-ratio distance matrix for Aitchison distances d_A

a better feel for simplex’s geometry, let us consider an illustrative example. Since there is a vector space structure within the simplex, one can define geometric elements such as straight lines. Figure 14.9a displays seven compositions and ‘straight’ lines between them. Red square: $c1 = (0.333, 0.333, 0.333)$ —center point; black star: $c2 = (0.446, 0.446, 0.107)$; magenta star: $c3 = (0.485, 0.485, 0.029)$; blue star: $c4 = (0.499, 0.499, 0.001)$; green star: $c5 = (0.091, 0.908, 0.001)$; red star: $c6 = (0.001, 0.972, 0.027)$; cyan star: $c7 = (0.067, 0.837, 0.097)$.

Figure 14.9b shows the log-ratio (see below) distance matrix. It is instructive to note that the largest distance is $d_A(c5, c6) = 5.93$, and the following holds

$$d_A(c4, c5) = 1.90 < d_A(c1, c3) = 2.30 < d_A(c3, c4) = 3.19$$

In real data, zero components, missing values are often present. Moreover, concentrations below instrument detection limit (DL) are routinely encountered in experiments. Usually such nondetects and missing values are erroneously replaced by zeros. Since the statistical analysis of compositional data is based on logratios it cannot be applied to data with zero components.

One approach to this problem is to transform N -element compositional data onto the surface of $(N-1)$ dimensional hypersphere [48], thus bringing the well-developed methods of directional data analysis to compositional data and allowing to deal with zero components. After applying such transforms, the multiscale methods developed for general compact manifolds and for a sphere in particular, will be especially useful for multiscale analysis of functions (like FOMs) defined on compositional space [6, 18, 20, 21, 50, 54–59].

Another recent approach is based on finding that logratio analysis is in fact intimately connected with correspondence analysis (CA) [23–25]. There exists a family of methods parameterized by a power-transformation of the original compositional data: when this power is equal to 1 the ensuing method is exactly CA, and when this power tends to zero the limiting method is exactly LRA. In between we have a continuum of interesting special cases, for example, square root and double square root transformations, but the main point is that these two apparently unrelated and competing methods are really members of a wider common family [24, 25].

14.6.6 Interpolation of Compositional Data: Composition Profiles from Sputtering

In this section we use spatial interpolation of composition measurements, a standard operation in combinatorial research, to demonstrate how behavior of logratio variables differs from that of raw compositional variables. To create a synthetic dataset, we employ a common combinatorial synthesis technique, multi-source co-sputtering of a composition spread thin film. Combinatorial sputtering is commonly used for synthesis of binary, ternary, and quaternary thin film libraries. See for more details [61]

We will assume that accurate, noise-free measurements of the 4 compositional variables are made on a set of 25 substrate positions chosen as a 5x5 square grid with 25 mm spacing. An appropriate spatial interpolation method should at least guarantee that the non-negativity and constant-sum constraints are satisfied. In fact, among the conventional unconstrained interpolation techniques, linear interpolation satisfies these requirements. However, usual straightforward approaches, even if they satisfy the constraints, interpolate each component x_i independently, thus ignoring the inner relationships between the compositional elements. Since our end goal is to enable

the analysis of covariance structure of compositions without the artifacts of induced correlation and sub-compositional incoherence, an approach that leads to accurate logratio values employed by the simplicial distance (14.16) and compositional covariance matrix T (14.18) is required. To achieve this, we utilize a broadly-applicable and highly versatile technique based on kriging. The kriging-based interpolation was computed with **R** language and environment for statistical computing by applying the R-package “compositions” [65] van der Boogaart and Tolosana-Delgado 2013). The method exploits codependences in the composition and takes into account the spatial covariance structure by modeling the set of variograms for all possible pairwise balances (14.17). It takes into account various effects and parameters including the nugget effect, the choice of exponential and spherical variograms which parameters we chose to be 62.5 and 162.0 respectively. Since this interpolation technique is specialized for compositional data, we refer to it as “compositional interpolation” and represent the result of compositional interpolation of the 25 z_i measurements as $\text{CompInterp}(z_i)$. To attain an analogous result using traditional linear interpolation, x_i and x_N can be independently interpolated followed by calculating z_i (14.12), resulting in a spatial map of z_i referred to as $\text{LinInterp}(z_i)$. The results of compositional and linear interpolations and their comparisons with the “perfect” data calculated from the model compositions are shown in Fig. 14.10. By definition, both interpolation methods produce exact values at each of the 25 locations in the sampling grid. The assessment of the performance of a given interpolation is thus performed by evaluat-

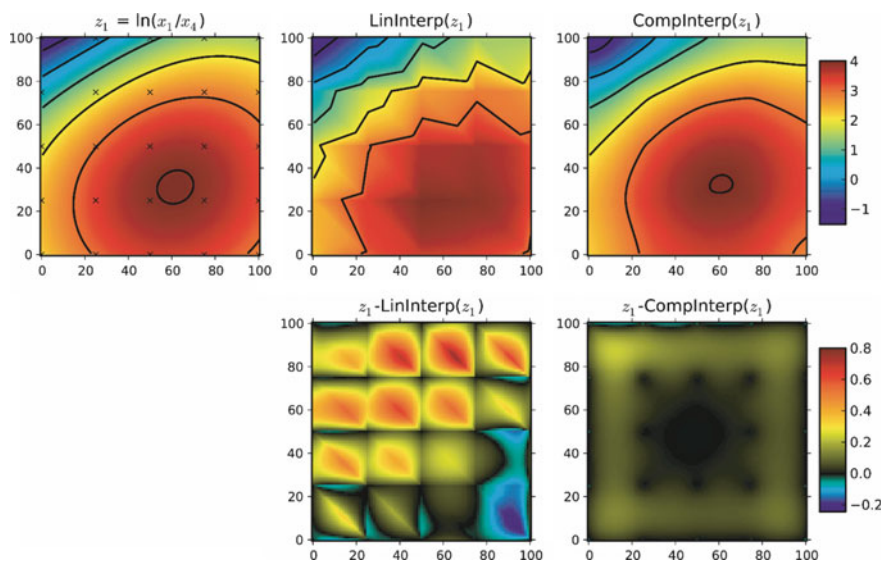


Fig. 14.10 Interpolation of a 4-element composition. z_1 : logratio of compositions x_1 and x_4 with the 25 sampling points marked by “x”; $\text{LinInterp}(z_1)$: logratio of the linearly interpolation of x_1 and x_4 ; $\text{CompInterp}(z_1)$: logratio of the compositionally interpolated z_1 ; the difference between the model data and its linear interpolation; the difference between the model data and its compositional interpolation [61]

