

UNCERTAINTY QUANTIFICATION WITH SURROGATE MODELS IN ALLOY MODELING

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Jesper Toft Kristensen

May 2015

© 2015 Jesper Toft Kristensen
ALL RIGHTS RESERVED

UNCERTAINTY QUANTIFICATION WITH SURROGATE MODELS IN ALLOY MODELING

Jesper Toft Kristensen, Ph.D.

Cornell University 2015

The success of computational materials science in designing the materials of the future relies on the computation of materials properties using temporally expensive computer codes. Surrogate models, or simply surrogates, are employed to approximate the computer code and simultaneously decrease the computational effort by orders of magnitude enabling characterization, optimization, and design of materials *in silico*. To learn a surrogate, a few observations must be obtained from the expensive code, but once learned, the surrogate can predict the computer code output for *any* input. In most cases, one is faced with the task of choosing a narrow subset of best candidates among a larger set of potential surrogates for a given task.

In this thesis, we present two new methods which improve surrogate modeling, and a third method capable of quantifying the uncertainty in materials properties predicted from the approximating surrogate. Our methods are general, but will be applied to the subset of materials called alloys.

We start out by applying information theory to improve thermodynamic characterization of materials using surrogates. We explore employing relative entropy to measure “distance” between Boltzmann distributions, and find that our approach greatly improves current thermodynamic characterization in much better agreement with experiments.

Then, we provide a rigorous way to propagate the uncertainty from not know-

ing the best surrogate candidate to use and from having observed only a limited number of computer code observations during the process of lining up a set of relevant candidates. We show that the surrogate indeed captures properties to a satisfactory degree, as has been tacitly assumed so far in the field, with a predictive variance on the order of 5-10 %, which can be further improved if needed.

Finally, we consider a materials design problem using surrogates. We consider the task of minimizing the thermal conductivity in nanowires. The low symmetry of these wires motivates us to introduce a modification to the traditional surrogate model employed to learn the property of such systems. We call the new technique the cluster expansion ghost lattice method and show how the method successfully predicts the nanowire structure with lowest thermal conductivity.

BIOGRAPHICAL SKETCH

Jesper Toft Kristensen was born near Copenhagen in Denmark on April 27 1987. He earned his Bachelor of Science degree in Physics and Nanotechnology from the Technical University of Denmark in 2010. He received his Master of Engineering (MEng) degree from Cornell University (Cornell) located in the United States in Ithaca New York in 2011 from the school of Applied and Engineering Physics (AEP). He then joined the doctoral AEP program at Cornell in 2011 where, as part of this program, he received a Master of Science degree in 2014.

Dr. Kristensen has received numerous awards including the Henri S. Sack Memorial Award for the best academic performance in the AEP MEng program at Cornell. He entered the AEP MEng program at Cornell as a MEC fellow for being an outstanding applicant to the program. He is the recipient of the Frants Alling scholarship, the Vera and Carl Johan Michaelsen scholarship, the Reinholdt W. Jorck and Wife scholarship, the Otto Mønsted scholarship, and the Oticon Foundation scholarship.

While pursuing his degree, Dr. Kristensen worked as a tutor in the Learning Strategies Center, part of the physics department at Cornell. He also tutored for the Office of Academic Diversity Initiatives and was a teaching assistant in the AEP department in the first semester of his doctorate program.

Dr. Kristensen has attended multiple international conferences and seminars to present his research and has published his work in the international Physical Review B and Elsevier journals.

Dr. Kristensen's dissertation on Uncertainty Quantification with Surrogate Models in Alloy Modeling was supervised by Prof. Dr. Nicholas John Zabaras.

To my mother Lisbeth, my father Jens,
and my siblings Mette, Lene, and Ida.

ACKNOWLEDGEMENTS

First I would like to thank Prof. Dr. Nicholas John Zabaras for guiding me through the work in this thesis and for always expecting and demanding the best from me. The environment in my research group has prepared me to work both by myself but also in highly motivated teams striving for nothing but the best. I am grateful to Prof. Zabaras for letting me work on what I found interesting and for allowing me to develop into an independent researcher. Prof. Zabaras has an extreme breadth and depth of knowledge and understanding of what topics are relevant to work on, which helped me choose specific problems in the beginning of my doctorate degree and later served as a reliable source of advice on potential projects.

I am grateful to Prof. Dr. Richard Hennig for offering me guidance and advice whenever I needed it and for always being supportive. I am thankful that he let me work with his group on a collaborative project to further expand my knowledge of computational materials science. I also appreciate technical support from him on density functional theory during my first project and for providing access to the HiPerGator supercomputer in Florida.

I am thankful to Cornell University for providing world class engineering lectures and guidance from the faculty body, which has helped form me as a scientist. In particular, I would like to thank Prof. Dr. Thomas Stiadle for his teaching, Prof. Dr. Bruce Kusse for offering, and also teaching, the course Mathematical Physics in the AEP department, which changed my view on mathematical physics entirely. The same goes for Prof. Kusse's amazing class on statistical physics, which greatly improved my fundamental knowledge on this topic.

I appreciate computing resources provided by Pacific Northwest National Laboratory (PNNL) and help from the PNNL Institutional Computing team. I also

received computing resources from the National Energy Research Scientific Computing Center supported by the Office of Science of the US Department of Energy under contract No. DE-AC02-05CH1123. I am thankful for the support from the AFOSR Computational Mathematics Program, an OSD/AFOSR MURI09 award on uncertainty quantification, the US Department of Energy, the Office of Science, the Advanced Scientific Computing Research program, and the National Science Foundation (NSF) (award DMS-1214282).

Finally, I am thankful for the friends I have met and the network I have created both in my program, in my research group, and also in the wider community.