ing the absolute magnitude and pattern of interpolation error in the regions between the sampling points (the grid). Compared to linear interpolation, the compositional interpolation provides more accurate results and the discrepancy varies smoothly over the entire interpolation region. It is important to note that artificial ‘patchiness’ of the linear interpolation would distort simplicial distances (14.16) between different compositions. Such distortions would lead to artificial associations in the analysis of correlational structure of compositions and, more generally, to erroneous results in all calculations that involve distances, e.g. mean and standard deviation.

Kriging assumes that the observed values are a realization of a stochastic process, so the quantitative advantages of compositional interpolation based on kriging should become more pronounced as variation of the composition variables increases. It is worth noting, that there are other interpolation methods that preserve the non-negativity and constant-sum constraints such as local sample mean, inverse distance interpolation, and triangulation (since the weights they use range from 0 to 1, and sum to unity). However, unlike the approach utilized here (Tolosana-Delgado and van den Boogaart, [53, 71], those methods do not take into account the spatial covariance structure which may be critical for statistical analysis. As combinatorial materials science continues to expand into high order compositions spaces, the prudent application of statistical methods developed specifically for CDA will be required to enable accurate data mining.

14.7 Summary and Conclusions

A central premise of high throughput combinatorial science is that systematic measurements of material libraries can reveal relationships among material composition, structure, performance and other properties. To facilitate accurate and effective extraction of information from large, complex data sets created by high throughput experiments, the materials scientists must engage in close interdisciplinary communication/collaboration with researchers in other disciplines such as statistics, computer science, applied mathematics, and artificial intelligence. One step in fostering a close interdisciplinary communication/collaboration between materials scientists and researchers in other disciplines would be to follow the example of top journals in other fields and have statisticians as members of the Editorial Boards of journals concerned with high throughput and combinatorial materials science. This would ensure that the adequacy of the statistical analysis used in papers is properly evaluated and, more importantly, will enable journals to formulate statistics guidelines for contributors [3, 30, 67]. Moreover, it is important to emphasize that statisticians should not only be consulted after data had already been generated, but rather should be involved in the design of experiments. It is only through prudent incorporation of informatics in high throughput workflows that combinatorial materials science can be fully realized. This chapter introduced high throughput experimental pipelines and example data to illustrate 2 areas of informatics that are central to combinatorial materials science. The high throughput strategy of tiered screening and resulting

complex datasets demonstrate the need for new statistical techniques that enable the generation of information rich databases and provide accurate assessment of composition-property relationships. While further research is required to assimilate and advance these informatics techniques, foundational work in these research areas is presented.

Acknowledgments The authors would like to thank Prof. Alfred Ludwig for stimulating discussions. This work is performed by the Joint Center for Artificial Photosynthesis, a DOE Energy Innovation Hub, supported through the Office of Science of the U.S. Department of Energy under Award Number DE-SC000499.

References

1. J. Aitchison, The statistical analysis of compositional data (with discussion). *J. R. Stat. Soc. Ser. B Stat. Methodol.* **44**, 139–177 (1982)
2. J. Aitchison, *The Statistical Analysis of Compositional Data*. Monographs on Statistics and Applied Probability (Chapman & Hall Ltd., London, 1986) (2d edition with additional materials, The Blackburn Press, 2003)
3. D. Altman et al., Statistical guidelines for contributors to medical journals. *BMJ* **286**, 1489–1493 (1983)
4. P.V. Balachandran, S.R. Broderick, K. Rajan, Identifying the inorganic gene for high-temperature piezoelectric perovskites through statistical learning. *Proc. R. Soc. Math. Phys. Eng. Sci.* **467**, 2271–2290 (2011). doi:[10.1098/rspa.2010.0543](https://doi.org/10.1098/rspa.2010.0543)
5. S. Bandyopadhyay, U. Maulik, Nonparametric genetic clustering: comparison of validity indices. *IEEE Trans. Syst. Man Cybern. Part C (Applications and Reviews)* **31**, 120–125 (2001). doi:[10.1109/5326.923275](https://doi.org/10.1109/5326.923275)
6. S. Bernstein, I. Pesenson, Crystallographic and Geodesic Radon Transforms on $SO(3)$: motivation, generalization, discretization. in *Geometric Analysis and Integral Geometry, Contemporary Mathematics*, vol 598 (2013) (a volume dedicated to 85th birthday of S. Helgason)
7. M. Borenstein, L. Hedges, J. Higgins, H. Rothstein, *Introduction to Meta-Analysis* (Wiley, New York, 2009)
8. N. Boric, P.A. Estévez, Genetic Programming-based Clustering Using an Information Theoretic Fitness Measure. in *Proceedings of the IEEE Congress on Evolutionary Computation (CEC 2007)*, pp. 31–38 (2007)
9. S.R. Broderick, K. Rajan, Eigenvalue decomposition of spectral features in density of states curves. *EPL (Europhysics Letters)* **95**, 57005 (2011). doi:[10.1209/0295-5075/95/57005](https://doi.org/10.1209/0295-5075/95/57005)
10. P.J.S. Buenconsejo, A. Ludwig, Composition-structure-function diagrams of Ti-Ni-Au thin film shape memory alloys. *ACS Comb. Science* **16**, 678–685 (2014)
11. C.M. Caskey, R.M. Richards, D.S. Ginley, A. Zakutayev, Thin film synthesis and properties of copper nitride, a metastable semiconductor. *Mater. Horiz.* **1**, 424 (2014). doi:[10.1039/c4mh00049h](https://doi.org/10.1039/c4mh00049h)
12. J.N. Cawse, *Experimental Design for Combinatorial and High Throughput Materials Development* (Wiley, New York, 2002)
13. T. Chikyow, P. Ahmet, K. Nakajima, T. Koida, M. Takakura, M. Yoshimoto, H. Koinuma, A combinatorial approach in oxide/semiconductor interface research for future electronic devices. *Appl. Surf. Sci.* **189**, 284–291 (2002). doi:[10.1016/S0169-4332\(01\)01004-2](https://doi.org/10.1016/S0169-4332(01)01004-2)
14. Committee on the Analysis of Massive Data, *Frontiers in Massive Data Analysis* (The National Academies Press, Washington, 2013)
15. G. Cummings, *Understanding the New Statistics. Effect Sizes, Confidence Intervals, and Meta-Analysis* (Routledge, London, 2012)

16. L. Cwiklik, B. Jagoda-Cwiklik, M. Frankowicz, Influence of the spacing between metal particles on the kinetics of reaction with spillover on the supported metal catalyst. *Appl. Surf. Sci.* **252**(3), 778–783 (2005). doi:[10.1016/j.apsusc.2005.02.107](https://doi.org/10.1016/j.apsusc.2005.02.107)
17. Data-Enabled Science in the mathematical and Physical Sciences, A workshop funded by the national Science Foundation (2010), <https://www.nsf.gov/mps/dms/documents/Data-EnabledScience.pdf>
18. C. Durastanti, Y. Fantaye, F. Hansen, D. Marinucci, I. Pesenson, A simple proposal for radial 3D needles. *Phys. Rev. D.* (Accepted) (2015)
19. J. Fan, F. Han, H. Liu, Challenges in big data. *Natl. Sci. Rev.* **1**, 1–22 (2014)
20. D. Geller, D. Marinucci, Spin wavelets on the sphere. *J. Fourier Anal. Appl.* **16**, 840–884 (2010)
21. D. Geller, I. Pesenson, Bandlimited localized Parseval frames and Besov spaces on compact homogeneous manifolds. *J. Geom. Anal.* **21**(2), 334–371 (2011)
22. E. Gokcay, J.C. Principe, Information theoretic clustering. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**, 158–171 (2002). doi:[10.1109/34.982897](https://doi.org/10.1109/34.982897)
23. M.J. Greenacre, *Correspondence Analysis in Practice* (Chapman & Hall, London, 2007)
24. M.J. Greenacre, Log-ratio analysis is a limiting case of correspondence analysis. *Math. Geosci.* **42**, 129–134 (2010)
25. M.J. Greenacre, Measuring subcompositional incoherence. *Math. Geosci.* **43**, 681–693 (2011)
26. J. Gregoire, J. Haber, S. Mitrovic, C. Xiang, S. Suram, P. Newhouse, E. Soedarmadji, M. Marcin, K. Kan, D. Guevarra, Enabling solar fuels technology with high throughput experimentation, in Paper presented at the MRS Proceedings (2014)
27. J.M. Gregoire, D. Dale, A. Kazimirov, F.J. DiSalvo, R.B. van Dover, High energy x-ray diffraction/x-ray fluorescence spectroscopy for high-throughput analysis of composition spread thin films. *Rev. Sci. Instrum.* **80**, 123905 (2009). doi:[10.1063/1.3274179](https://doi.org/10.1063/1.3274179)
28. J.M. Gregoire, D.G. Van Campen, C.E. Miller, R. Jones, A.M. Suram SK, High throughput synchrotron X-ray diffraction for combinatorial phase mapping. *J. Synchrotron Radiat.* **21**(6), 1262–1268 (2014)
29. J.M. Gregoire, C.X. Xiang, X.N. Liu, M. Marcin, J. Jin, Scanning droplet cell for high throughput electrochemical and photoelectrochemical measurements. *Rev. Sci. Instrum.* **84**(2) (2013). doi:[10.1063/1.4790419](https://doi.org/10.1063/1.4790419)
30. Guidelines for Using Confidence Intervals for Public Health Assessment, Washington State Department of Health (2012)
31. J.A. Haber, Y. Cai, S. Jung, C. Xiang, S. Mitrovic, J. Jin, A.T. Bell, J.M. Gregoire, Discovering Ce-rich oxygen evolution catalysts, from high throughput screening to water electrolysis. *Energy Environ. Sci.* **7**(2), 682 (2014a). doi:[10.1039/c3ee43683g](https://doi.org/10.1039/c3ee43683g)
32. J.A. Haber, D. Guevarra, S. Jung, J. Jin, J.M. Gregoire, Discovery of new Oxygen evolution reaction electrocatalysts by combinatorial investigation of the Ni–La–Co–Ce Oxide composition space. *ChemElectroChem* 1613–1617 (2014). doi:[10.1002/celec.201402149](https://doi.org/10.1002/celec.201402149)
33. A. Shinde, R.J. Jones, D. Guevarra, S. Mitrovic, N. Becerra-Stasiewicz, J.A. Haber, J. Jin, J.M. Gregoire, High-throughput screening for acid-stable oxygen evolution electrocatalysts in the (*Mn* – *Co* – *Ta* – *Sb*) Ox composition space, *Electrocatalysis* **6**(2), 229–236 (2015)
34. J.A. Haber, C. Xiang, D. Guevarra, S. Jung, J. Jin, J.M. Gregoire, High throughput mapping of electrochemical properties of (Ni-Fe-Co-Ce)Ox Oxygen evolution catalysts. *Chem. Electro. Chem.* **1**(3), 524–528 (2014)
35. J.R. Hatrick-Simpers, W.S. Hurst, S.S. Srinivasan, J.E. Maslar, Optical cell for combinatorial in situ Raman spectroscopic measurements of hydrogen storage materials at high pressures and temperatures. *Rev. Sci. Instrum.* **82**, 033103 (2011). doi:[10.1063/1.3558693](https://doi.org/10.1063/1.3558693)
36. E. Jaynes, Information theory and statistical mechanics. *Phys. Rev.* **106**, 620–630 (1957). doi:[10.1103/PhysRev.106.620](https://doi.org/10.1103/PhysRev.106.620)
37. R. Jenssen, D. Erdogmus, K. Hild, J.C. Principe, T. Eltoft, Optimizing the Cauchy-Schwarz PDF distance for information theoretic, non-parametric clustering, in *Int'l Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, pp. 34–35 (2005)
38. R. Jenssen, J.C. Principe, D. Erdogmus, T. Eltoft, The cauchy-schwarz divergence and parzen windowing: connections to graph theory and mercer kernels. *J. Franklin Inst.* **343**, 614–629 (2006). doi:[10.1016/j.jfranklin.2006.03.018](https://doi.org/10.1016/j.jfranklin.2006.03.018)