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	iv
Acknowledgements	v
Table of Contents	vi
List of Figures	viii
1 Introduction	1
2 The Cluster Expansion	5
2.1 Alloy Configuration Space	5
2.2 Functions over Configuration Space	6
2.3 Expanding Functions over Configuration Space in a Basis	7
2.3.1 Orthonormal Basis of a Single Lattice Site	7
2.3.2 Orthonormal Basis of N_c Lattice Sites	9
2.4 The Cluster Expansion and its Effective Cluster Interactions	11
2.4.1 Clusters	12
2.4.2 Symmetries and Cluster Orbits	12
2.4.3 Infinite Crystals	14
2.4.4 Truncating the Cluster Expansion	14
2.5 Cluster Expanding M -nary Alloys	15
2.6 Learning the Cluster Expansion	16
2.7 Alternative Notations of the Cluster Expansion	17
2.8 Summary	18
3 Markov Chain Monte Carlo Techniques	19
3.1 State Space	20
3.2 Reversible Jump Markov Chain Monte Carlo	22
3.2.1 State-Dependent Move Types	24
3.2.2 Developing a Reversible Jump Markov Chain Monte Carlo Algorithm for Linear Regression	26
3.2.3 Acceptance Probability of an Update Move in Regression	26
3.2.4 Acceptance Probability of a Birth Move in Regression	28
3.2.5 Acceptance Probability of a Death Move in Regression	29
4 Improving Thermodynamic Characterization of Alloys using In- formation Theory	31
4.1 Chapter Outline	32
4.2 Relative Entropy	33
4.3 Numerical Results	36
4.3.1 Magnesium Lithium Toy Model	40
4.3.2 Transition from Two-Phase Coexistence to Disorder in Dia- mond Silicon Germanium	41
4.3.3 Order to Disorder Phase Transition in bcc Magnesium Lithium	43

4.4	Conclusion	45
5	Bayesian Uncertainty Quantification in using Surrogates for Predicting Materials Properties	48
5.1	Chapter Outline	51
5.2	Background	52
5.3	Bayesian Uncertainty Quantification	54
5.4	Parametrizing the Response Surface with the Cluster Expansion . .	57
5.5	Bayesian Treatment of the Response Surface Parametrization	58
5.6	Sampling from the Posterior: Reversible Jump Markov Chain Monte Carlo	60
5.7	Uncertainty Quantification of the Magnesium-Lithium Ground State Line	62
5.8	Uncertainty Quantification of Disorder to Two-Phase-Coexistence Phase Transition	64
5.9	Data	65
5.10	Results	66
5.11	Discussion	68
5.12	Conclusion	73
6	Using Surrogates for Materials Design	75
6.1	Chapter Outline	76
6.2	The Cluster Expansion Ghost Lattice Method	77
6.3	Application: Predicting Lattice Thermal Conductivity of Si-Ge Nanowires	82
6.3.1	Nanowires as Thermoelectric Devices	82
6.3.2	Creating the Nanowire Data Set	83
6.3.3	Computing the Nanowire Thermal Conductivities and the Design Matrix	87
6.4	Results	90
6.5	Conclusion	94
7	Conclusion	96

LIST OF FIGURES

2.1	(Color) Demonstration of two different configurations of some binary alloy on a face-centered cubic lattice shown via the non-conventional unit cell. Each lattice site can host either of two chemical species: a red species, which could be, e.g., Mg or a blue species which could be, e.g., Li. σ_i is the configuration vector of configuration i identifying which chemical species occupies which lattice site.	5
4.1	(Color) ECI obtained for (a) the Mg-Li and (b) the Si-Ge system using least squares LOOCV as implemented in ATAT. Left-right arrows, with colors and text, group the number of sites in the clusters associated with the ECI. “ i -pt” is shorthand for i point clusters. Within each group, spatially larger clusters are found further to the right. The total number of clusters reported excludes the empty and 1-pt cluster. The inset in (a) shows the ECI for all clusters except the nearest neighbor 2-pt cluster (n.n. 2-pt) for a more detailed variation. Notice the difference in scale in energies between the two systems.	38
4.2	(Color) Canonical ordered phase transition from C11 _b to disorder in a toy model based on the Mg-Li system at 66 % Mg. The left-right arrows with associated temperatures identify the peak-to-peak distance between predicted transitions of the two considered methods and the full model. Relative entropy is measured on the right scale as indicated by the dashed arrow	41
4.3	(Color) Si-Ge system. (a) ECI obtained with the relative entropy method at 2000 K. Compare to Fig. 4.1(b). (b) shows the difference in ECI versus T between least squares with LOOCV (LS LOOCV) and relative entropy. Circled numbers in (b) identify the clusters in (a) selected for plotting.	42
4.4	(Color) Prediction of two-phase coexistence to disorder phase transition as indicated by a “divergence” in constant pressure heat capacity C_p for Si-Ge on a diamond lattice at 50 % composition. The least squares LOOCV method is compared to relative entropy. Both methods give (almost) identical curves for the heat capacity. The rectangular box near the peak reports the temperature at which the maximum value of the heat capacity was obtained which, to within 1 K, was the same for both methods.	43
4.5	(Color) Mg-Li system. (a) ECI obtained with the relative entropy method at 2000 K. Compare with Fig. 4.1(a). (b) shows the difference in ECI versus T between least squares with LOOCV (LS LOOCV) and relative entropy. Circled numbers in (b) identify the clusters in (a) selected for plotting.	44