39. R.J. Jones, D. Guevarra, A.S. Shinde, C. Xiang, J.A. Haber, J. Jin, J.M. Gregoire, Parallel electrochemical treatment system. *ACS Comb. Sci.* **17**(2), 71–75 (2015)
40. D. Kan, C.J. Long, C. Steinmetz, S.E. Lofland, I. Takeuchi, Combinatorial search of structural transitions: systematic investigation of morphotropic phase boundaries in chemically substituted BiFeO₃. *J. Mater. Res.* **27**, 2691–2704 (2012). doi:[10.1557/jmr.2012.314](https://doi.org/10.1557/jmr.2012.314)
41. C.S. Kong, W. Luo, S. Arapan, P. Villars, S. Iwata, R. Ahuja, K. Rajan, Information-theoretic approach for the discovery of design rules for crystal chemistry. *J. Chem. Inf. Model.* **52**, 1812–1820 (2012). doi:[10.1021/ci200628z](https://doi.org/10.1021/ci200628z)
42. J. Kruschke, Bayesian estimation supersedes the t-test. *J. Exp. Psychol. Gen.* (2012)
43. J. Kruschke, *Doing Bayesian Data Analysis*, 2nd edn. (Academic Press, Waltham, 2014)
44. A.G. Kusne, T. Gao, A. Mehta, L. Ke, M.C. Nguyen, K.-M. Ho, V. Antropov, C.-Z. Wang, M.J. Kramer, C. Long, I. Takeuchi, On-the-fly machine-learning for high-throughput experiments: search for rare-earth-free permanent magnets. *Sci. Rep.* **4**, 6367 (2014). doi:[10.1038/srep06367](https://doi.org/10.1038/srep06367)
45. R. Lebras, T. Damoulas, J.M. Gregoire, A. Sabharwal, C.P. Gomes, R.B. Dover, Constraint Reasoning and Kernel Clustering for Pattern Decomposition With Scaling, in *Proceedings of the 17th international conference on Principles and practice of constraint programming*, pp. 508–522 (2011)
46. J. Leek, R. Scharpf, H. Bravo, Tackling the widespread and critical impact of batch effects in high-throughput data. *Nat. Rev.* **1**, 733–739 (2010)
47. H. Li, Y. Liang, Q. Xu, Support vector machines and its applications in chemistry. *Chemom. Intell. Lab. Syst.* **95**, 188–198 (2009). doi:[10.1016/j.chemolab.2008.10.007](https://doi.org/10.1016/j.chemolab.2008.10.007)
48. K.V. Mardia, P.E. Jupp, *Directional Statistics*, 2nd edn. (Wiley, New York, 2000), p. 160
49. W.F. Maier, K. Stowe, S. Sieg, Combinatorial and high-throughput materials science. *Angew. Chem. Int. Ed.* **46**, 6016–6067 (2007)
50. D. Marinucci, G. Peccati, *Random Fields on the Sphere*. London Mathematical Society Lecture Note Series (2011)
51. S. Mitrovic, E. Soedarmadji, P.F. Newhouse, S. Suram, J.A. Haber, J. Jin, J.M. Gregoire, Colorimetric screening for high-throughput discovery of light absorbers. *ACS Comb. Sci.*
52. D.P. Muni, N.R. Pal, J. Das, A novel approach to design classifiers using genetic programming. *IEEE Trans. Evol. Comput.* **8**, 183–196 (2004). doi:[10.1109/TEVC.2004.825567](https://doi.org/10.1109/TEVC.2004.825567)
53. V. Pawlowsky-Glahn, A. Buccianti (eds.), *Compositional Data Analysis: Theory and Applications* (Wiley, New York, 2011)
54. I. Pesenson, Sampling of Paley-Wiener functions on stratified groups. *J. Fourier Anal. Appl.* **4**(3), 271–281 (1998)
55. I. Pesenson, Paley-wiener approximations and multiscale approximations in sobolev and besov spaces on manifolds. *J. Geom. Anal.* **19**(2), 390–419 (2009)
56. I. Pesenson, A sampling theorem on homogeneous manifolds. *Trans. Am. Math. Soc.* **352**(9), 4257–4269 (2000)
57. I. Pesenson, Springer Handbook of Geomathematics, *Splines and Wavelets on Geophysically Relevant Manifolds* (Springer, Berlin, 2015), pp. 1–32
58. I. Pesenson, Multiresolution Analysis on Compact Riemannian Manifolds, in *Multiscale Analysis and Nonlinear Dynamics: From Genes to the Brain*, ed. by M. Pesenson (Wiley-VCH, Weinheim, 2013), pp. 65–82
59. M.Z. Pesenson, I.Z. Pesenson, Adaptive multiresolution analysis based on synchronization. *Phys. Rev. E* **84**, 045202(R) (2011)
60. M.Z. Pesenson, Multiscale Analysis—Modeling, Data, Networks, and Nonlinear Dynamics, in *Multiscale Analysis and Nonlinear Dynamics*, *Wiley Reviews of Nonlinear Dynamics and Complexity*, ed. by M.Z. Pesenson (Wiley-VCH, Weinheim, 2013), pp. 1–19
61. M.Z. Pesenson, S. Suram, J.M. Gregoire, Statistical analysis and interpolation of compositional data in materials science. *ACS Comb. Sci.* **17** (2), 130–136 (2015)
62. M.Z. Pesenson, S. Suram, J. Haber, D. Guevara, P. Newhouse, E. Soedarmadji, J.M. Gregoire, Correlation Structure of High Throughput Composition Screening Libraries (in preparation) (2015)