4.6	(Color) Prediction of the order to disorder phase transitions for Mg-Li bcc as indicated by “divergence” in the constant pressure heat capacity C_p at compositions (a) 33 %, (b) 50 %, and (c) 66 % Mg. In each plot least squares LOOCV (A: blue, associated with the left y -axis) are compared to relative entropy (B: red, associated with the right y -axis). The rectangular boxes on each peak report the temperature at which the maximum value of the heat capacity was obtained.	46
5.1	Data sets obtained via VASP. (a) Shows the quantum mechanical formation energies of 82 Mg-Li structures and (b) for 38 Si-Ge structures. The line of zero formation energy is shown as well. . . .	68
5.2	BL-RJMCMC simulation details for the Mg-Li system. In (a) we monitor the model complexity, i.e., k in Eq. (5.9), where the maximum complexity k_{\max} is 21 (not counting the empty cluster orbit). In (b) we show the estimated data noise ϵ in Eq. (5.5). Insets show similar results for the diamond Si-Ge system where, in (a), the maximum complexity is 16 (not counting the empty cluster orbit).	69
5.3	(Color) Uncertainty in the ground state line of bcc Mg-Li. Black dots are the formation energies obtained from VASP and their ground state line is the black line. Red dash-dot and green dash lines are LASSO-CV and least squares, respectively. The BL-RJMCMC result is plotted as a blue full line (median of the posterior) with the shaded areas representing the HPD.	70
5.4	(Color) Uncertainties in the disorder to two-phase-coexistence of Si-Ge at 50 % composition identified by a peak in Eq. (5.12). Five different runs are shown each in a different color. The inset summarizes the main figure via the posterior median as a black vertical line surrounded by the HPD as the shaded area.	70
5.5	(Color) ECI obtained via the BL-RJMCMC method (box plots), least squares (red stems), and LASSO-CV (green stems) plotted for the (a) Mg-Li bcc and (b) Si-Ge diamond system. Insets show the single site cluster orbit for increased clarity.	71
6.1	(Color) Demonstration of the CE-GLM for some hypothetical single-layer slab of Si (Silicon; red) and Ge (Germanium; blue) on a simple cubic lattice. Ghost sites are white. The black square surrounding the boxes in (a) and (b) show where a cross-section is taken in (a) to reveal the slab in (b). A black ellipse in (b) delineates a 3-pt cluster containing ghost sites. The sites in the cluster are colored green, purely as a visual construct, to associate them with the cluster.	78

6.2	(Color) (a) Number of unknown ECI to be determined for an fcc lattice, representing some hypothetical alloy, versus the maximum spatial extent L of any included cluster in the cluster expansion, when not using any symmetries in the system (N_{none}) and when using all 48 bulk space group symmetries (N_{bulk}). The inset shows the ratio of the main plot curves versus the same L . The parenthesized numbers above the abscissa at L equal to 4, 5, and 7 Å are matched with those in (b). (b) The fcc lattice considered with a 5 Å lattice constant and its lattice sites represented as green balls. The cubic cell structure is delineated with black lines. The red lines with parenthesized numbers show the maximum allowed length of any 2-pt cluster included in the cluster expansion when choosing L at three different values along the abscissa in (a).	80
6.3	(Color) The ATAT representation of an arbitrary Si (red atoms) nanowire from the data set \mathcal{D} alloyed with 6.4 % Ge (blue atoms). Ghost sites are black atoms, made small so the wire becomes visible. The atoms in the two images of the wire along its axis are colored with a slightly lighter color to distinguish them from the main wire. The thick black circle identifies one of the two nanowire images. Ge atoms are made slightly larger than Si atoms for visual clarity.	85
6.4	(Color) The corresponding LAMMPS representation of the nanowire in Fig. 6.3. The black surrounding box is the LAMMPS simulation cell. (a) Shows the nanowire looking down its axis, and (b) illustrates the same wire from the side. The horizontal arrow inside the cell in (b) identifies the x -axis. Double headed arrows report sizes of the wire, and the smaller arrows in (b), over the wire, identify the atom types. Ge atoms are made slightly larger than Si atoms for visual clarity.	86
6.5	(Color) Demonstration of the method used to obtain the best τ_m in Eq. (6.2) called τ_m^* . All ordinates are in units of W/m.K. (a) Shows the pure Si wire and (b) the PPG wire. Top plots are κ versus τ_m in Eq. (6.2). In the bottom plots the μ/σ graph is the mean/standard deviation of 40 different unweighted running averages of window sizes ranging from 50 to 200 ps applied to the top graphs. The μ/σ graph is measured on the right/left ordinate. The red circles with vertical lines crossing the abscissa mark the minima of the σ graphs and hence the times τ_m^*	90

6.6	(Color) (a) Nanowire data set of random Si-Ge configurations plotted in a histogram with the number of elements versus their thermal conductivity κ . (b) Each type of wire considered in the fitting and predictions of thermal conductivity: wires with random Si-Ge configurations (RW), the pure planes of Ge wire (PPG), and wires that are similar to the PPG wire (SPPG). Red/blue atoms are Si/Ge. For the SPPG wire, black circles show where two random atoms, one atom from one of the pure Ge planes, and one atom from the non-plane region, were substituted to perturb the PPG wire.	92
6.7	(Color) Predictions of thermal conductivity from the cluster expansion ghost lattice method (CE-GLM) versus the results from molecular dynamics. In (a) the least squares fit on RW train is shown. Only one of the 10 train/test splits are shown (green/white circles). In (b) the CE-GLM is tested on three different test sets: RW test, the SPPG wires, and the PPG wire. Note that in (b), RW train (green circles) is the same as in (a). Error bars on the SPPG wires and the PPG wire stem from 10 different random splits of the RW data into train and test. For more information on the various nanowire data sets see, e.g., Fig. 6.6(b).	94

CHAPTER 1

INTRODUCTION

One of the biggest achievements of computational materials science would be to produce a full-fledged virtual materials design laboratory thus removing any need for costly real world experimental testing. We imagine inputting a set of desired materials properties into a computer code, wait for a reasonable amount of time, and receive the output material optimal for the application at hand. Only a single material, the optimal one, needs to be produced in a real world laboratory and we know immediately that all other possible representations of this material are inferior.

We still have a long way to go to reach this goal successfully, but in order to do so, we must be skillful in at least the following two tasks: We must be able to optimize materials properties, such as the energy, the band gap, the thermal conductivity, etc., represented *in silico*, and second, we should be able to accurately characterize materials. An essential characterization is of thermodynamic nature. That is, we need to know, e.g., in which phases the material will exist at various external conditions and also how stable these phases are. In both of these tasks, we must have a clear understanding, and quantification, of the uncertainties associated in our work.

One of the main reasons why this virtual materials laboratory is not a reality during the time of writing is the fact that computing the property of a material requires running temporally expensive computer codes. For example, if we want the accurate *ab initio* quantum mechanical energy of a single representation of a material, we can run the Vienna *ab initio* simulation package (VASP),^{1,2} which currently takes hours even on a supercomputer. This introduces a serious problem

since any optimization task needs to search typically many hundreds of millions different representations of the material in its way to the optimal answer. It seems that we are faced with an insurmountable computational effort. To make matters worse, for thermodynamic modeling of a material, we rely on statistical thermodynamics which demands the evaluation of ensemble averages requiring the materials property evaluated for the system found in the most likely (in principle *all*) states at a set of external conditions such as temperature and pressure. Again, although not an optimization task *per se*, we find ourselves facing the same problem of an infeasible computational task due to the expensive computer code. By now it should be clear that, overcoming the computational cost associated with obtaining materials properties is of central importance in computational materials science and represents one of the most relevant research areas.

The above discussion motivates the development of so-called *surrogate models*. A surrogate model, or simply, a surrogate, is a replacement for the accurate, expensive, computer code. In a nutshell, the surrogate attempts to learn the output, also called the *response*, of the computer code for *any* given input, in orders of magnitude less time than it takes to run the expensive code. The outputs of the computer code for all possible inputs are collectively called the *response surface*.

When employing a surrogate to learn the response surface, we necessarily introduce uncertainty into computed materials properties. This generates essential questions such as, how do these uncertainties affect our final predictions about which material is best for a given application? We may tell the experimentalist to produce material A when in fact material B is better because we relied too much on a single surrogate, this is the danger of using non-Bayesian methods. Bayesian methods on the other hand, as we will see in this thesis, accounts probabilisti-

cally for *all* possible surrogates consistent with a set of samples taken from the true response surface. Rigorously accounting for uncertainties with a probabilistic approach will tell us *how much* we believe the best material is A over B, leaving us with a much more informed approach to materials design. Furthermore, in the context of thermodynamic characterization, how does the approximate surrogate affect our uncertainty about which phase a material is most stable in for various temperatures and pressures? Disregarding this uncertainty can be detrimental to future materials design.

Although the methods developed in this thesis are general, we have decided to focus our attention on a subset of materials, namely alloys, i.e., materials composed of a set of different chemical elements of which at least one element is metallic. Alloys are very commonplace in society. Indeed, mercury mixed with silver, tin, copper, and zinc forms dental fillings, iron mixed with aluminum, nickel, cobalt, and other elements, creates the magnets in loudspeakers, copper mixed with zinc produces door locks and bolts, iron mixed with carbon and silicon is used to build bridges and cookware, copper mixed with nickel and manganese is used to create coins, aluminum mixed with copper, magnesium, and manganese forms materials used in automobiles, for aircraft body parts, and for military equipment. The list goes on.

In the case of alloys represented *in silico*, a particular surrogate model called the *cluster expansion* has been employed for many decades³ to represent alloy configurational properties, i.e., properties that depend on exactly where the atoms in the alloy are placed, called a configuration, on the lattice (for example a face-centered cubic (fcc) lattice with a basis)⁴ defining the geometry of the alloy. The cluster expansion is useful because of its computational speed and can, in prin-

ciple, be made arbitrarily accurate to the point where the true computer code is *exactly* represented. In practice, however, it is made approximate, as we discuss in Chapter 2.

In this thesis we address both the issue of optimization and that of thermodynamic characterization. In Chapter 4, we first set out to merge information theory with the cluster expansion to obtain a thermodynamic treatment in closer agreement with experiments.

Then, in Chapter 5, we address, in a rigorous way, the important issue of how much accuracy we sacrifice when replacing the expensive computer code with the surrogate. We present an approach capable of quantifying the uncertainties in properties computing using the surrogate. The framework is general, but will be applied to the cluster expansion.

Finally, in Chapter 6, we demonstrate a case of using surrogates for materials design. The cluster expansion surrogate model is employed to predict a material with a given specified property. We show how the cluster expansion is useful for materials design and in this process we modify the traditional cluster expansion applicable to bulk alloy systems to now handle low-dimensional systems of *arbitrary* shapes.

We will first need a theoretical framework which will be the topic of the following chapter.




Figure 2.1: (Color) Demonstration of two different configurations of some binary alloy on a face-centered cubic lattice shown via the non-conventional unit cell. Each lattice site can host either of two chemical species: a red species, which could be, e.g., Mg or a blue species which could be, e.g., Li. σ_i is the configuration vector of configuration i identifying which chemical species occupies which lattice site.

CHAPTER 2

THE CLUSTER EXPANSION

2.1 Alloy Configuration Space

Define an alloy via a lattice, which could but need not be of the Bravais type, and consider a set of atoms which can occupy each lattice site. A single configuration, mentioned briefly in Chapter 1, is represented by a vector σ and given by specifying which chemical species sits where on the lattice sites. Eee Fig. 2.1 for two different configuration examples if the alloy is binary and defined on the face-centered cubic (fcc) lattice. The set of all possible configurations of the lattice forms the configuration space, denoted $\{\sigma\}$. If considering an alloy of N_c possible lattice sites, the configuration space has 2^{N_c} configurations. To form the vector σ , the two chemical species are mapped to integers. A popular choice for the binary alloy, e.g., is ± 1 . For reasons that will become clear later we can call the integers the *spin values* of the site, but warn the reader that they have nothing directly to

do with quantum mechanical spins.

Most alloy properties are configuration-dependent, i.e., they depend on which chemical species occupy the lattice sites. We think of such properties as functions living in a function space over configuration space. These functions map configurations to alloy properties. Energy is a configuration-dependent property and we can think of associating with each configuration a real number—the energy of that alloy configuration. To simplify the discussion ahead we will now assume that the alloy property we are interested in is the energy.

2.2 Functions over Configuration Space

The alloy properties are computed via some computer code(s). This code is what brings us from a configuration to the alloy energy. We can think of it as a mathematical function mapping members of configuration space to energies. As discussed in Chapter 1, this function is complicated and temporally expensive to evaluate. We treat the computer code as a black box providing us energies/responses for inputs we provide. Our goal is to replace the expensive code with some alternative which is much faster. This is equivalent to saying that we are replacing Schrödinger’s equation, the exclusion principle, etc., with an emulator. At first, this seems impossible since we have no access to the internal workings of the code. On second thought, however, we are implicitly given information about how the code behaves each time it responds to a given input. This actually turns out to be enough for most cases when we are willing to sacrifice accuracy for speed.