63. R. Potyrailo, V.M. Mirsky, *Combinatorial Methods for Chemical and Biological Sensors* (Springer Science & Business Media, Berlin, 2009), p. 125
64. J. Principe, D. Xu, J. Fisher, *Information theoretic learning. Unsupervised adaptive filtering*, vol. 1 (Wiley, New York, 2000)
65. R Development Core Team. R: A Language and Environment for Statistical Computing; R Foundation for Statistical Computing: Vienna, Austria (2004)
66. K. Rajan, *Combinatorial Materials Sciences: Experimental Strategies for Accelerated Knowledge Discovery*. *Ann. Rev. Mater. Res.* **38**, 299–322 (2008)
67. H. Roediger, What's New at Psychological Science. An Interview with Editor in Chief (2013). <http://www.psychologicalscience.org/index.php/publications/observer/2013/november-13/whats-new-at-psychological-science.html>
68. X. Shi, J. Luo, N.P. Njoki, Y. Lin, T.-H. Lin, D. Mott, S. Lu, C.-J. Zhong, *Combinatorial Assessment of the Activity-Composition Correlation for Several Alloy Nanoparticle Catalysts*. *Ind. Eng. Chem. Res.* **47**, 4675–4682 (2008). doi:10.1021/ie800308h
69. S.K. Suram, J.A. Haber, J. Jin, J. Gregoire, *Generating information rich high-throughput experimental materials genomes using functional clustering via multi-tree genetic programming and information theory*. *ACS Comb. Sci.* **17** (4), 224–233 (2015)
70. R. Tolosana-Delgado, K. van den Boogaart, V. Pawlowsky-Glahn, *Geostatistics for Compositions*, in *Compositional Data Analysis: Theory and Applications*, eds. by V. Pawlowsky-Glahn, A. Buccianti (Wiley, Chichester, 2011), pp 73-86
71. K. van den Boogaart, R. Tolosana-Delgado, *Analyzing Compositional Data with R*, Use R! Series (Springer, Berlin, 2013)
72. D.W. van der Merwe, A.P. Engelbrecht, *Data clustering using particle swarm optimization*. *2003 Congr. Evol. Comput.* **1**, 215–220 (2003). doi:10.1109/CEC.2003.1299577
73. R. Wilcox, *Fundamentals of Modern Statistical Methods. Substantially Improving Power and Accuracy*, vol. 2 (Springer, New York, 2010)
74. C. Xiang, J. Haber, M. Marcin, S. Mitrovic, J. Jin, J.M. Gregoire, *Mapping quantum yield for (Fe-Zn-Sn-Ti)Ox photoabsorbers using a high throughput photoelectrochemical screening system*. *ACS Comb. Sci.* **16**(3), 120–127 (2014a). doi:10.1021/co400081w
75. C. Xiang, S.K. Suram, J.A. Haber, D.W. Guevarra, J. Jin, J.M. Gregoire, *A high throughput bubble screening method for combinatorial discovery of electrocatalysts for water splitting*. *ACS Comb. Sci.* **16**(2), 47–52 (2014b)
76. R. Zarnetta, P.J.S. Buenconsejo, A. Savan, S. Thienhaus, A. Ludwig, *High-throughput study of martensitic transformations in the complete TieNieCu system*. *Intermetallics* **26**, 98e109 (2012)

Index

A

Ab initio, 10, 69
Ab initio data, 187, 210
Absolute potts model, 119
Adaptive experimental design, 7
Additive manufacturing (AM), 141
AlSi10Mg, 143
Amorphous systems, 116
ANalysis Of VAriance (ANOVA), 143
Annealing parameter, 112
Approximate knowledge gradient (AKG), 67
Arrhenius plots, 182
Asymptotic representation, 85
Attribute-property relationships, 276
Automation of measurements, 244

B

Band structure, 11
Barsoum, 191
Bayes classifier, 78, 84
Bayes error, 78, 79, 86
Bayes' rule, 15, 88
Bayes-optimality, 45
Bayesian, 48, 77
Bayesian error estimate, 89
Bayesian experimental design, 21, 33, 70
Bayesian global optimization, 46
Bayesian inference, 14
Bayesian MMSE error estimation, 98
Bayesian optimization, 45–47, 56, 60, 61, 68–70, 175
Bayesian statistics, 46
Bernoulli, 68
Binary classification, 78

Binary oxide, 82
Bioinformatics, 103
Block structure, 108
BOD, 200
Body-centered cubic (bcc), 31
Bolstered resubstitution (bol), 89
Bonding, 190
Bonding-anti-bonding transition, 235
Bonding-antibonding transition, 224
Bootstrap, 6
Bootstrapping and ensemble averaging, 167
Breathing distortion, 220
Buckingham Pi theorem, 24
Bulk, 248
Bulk metallic glasses (BMGs), 209
Bulk modulus, 190, 223

C

Ca-Si-Al-hydrates, 209
Cahn-Hilliard equation, 22
Calcium silicate hydrate (CSH), 206
Canonical finite temperature phase diagram, 125
Capacitance, 259
Catalysis, 256
Catalytic, 260
Cathode materials, 175
Cauchy-Born, 191
Cauchy-Born criterion, 192
Cauchy-Schwarz divergence, 280
Cauchy-Schwarz divergence function, 281
CDA, 287
CDF transformation, 27
Cerium oxides, 158
CFS method, 146, 147