Given enough examples of how the code responds to various inputs, we can learn to *emulate* the code. This is the idea behind a surrogate model and is generally a

very successful approach in most engineering fields. The surrogate needs to both be a good interpolator, predicting energies well within the set of inputs used to learn it, but also a good extrapolator, predicting energies for inputs which were not part of choosing the surrogate. It is the extrapolating, also called predictive, capabilities that open up the possibility for materials design as we will show in Chapter 6.

2.3 Expanding Functions over Configuration Space in a Basis

There are many possible choices for surrogates, but a popular model used for alloys is the cluster expansion. The idea is to create an orthonormal and complete basis in the space of functions over the 2^{N_c} dimensional configuration space and then expand configuration-dependent properties in this basis.

2.3.1 Orthonormal Basis of a Single Lattice Site

The approach to creating this basis can be broken into two main parts. First, we build an orthonormal and complete set of basis functions in the space of functions over configuration space of a *single* lattice site. Then, we form the direct product of *all* N_c site bases. In this section, to simplify the exposition, we will consider binary alloys, but the approach extends easily to a general M -nary alloy.

In mathematical terms, consider two configuration-dependent, general func-

tions $f(\cdot)$ and $g(\cdot)$ and define the inner product:

$$\langle f, g \rangle = \frac{1}{2} \sum_{\sigma \in \{-1, 1\}} f(\sigma) g(\sigma),$$

where we use scalar σ to represent a particular spin value on the given site, not to be confused with vector $\boldsymbol{\sigma}$ representing an alloy configuration. The configuration space of the single site has just two points, one for each spin value the site can be found in. We are free to represent any function over this configuration space of the single lattice site, say, $h(\cdot)$ as the vector:

$$\mathbf{h} \rightarrow \frac{1}{\sqrt{2}} \begin{pmatrix} h(1) \\ h(-1) \end{pmatrix}$$

where we have used the spin values directly instead of σ and we explicitly see that for a binary alloy the single-site configuration space is isomorphic to \mathbb{R}^2 and thus it is two-dimensional. For slight increased notational convenience later, we can re-write the inner product as the dot product

$$\langle f, g \rangle = \mathbf{f}^T \mathbf{g},$$

where the superscript T identifies the transposition operator.

To form an orthonormal basis in this two-dimensional configuration space, we can start with a standard polynomial complete basis guess and then make it orthonormal. Consider therefore the following two basis vectors (written first in function notation) taking as input the spin value σ on the site:

$$\boldsymbol{\Theta}_0(\sigma) = 1 \rightarrow \frac{1}{\sqrt{2}} \begin{pmatrix} \Theta_0(1) = 1 \\ \Theta_0(-1) = 1 \end{pmatrix} \quad (2.1)$$

and

$$\boldsymbol{\Theta}_1(\sigma) = \sigma \rightarrow \frac{1}{\sqrt{2}} \begin{pmatrix} \Theta_1(1) = 1 \\ \Theta_1(-1) = -1 \end{pmatrix}. \quad (2.2)$$

They are both associated with the same site and we can collect them in a matrix \mathbf{B} :

$$\mathbf{B} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

We then compute the QR factorization $\mathbf{B} = \mathbf{Q}\mathbf{R}$ and the k th column of \mathbf{R}^{-1} , the inverse matrix of \mathbf{R} , tells us how much each polynomial basis function contributes to form the new k th orthonormal basis vector. In this case (aside from an overall negative multiplicative constant):

$$\mathbf{R}^{-1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

The first column is $(1, 0)^T$ which means that the first orthonormal basis vector of this single site is just $\Theta_0(\cdot)$ (plus zero times $\Theta_1(\cdot)$). The second column $(0, 1)^T$ informs us that the second orthonormal basis vector of the single site is simply $\Theta_1(\cdot)$. Thus, the original polynomial basis (Eqs. 2.1 and 2.2) guess happens to be already an orthonormal complete basis for the configuration space of a single lattice site. Of course, the alloy would be very boring if only containing a single site, so let us extend these developments to N_c sites.

2.3.2 Orthonormal Basis of N_c Lattice Sites

Configuration space now has 2^{N_c} points. Define the inner product of the configuration space of all N_c lattice sites:

$$\langle f, g \rangle = 2^{-N_c} \text{Tr}^{(N_c)} \mathbf{f}^T \mathbf{g},$$

where the generalized trace over N_c sites is the operator

$$\text{Tr}^{(N_c)} = \sum_{\sigma^{(1)}} \sum_{\sigma^{(2)}} \cdots \sum_{\sigma^{(N_c)}} ,$$

and $\sigma^{(j)}$ is the possible spin values of site j , i.e., $\sigma^{(j)} \in \{-1, 1\}$ for all j . $\sigma^{(i)}$ should not be confused with the (fixed) i th element of $\boldsymbol{\sigma}$. $\sigma^{(i)}$ is just a summation variable representing site i . The key to forming the configuration space basis over N_c sites is to take the direct product of each single-site basis sets found and to realize that such a product keeps the resulting basis functions orthonormal as long as the single-site basis set is orthonormal.⁵

The result, as shown in Ref. [5], is that each basis function over N_c sites is formed by taking all possible combinations of the single-site basis functions. One particular basis function over all N_c sites is obtained if taking the product of the $\Theta_0(\cdot)$ functions from each of the N_c sites which gives a constant function over configuration space. Another basis function over the N_c sites is found by taking the $\Theta_0(\cdot)$ basis function from all N_c sites except, say, site 1 where we take the $\Theta_1(\cdot)$ basis function instead, and so on, for *all* possible choices of basis functions.

Mathematically, we can define the vector \mathbf{b} with the i th element being k if we want to use basis function Θ_k for site i , and remind ourselves of the configuration vector $\boldsymbol{\sigma}$ defined in Section 2.1 whose i th element identifies the spin value on site i now to be parsed to the corresponding basis function identified by b_i . As an example, say spin value +1 occupies site i , i.e., $\sigma_i = 1$, and that we want to use basis function $\Theta_1(\sigma) = \sigma$ for this site, i.e., $b_i = 1$. Then, to get the basis function value for the particular configuration we input s_i to the basis function identified by b_i thus obtaining $\Theta_{b_i}(\sigma_i) = 1$. Another configuration of the same alloy could (but need not) change σ_i leading to a different value of this same basis function if it happens that the opposite spin value occupies site i now: $\sigma_i = -1$.

2.4 The Cluster Expansion and its Effective Cluster Interactions

In the case of the binary alloy, we are now ready to define the symbol Γ , called the *cluster function*, to represent the basis function over N_c lattice sites and write it as:

$$\Gamma_{\mathbf{b}}(\boldsymbol{\sigma}) = \prod_{j=1}^{N_c} \Theta_{b_j}(\sigma_j), \quad (2.3)$$

where the product is over all the lattice sites. We will see in the next section how we need not, in general, take the product over *all* lattice sites in Eq. (2.3). Note that our Eq. (2.3) is the same as Eq. (3) in Ref. [6]. We have followed the notation closely in Ref. [6] on purpose to provide a detailed explanation of how the cluster expansion notation is derived. This should bring the unfamiliar reader quickly up to speed with reading the cluster expansion literature.

With Eq. (2.3) we can finally write down the cluster expansion of some general configuration-dependent total property value $Q(\cdot)$ (an extensive quantity) as the following sum over cluster functions:

$$Q(\sigma) = \sum_{\mathbf{b}} V_{\mathbf{b}} \Gamma_{\mathbf{b}}(\boldsymbol{\sigma}), \quad (2.4)$$

where $V_{\mathbf{b}}$ are unknown expansion coefficients called *effective cluster interactions* (ECI). This is not the final form to be used in this thesis, but provides a starting point for further developments in the following sections.

2.4.1 Clusters

The presence of the $\Theta_0(\cdot)$ basis function, which is just 1, allows the following geometric interpretation of the basis functions. Since any site in the lattice which hosts the basis function $\Theta_0(\cdot)$ contributes with a factor of 1 in Eq. (2.3), we can, effectively, ignore the site. In other words, the product becomes over just the sites with non- $\Theta_0(\cdot)$ basis functions. It is this realization that leads us to only worry about collection of lattice sites with non- $\Theta_0(\cdot)$ basis functions. Each such collection is called a *cluster*, the very namesake of the expansion we will be interested in as a surrogate in the present thesis. We can think of visualizing clusters, e.g., as sticks connecting these lattice sites. Eq. (2.4) looks deceptively simple, but the cluster functions rely on the computation of the clusters (satisfying certain symmetries to be discussed in Section 2.4.2) which presents the main computational effort underlying the expansion.

As a final note, to those familiar with the Ising model,⁷ the cluster expansion is equivalent to an extended Ising model accounting for next-nearest neighbor interactions, interactions of groups of four sites, groups of five sites, etc., as justified in Ref. [8].

2.4.2 Symmetries and Cluster Orbits

Essential to the success of cluster expansions are the *space group symmetries* of the “empty” lattice. By empty, we mean the lattice with no particular configuration, e.g., an fcc lattice with no atoms occupying the sites. Alternatively, we can get the same result by determining the space group symmetries of the lattice with the same chemical species occupying all sites. Consider two different clusters \mathbf{b} and

\mathbf{b}' . They have the same ECI if they can be brought to lie on top of each other, i.e., if they are equivalent, under a lattice space group operation. Mathematically, $V_{\mathbf{b}} = V_{\mathbf{b}'}$ in Eq. (2.4). We will take $\mathbf{b}' \in \mathbf{b}$ to mean that \mathbf{b}' is equivalent to \mathbf{b} under the space group operations of the lattice. Since we can group clusters under the same ECI value we name each such group. In this thesis we will call it a *cluster orbit*, or simply an orbit for short. Some references also use the term “cluster family”.

This discussion changes the way we present Eq. (2.4) in that we can now break the sum into one over the orbits and, for each orbit, a sum over the equivalent clusters. Let α denote an orbit, then Eq. (2.4) takes the form:

$$Q(\sigma) = \sum_{\alpha} V_{\alpha} \sum_{\mathbf{b} \in \alpha} \Gamma_{\mathbf{b}}(\sigma), \quad (2.5)$$

which is already in a much more useful form. The importance of employing symmetries comes from the fact that, for orbit α , we have reduced the number of unknowns to be determined by the number of equivalent clusters in this orbit. This number is denoted M_{α} .

But there is a more subtle point to be made as well. The expensive computer code, thought of as a function, belongs to a subspace in the space of functions over configuration space. This subspace contains functions that satisfy the space group symmetries of the system. Loosely speaking, when employing symmetries in our surrogate, we “approach” this subspace and more quickly converge to a good approximation to the code.

We will return to the importance of symmetries in Chapter 6.

2.4.3 Infinite Crystals

If $Q(\cdot)$ is the total energy of some bulk crystal extending to infinity along all directions then $Q(\cdot)$ would be infinite itself (just think of adding the self energy of each atom to the total energy). Therefore, it is much more useful to make $Q(\cdot)$ an intrinsic quantity. We make it the energy *per* unit cell.

Eq. (2.4) takes the form

$$q(\boldsymbol{\sigma}) = \sum_{\alpha} m_{\alpha} V_{\alpha} \phi_{\alpha}(\boldsymbol{\sigma}), \quad (2.6)$$

where $q(\cdot)$ is the per-unit-cell version of $Q(\cdot)$, m_{α} is the per-unit-cell version of M_{α} introduced at the end of Section 2.4.2 and $\phi(\cdot)$ is a *correlation function* defined as the average over cluster functions:

$$\phi_{\alpha}(\boldsymbol{\sigma}) = \frac{\sum_{b \in \alpha} \Gamma_b(\boldsymbol{\sigma})}{N_u m_{\alpha}}, \quad (2.7)$$

where N_u is the number of unit cells in the structure from which the property was computed using an appropriate computer code. This makes the correlation function a real number in the range $[-1, 1]$. Eq. (2.6) is a standard way of presenting the cluster expansion in the literature.

2.4.4 Truncating the Cluster Expansion

In practice, a truncation of Eq. (2.6) to some upper limit α_{\max} is necessary. This is achieved by fixing the maximum number of crystal sites present in any cluster as well as its maximum spatial extent defined throughout this thesis as the largest distance between any two crystal sites in the cluster.⁹

Much of the cluster expansion literature has been, and is still, devoted to finding the best set of clusters to be used in the summation Eq. (2.6), see, e.g., Ref. [10].