- Charge disproportionation, 216
Chemistry-modulus relationship, 235
Chemoinformatics, 174
China, 243
Cholesky decomposition, 61
Clarke formula, 201
Class-conditional densities, 86, 91
Class-conditional distribution, 79, 87
Classification, 5, 78
Classification and clustering problems, 112
Classification rule, 80
Classifier, 77
Closure operation, 291
Clustering, 113
Clustering techniques, 279
Combinatorial materials science, 241, 273, 274, 286
Combinatorial optimization, 126
Combinatorial powder synthesis, 249
Community detection, 115, 118, 119, 125
Complex materials, 128
Complex multi-component systems, 254
Complex system, 115, 116
Composite ceramics, 254
Composition, 265
Composition clusters, 286
Composition library, 281
Composition-(micro)structure-experimental parameters-property relationships, 279
Composition-property function, 283
Composition-structure-property correlation maps, 288
Compositional covariance matrix, 295
Compositional interpolation, 296
Computational genomics, 188
Computational screening, 158
Computational susceptibility, 126
Conjugate priors, 92
Consistent, 80
Correlation, 233
Correlation coefficients, 198
Correlation structure, 289
Correlation-based Feature Selection (CFS), 147, 148
Correlations between variables, 164
Correspondence analysis (CA), 273, 294
Cost function, 118
Countries and institutions, 247
Covariance function, 48, 49
Covariance kernel, 48, 60
Covariance matrix, 225
Covariance structures of screening libraries, 292
Covariances, 80
(Cr₂Hf)₂Al₃C₃, 206
Cross-validation, 87, 89, 227, 234
Cross-validation error estimator, 82
Cubic perovskites, 82
Cumulative distribution function, 62
- D**
3D, 265
Data, 263
Data analysis, 168, 275
Data base, 261
Data clustering algorithm, 112
Data management, 275
Data mining, 188, 198, 260
Data mining and machine learning, 273
Data point, 263
Data point glyphs, 265
Data-driven approaches, 158
Data-driven machine learning, 69
Data-fit models, 19
Decision boundary, 96
Decision theory, 46
Decision tree, 4, 148
Decomposed the DOS spectra, 232
Deconvolution, 232
D-electron valency, 229
Delithated, 175
Density function, 88
Density functional theory (DFT), 185, 190, 223
Density of states (DOS), 195, 223
Descriptor, 164, 189, 191, 197, 200, 210, 214
Design cost, 80
Design of experiments, 275
Dielectrics, 251
Dilute magnetic semiconductor, 253
Dimensionality reduction, 165
Directional data analysis, 294
Dirichlet priors, 98
Discrete histogram rule, 84
Distance and density based clustering approaches, 279
Distance matrix, 103
Distortion-mode decomposition analysis, 214
Distribution-free, 81, 84
DOS spectra, 230
Double asymptotic representations, 85
Down-selection methods, 277

Dynamical heterogeneities, 132

E

Eagar-Tsai (E-T) model, 145, 146, 148, 150

Edge, 118

Effective class-conditional density, 91

Efficient Global Optimization (EGO), 9, 46, 64

Electrocatalyst screening, 277

Electrochemical, 260

Electronic bandwidth-controlled, 217

Electronic nose, 260

Electronic structure, 190

Electronic structure-crystal structure-property relationships, 237

Empirical Bayes, 57

Empirical risk minimization, 81

Energy landscape, 117

Energy-storage, 176

Ensemble, 117

Error estimation, 7, 78

Estimated error, 84

Estimation rule, 82

Estimator, 18

E-T, 146, 149–151

Euclidean distance, 50

Euclidean distance matrix, 104

Euclidean metric based Gaussian kernels, 286

Euclidean-based correlation structure, 287

Evanescence microwave probe (EMP), 258

Evidence, 18

Evolutionary operations of genetic programming, 279

Expected improvement, 62, 68

Expected improvement method, 69

Expected information gain, 17

Expected KL divergence, 33

Expected utility, 14

Expected value of sample information (EVSI), 68

Experimental design, 14

Experimental noise, 56

Experimental uncertainty, 35

Exploration and exploitation, 7

Exploration vs. exploitation tradeoff, 63

Eyre's method, 22

F

Face-centered cubic (fcc), 31

Feature, 146

Feature-label distribution, 78, 86, 88, 98

Feature reduction, 113

Features, 9, 78, 214

Feature selection, 160, 164

Features or fingerprint, 164

Features Space, 104, 107

Feature vector, 78

Feedback, 261

Fermi level (E_F), 195

Ferroelectric, 113, 251

Ferromagnetic, 121

First principles, 168

First principles methods, 158

First-order phase transition, 112

Fisher, R.A., 97

Fourier weights, 124

Frequentist, 26, 47

Full factorial experimental design, 143

Full fractional design, 146

Functional polymers, 10

G

Gas sensing, 259

Gaussian kernel, 49

Gaussian model, 79, 89, 91, 94, 98

Gaussian processes, 181

Gaussian process regression, 45–48, 56, 58–60, 69

Gaussian random field, 21

General character and scale, 132

Genetic, 118

Genetic operators, 285

Genetic trees, 282

Genomics, 8

Glasses, 116

Glyph, 263

Gradient boosting, 4

Gradient descent, 58

Graphic, 263

Greedy algorithm, 120

Guided screening strategy, 164

Guidelines, 159

H

Hamiltonian, 119, 121, 125, 126

Heat of formation (HoF), 191, 192

Heterophase interfaces, 29, 30

Heteroscedastic noise, 56

Hierarchical models, 19

Hierarchical structure, 121

Hierarchical surrogates, 38

High-dimensional hyperplane, 227

High dimensional transformation, 166

High entropy alloys, 209
 High order composition spaces, 273, 296
 High-throughput, 4
 High throughput (HiTp), 273
 High-throughput analysis, 242
 High throughput combinatorial science, 296
 High throughput experimental pipeline, 272, 275, 296
 High-throughput measurement and analysis, 254
 High-throughput screening, 174
 High throughput synthesis and evaluation, 272
 High-throughput topics, 246
 High-throughput workflow, 279
 HiTp experimentation, 277
 HiTp materials characterization, 276
 Holdout estimate, 82
 Homoscedastic, 55
 Homoscedasticity, 55
 Human Genome Project, 187
 Hungarian algorithm, 106
 Hybrid approach, 234
 Hybrid informatics approach, 237
 Hydrogen storage, 259
 Hydrothermal, 253
 Hyperparameters, 52, 57, 87, 96
 Hysteresis loop, 258, 264