2.5 Cluster Expanding M -nary Alloys

In this section we generalize the developments leading to Eq. (2.6). The presentation follows closely that of Ref. [11]. Any site can now host M_i species. So far, we have assumed $M_i = 2$. Note that this also allows for site i hosting at maximum two different species while site j hosts, say, three different species and so on.

The reader will find the cluster expansion in many different forms in the literature. In this section we will use a notation slightly different than Eq. (2.6). This is mainly to allow the reader, who may not be an expert on cluster expansions, to more quickly adapt to changes in the cluster expansion notation in the literature.

The expansion now takes the form

$$q(\boldsymbol{\sigma}) = \sum_{\boldsymbol{\alpha}} m_{\boldsymbol{\alpha}} J_{\boldsymbol{\alpha}} \langle \Gamma_{\boldsymbol{\alpha}'}(\boldsymbol{\sigma}) \rangle_{\boldsymbol{\alpha}}, \quad (2.8)$$

where $J_{\boldsymbol{\alpha}}$ is now the symbol for the ECI and the correlation function, which was called $\phi(\cdot)$ in Eq. (2.6), is now written explicitly as an average over the clusters $\boldsymbol{\alpha}'$ symmetrically equivalent to $\boldsymbol{\alpha}$ (\boldsymbol{b} in Eq. (2.7)).

In Eq. (2.8) we have changed the notational meaning of the vector $\boldsymbol{\alpha}$ representing a cluster. If M_i different chemical species can now occupy lattice site i , α_i can take values from zero to $M_i - 1$. If $\alpha_i = 0$, site i is not contained in the cluster, i.e., the basis function of site i is $\Theta_0(\cdot)$. Importantly, sites hosting different sets of species are to be considered symmetrically distinct when determining the space group symmetries.

The cluster function $\Gamma_{\alpha}(\boldsymbol{\sigma})$ is now defined similar to Eq. (2.3) as the following product over all N_c lattice sites in the system:

$$\Gamma_{\alpha}(\boldsymbol{\sigma}) = \prod_{i=1}^{N_c} \gamma_{\alpha_i, M_i}(\sigma_i),$$

where

$$\gamma_{0, M_i}(\sigma_i) = 1 \tag{2.9}$$

(what we called $\Theta_0(\sigma_i)$ in Eq. (2.3)) and

$$\frac{1}{M_i} \sum_{\sigma_i=0}^{M_i-1} \gamma_{\alpha_i, M_i}(\sigma_i) \gamma_{\beta_i, M_i}(\sigma_i) = \begin{cases} 1 & \text{if } \alpha_i = \beta_i \\ 0 & \text{otherwise,} \end{cases} \tag{2.10}$$

which implies that the cluster functions are orthonormal, i.e., $\langle \Gamma_{\alpha}, \Gamma_{\beta} \rangle$ is one if $\alpha = \beta$ and zero otherwise. The cluster expansion is thus an expansion over the average cluster function in all cluster orbits. Ref. [11] provides more details and discusses a particular implementation of $\gamma_{\alpha_i, M_i}(\cdot)$.

2.6 Learning the Cluster Expansion

The cluster expansion has learned a given computer code once the retained clusters in the truncation along with the values for the ECI have been identified. The approach is to obtain a data set $\mathcal{D} = \{(\boldsymbol{\sigma}_i, q(\boldsymbol{\sigma}_i))\}_{i=1}^N$ of N properties expensive to obtain. How to choose \mathcal{D} is an open question that our current research is trying to address. We will, in this thesis, take \mathcal{D} for granted. Typically, N is on the order of 50-200. A fitting method is employed such as least squares¹² with leave-one-out cross-validation (LOOCV) (in some cases modified with *ad hoc* weights)¹³ potentially coupled with genetic algorithms,¹⁴ or compressive sensing,¹⁵ to learn both the best clusters and associated ECI. For typical systems, 20-80 clusters

are retained, but can vary greatly depending on the property and the system in question.

2.7 Alternative Notations of the Cluster Expansion

In Chapter 4, we cast the cluster expansion Eq. (2.6) in the following form:

$$q(\boldsymbol{\sigma}|\boldsymbol{\gamma}) = \sum_{i=1}^M \gamma_i \phi_i(\boldsymbol{\sigma}) \quad (2.11)$$

(read the left hand side as: property q given the particular set of ECI $\boldsymbol{\gamma}$), where the sum is over the cluster orbits, indexed by i , chosen (somehow) for the system and γ_i is the i th element of the ECI now collected in the vector $\boldsymbol{\gamma}$. If we let $\boldsymbol{\alpha}_i$ be the cluster $\boldsymbol{\alpha}$ in Eq. (2.8) chosen to represent the i th cluster orbit (such a cluster is also called a prototype) then γ_i in Eq. (2.11) would be equivalent to $J_{\boldsymbol{\alpha}_i}$ in Eq. (2.8), and $\phi_i(\boldsymbol{\sigma})$ is simply $m_{\boldsymbol{\alpha}_i} \langle \Gamma_{\boldsymbol{\alpha}'_i}(\boldsymbol{\sigma}) \rangle_{\boldsymbol{\alpha}_i}$, where $\boldsymbol{\alpha}'_i$ denotes a cluster in orbit i which is symmetrically equivalent to $\boldsymbol{\alpha}_i$. This definition of $\phi(\cdot)$ is very close to the one in Eq. (2.6) except for the $m_{\boldsymbol{\alpha}}$ and the fact that the subscript notation now indexes a cluster orbit in some integer.