I

Ideal perovskite, 113
 Image, 263
 Image segmentation, 128
 Impedance screening, 260
 Improvement, 62
 In silico, 69
 Individual regression models, 167
 Induced correlations, 290
 Industry, 244
 Inference, 51
 Informatics, 224
 Informatics tools, 272
 Information gain, 18, 27
 Information science, 4
 Information theoretic, 117
 Information-theoretic approach, 280
 Information theoretic objectives, 37
 Information theory, 118, 125, 127
 Information theory correlations, 116
 Initialization, mutation, selection, crossover and termination, 282
 Ink-jet printing, 250

Inverse design, 223
 Inverse-Wishart distribution, 93
 Ionic conductors, 251
 Irreducible representation, 215
 Island scanning, 152

J

Jahn-Teller distortions, 220
 Japan, 243
 Jeffreys prior, 27

K

Kernel width, 280
 KG factor, 66
 Knowledge-gradient, 9, 60, 65, 68
 Knowledge-gradient algorithm, 65, 66
 Knowledge-gradient (KG) factor, 66
 Knowledge-gradient method, 69
 Kob-Andersen binary system, 130
 Kriging, 47, 181, 296
 Kullback-Leibler (KL) divergence, 17

L

Labels, 78
 Landau theory, 215
 Laser powder-bed fusion, 141
 Lattice parameters, 32
 Lattice thermal conductivity, 201
 Layered metal composites, 29
 Learning, 175
 Leave-one-out cross validation, 58
 Leave-one-out error, 84
 Leave-one-out estimator, 82, 85
 Lennard-Jones, 129
 Likelihood function, 34
 Lindemann rule, 178
 Linear data-dimensionality reduction, 218
 Linear discriminant analysis (LDA), 80, 81
 Linear regression models, 227
 Lithium ion batteries, 175
 Lithium Super Ionic CONductors (LISICON), 182
 Loadings and scores matrices, 225
 Loadings plot, 229
 Local minimum, 120
 Logratio analysis, 294
 Logratio analysis to interpolation, 288
 Logratio transforms, 288
 Low O vacancy formation energy, 161
 Low valence state, 161
 316L stainless steel, 142, 144, 145, 154

M

Machine learning, 4, 116, 164, 178, 188
 Machine learning algorithms, 112
 Magnetic force microscopy (MFM), 256
 Magnetoelectric, 254
 Many-body, 130
 Marginal likelihood, 16, 34, 58
 Markov decision process, 68
 Markov-chain-Monte-Carlo (MCMC), 98
 Matérn covariance kernel, 50
 Materials design, 45
 Materials Genome, 4
 Materials genome initiative, 116
 Materials informatics, 4, 103
 Materials Space, 104, 107
 MAX database, 197
 MAX phases, 187, 189
 MAX solid solutions, 205
 Maxenes, 191
 Maximum membership class, 284
 MAX ($M_{n+1}AX_n$) phase, 188
 Mean function, 51
 Mean-square error (MSE), 82, 147–149
 Mechanical properties, 190
 Mechanistic insights, 214
 Merit, 146, 147
 Metal to insulator transitions, 216
 Metallic glass, 130
 Metals, oxides and ceramics, 247
 Metric vector space structure, 291
 Minimization Procedure, 106
 Minimum thermal conductivity κ_{min} , 203
 Minimum-mean-square-error (MMSE), 77
 Mining and extraction, 164
 Misclassification, 78
 Misfit, 31
 Mixed moments, 85
 MMSE estimate, 89
 Model-based optimal experimental design, 14
 Model discrimination, 17, 29
 Model selection, 14, 33
 Model space, 18
 Moiré pattern, 32
 Molecular beam epitaxy (MBE), 248
 Monte Carlo, 35
 Monte Carlo optimization methods, 112
 MT-GP algorithm, 284, 285
 MT-GP approach, 282
 Multi scale community detection, 116
 Multi-scale “inverse design”, 235
 Multi-scale community detection, 129
 Multi-tree genetic programming, 280

Multinomial discrimination, 85
 Multinomial model, 85
 Multiresolution analysis (MRA), 274
 Multivariate model, 85
 Multivariate normal, 48
 Multivariate normal random, 51
 Multivariate student's t -distribution, 93, 95
 Multi-resolution analysis, 130
 Mutual information, 17
 MXenes, 205

N

National institute of standards and technology (NIST), 244
 $N(E_F)$, 195
 Neighborhood, 105
 Neighborhood ordering, 107, 109
 Newton's method, 58
 Ni_xNb_{1-x} , 209
 Node, 118, 120
 Non-Euclidean, 272
 Non-Euclidean space, 274
 Normal-inverse-Wishart distribution, 93
 Normal-inverse-Wishart prior, 89, 96
 Normalized mutual information (NMI), 122
 NP hard, 126
 Numerical quadrature, 35

O

Objective cost of uncertainty, 98
 Objective function, 60, 281
 Occam's razor, 16
 OER catalyst activity, 276
 OLCAO, 192
 Open Quantum Materials Database, 173
 Optical, 256
 Optimal Bayesian classifier, 90, 94
 Optimal Bayesian experimental design, 37
 Optimal classifiers, 77
 Optimal constrained classifier, 81
 Optimal experimental design, 14
 Optimal learning, 70
 Optimal membership set, 284
 Optimal minimum-mean-square-error (MMSE), 87
 Optimistic bias, 84
 Optimization, 8, 11, 19, 181
 Order parameter, 215
 Order parameter field, 22
 OS supersedes, 167
 OS, IR and EA, 165
 Over-fitting, 230

Overfitting, 81
 Oxidation state, ionic radius and electron affinity, 168

P

Pair density correlations, 132
 Parallel tempering, 112
 Parameter inference, 14, 17
 Partial and infinite swapping, 112
 Partial least squares, 225
 Pattern recognition model, 84
 PBOD, 191
 PCA analysis, 232
 PCA on the DOS spectra, 228
 Pearson correlation coefficient, 289
 Permittivity, 264
 Permutation matrix, 106
 Perovskite, 216
 Perovskite crystal structure, 106
 Perturbation, 291
 Phase diagram, 126
 Phosphors, 250
 Photocatalytic, 250
 Physical vapor deposition, 30
 Pie charts, 264
 Piezoelectric data, 106
 Piezoelectrics, 253
 Piezoresponse force microscopy (PFM), 256
 Poisson's ratio, 190
 Polynomial chaos expansions, 20
 Portland cement, 206
 Posterior, 98
 Posterior distribution, 62, 91
 Posterior probability, 16
 Posterior probability distribution, 48
 Potts model, 121
 Powder-based AM, 141
 Powering, 291
 Principal component analysis (PCA), 165, 218, 224, 273
 Principal components (PCs), 165, 225
 Prior distribution, 48, 49, 87
 Prior knowledge, 8, 86
 Prior probability distribution, 48
 Probabilistic model, 290
 Promising candidates, 164
 Proper orthogonal decomposition (POD), 20
 Pugh moduli, 190
 Pulsed laser deposition (PLD), 248

Q

QSPR for the bulk modulus, 233

Quadratic discriminant analysis (QDA), 80, 81
 Qualitatively, 265
 Quantitative, 265
 Quantitative structure-activity relationship, 174
 Quantitative structure-property relationship (QSPR), 225, 230
 Quaternary alloys, 206

R

Radial-basis-function support vector machine (RBF-SVM), 83
 Random forests, 165
 Ranking and selection, 68
 Rational strategies, 158
 Reduced-order, 14
 Reduced-order models (ROMs), 19, 37
 Regression, 6
 Regression tree, 148, 149
 Relative volume-change, 176
 Replicas, 122
 Resolution parameter, 120, 121
 Resubstitution, 89
 Resubstitution error estimate, 85
 Resubstitution estimator, 85
 Reverse Monte Carlo methods, 132
 Robot, 261
 Robustness, 69
 ROC curves, 84
 Root-mean-square (RMS), 82
 RTMS XRD detectors, 257

S

Sabatier's principle, 160
 Sc, Cr, Y, Zr, Pd and La, 168
 Sc, Cr, Zr, La, Pd and Y, 162
 Scanning evanescent microwave microscope (SEMM), 258
 Scanning magneto-optical kerr effect (SMOKE), 260
 Scanning probe microscopy (SPM), 256
 Scanning SQUID, 260
 Scree plot, 219
 Screening criteria, 160
 Screening framework, 168
 Segments, 123
 Selective laser melting (SLM), 141
 Separate sampling, 80
 Shannon entropy, 125–127
 Shannon information, 17
 Shear modulus, 190

- Side-to-Side, 105
- Side-to-Side ordering, 107, 109
- Signatures of the DOS, 234, 235
- Signatures of the DOS spectra, 224
- Similar thermodynamic activity, 166
- Simplicial distance, 295
- Simulation codes, 11
- Single-track experiments, 150
- Sintered combinatorial libraries, 251
- Slack, 201
- SOFC, 253
- Solid state materials, 272
- Solvable hard, 126
- Solvable phase, 127
- Space-time, 128
- Spatial interpolation of composition measurements, 294
- Spectral parameters, 228
- Spread, 292
- Squared exponential, 49
- Squared exponential covariance kernel, 50
- Squared exponential kernel, 50
- SrM, 264
- Stars, 164
- State-constrained optimal Bayesian classifier (SCOBC), 96
- Statistical, 8
- Statistical correlations, 214
- Statistical mechanics, 116
- Stochastic approximation, 19
- Stochastic block models, 126
- Structural glasses, 116
- Structure-modulus relationship, 235
- Structure-property relationship clustering, 285
- Structure-property relationships, 213
- Sub-compositional coherence, 290
- Subcomposition, 290
- Subcompositional incoherence, 290
- Substrate length, 23
- Suolunite, 209
- Support vector machines, 4
- Surrogate, 38
- Surrogate models, 14, 19
- Swarm intelligence, 118
- Symmetry-breaking, 214
- Symmetry-mode analysis, 215
- T**
- TBDO, 200
- Temperature, 126
- Ternary system, 130, 264
- Thermal, 260
- Thermal conductivity, 178
- Thermal fluctuations, 121
- Thermochemical splitting of water, 158
- Thermodynamic functions, 128
- Thermoelectrics, 177
- Thick film ceramics, 248
- Thin-film, 214, 248
- $\text{Ti}_2\text{Al}(\text{C}_x\text{N}_{1-x})$, 206
- Tiered pipeline screening, 279
- Tolerance factor, 112, 215
- Total bond order (TBO), 191
- Total bond order density (TBOD), 189, 191, 200, 209, 210
- Training set, 199
- Transition metals, 30
- Triangulation, 296
- True error, 84
- U**
- Uncertainties, 6
- Uncertainty class, 86, 87
- Uncertainty propagation, 20
- Uniform prior, 88
- Universal elastic anisotropy (A^U), 203
- Universally consistent, 81
- Unsolvable, 126
- USA, 243
- Utility function, 16
- V**
- Vapnik-Chervonenkis (VC) dimension, 81
- Variation matrix, 292
- VASP, 191
- Virtual DOS, 237
- Virtual DOS spectra, 224
- Virtual screening, 173, 178
- Visualizations algorithms, 103
- Vitreloy, 209
- VRH polycrystalline approximation, 190
- W**
- Water splitting/ H_2 , 256
- WEKA, 199
- Wisdom of the crowds, 117, 118
- Y**
- Young's modulus, 190
- Z**
- $\text{Zr}_x\text{Cu}_{1-x}$, 209
- $\text{Zr}_x\text{Cu}_y\text{Al}_z$, 209