In Chapter 5 we shall view the cluster expansion purely from a linear regression perspective. From Eq. (2.8), the cluster expansion of a configuration dependent alloy property $q(\cdot)$ can be expressed in the linear regression form, with added noise $\boldsymbol{\epsilon}$, as:

$$\mathbf{q} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}, \quad (2.12)$$

where any dependence on $\boldsymbol{\sigma}$ is suppressed, the i th element of vector \mathbf{q} is $q(\cdot)$ of structure i with its associated noise in the i th position of $\boldsymbol{\epsilon}$, \mathbf{X} is called the *design matrix* and contains, in the i th row and j th column, the term $m_{\boldsymbol{\alpha}_j} \langle \Gamma_{\boldsymbol{\alpha}'_j}(\boldsymbol{\sigma}) \rangle_{\boldsymbol{\alpha}_j}$,

discussed in the previous paragraph, of structure i for cluster orbit j , and the ECI are now collected in the vector β , not to be confused with the (scalar) inverse temperature $\beta = 1/k_{\text{B}}T$ used in Chapter 4 or the notation for some generic cluster in Eq. (2.10).

2.8 Summary

The cluster expansion is a surrogate model for alloys which expands the expensive computer code in a basis set with associated expansion coefficients called ECI. The cluster expansion provides an exact representation of any configuration dependent property $q(\cdot)$ of alloys if untruncated, but in practice we truncate it to M clusters/terms. The expansion is used to obtain the design matrix in a linear regression setting where the properties are obtained via an expensive computer code. We learn the clusters and associated ECI from a training data set \mathcal{D} and given these, the property of *any* alloy configuration, even one not part of \mathcal{D} , can be obtained at, essentially, no computational cost, enabling alloy optimization and thermodynamic modeling.

CHAPTER 3

MARKOV CHAIN MONTE CARLO TECHNIQUES

In typical Bayesian settings we need to maximize some unnormalized joint posterior distribution and to collect samples from the target distributions later used for inference. This is a very difficult task leaving Bayesian probability theory essentially unemployed in practice until around the 1980s where Markov Chain Monte Carlo (MCMC) methods finally¹ gained popularity in the statistical community. MCMC can be used as a numerical sampling method to obtain random variables distributed according to “custom” *unnormalized* probability distributions, exactly as needed in typical Bayesian tasks. It scales well with high dimensions which makes it useful for problems with hundreds or even thousands of unknowns.

It is not just in Bayesian probability theory *per se* that MCMC finds use. Consider statistical mechanics. Here, the Boltzmann (probability) distribution plays an essential role. Indeed, if we consider a system in contact with a heat bath at temperature T , the probability $p(\cdot)$ of finding the system in a state of energy E is:

$$\frac{1}{Z} \exp\left(-\frac{E}{k_{\text{B}}T}\right),$$

where k_{B} relates energy to temperature and is called the Boltzmann constant, and Z is the normalization constant of the Boltzmann factor for the system also called the *partition function*. To compute it, we need to sum over all possible states of the system

$$Z = \sum_{j \in \text{all states}} \exp\left(-\frac{E_j}{k_{\text{B}}T}\right),$$

¹The MCMC method was devised in the 1950s.

where E_j is the energy of state j . The partition function captures a wealth of information. For our purposes, in short, the MCMC method is a way to sample from custom probability distributions.

3.1 State Space

We eventually want to evolve the system in a so-called Markov chain, the namesake of MCMC, for reasons that become clear as we go along.¹⁶ We need this chain to satisfy certain conditions, which will be discussed here as well. To start our treatment, first define a state space as a space containing all possible states of the system, i.e., we need some way of associating a meaning to the word “state” when saying “move from state i to state j ”. For the alloy, e.g., this space could be the configuration space discussed in Section 2.1, where we move between states by altering the identities of the chemical species occupying the (fixed) lattice sites, but it all depends on the application at hand.

We want the ability to move the system from state to state at random. This is in order to form a *stochastic process*. For such a process, the system can evolve in many different directions from the current state to some next state generally dependent on the entire past of the system. Contrast this to solving, say, a reasonably defined ordinary differential equation where we, given the initial condition, can find a deterministic solution. For some subset of stochastic processes the way we move to the next state only depends on the current state of the system, not on any other past states. In other words, the move can only be influenced by information associated with the state we are in at the very moment we are making the move decision. This subset of processes contains what are called Markov chains and what we will be interested in. For an example of a Markov chain please refer

to the “Drunkard’s walk” in Chapter 2 of Ref. [17].

If not careful, the Markov chain will not end up sampling from the distribution we need it to sample from. Concepts such as the chain having multiple so-called stationary distributions, transient states and cycles can all lead to invalid samples. There is, however, a way to bypass these issues. The trick is to ensure that the chain is *ergodic*.¹⁷ This means that we must be able, with some finite probability, to transition from state α to state β after some finite time. Ergodic chains are, however, still not enough, we need to study a subclass of Markov chains, namely those satisfying *detailed balance*. Detailed balance implies that the equilibrium probability for the state in a general set A to move to a state in set B is the same with A and B reversed.²² So in summary, the Markov chain converges to a custom distribution of our choosing if it is Markovian, it is ergodic, and it satisfies detailed balance. The question is then *how* to ensure this and the answer is to develop an algorithm which, taking as input the current state, tells us the probability of moving to any other state. A famous such algorithm is the Metropolis algorithm¹⁸ which was later generalized by W. K. Hastings to the Metropolis-Hastings algorithm.¹⁹ In 1995, the algorithm was further generalized and we will employ this in Section 5.6. The algorithm was devised by P. Green and is called the Metropolis-Hastings-Green algorithm.^{20,21} We will discuss this algorithm from a theoretical standpoint in the following section.

As a final note, the name MCMC is used to describe the class of the algorithms, including the ones just mentioned, which uses a Markov chain to sample from some custom distribution of interest.

3.2 Reversible Jump Markov Chain Monte Carlo

For some problems it can be beneficial, or directly necessary, to let the state space change dimensions as the Markov chain moves between states. For example, we want to allow a move from a state space where states are characterized by two-dimensional vectors to another state space where they are characterized by four-dimensional vectors. Regression problems come to mind. What if we want to infer, from the data, the number of basis functions needed in a given basis expansion regression problem, and not simply the values of their associated expansion coefficients? In that case, the Metropolis-Hastings-Green algorithm, briefly mentioned in the previous section, can be employed. The resulting technique is called the reversible jump MCMC (RJMCMC) method. To present it, we follow the approach, not requiring a measure theoretic approach defining dominant measures and Radon-Nikodym derivatives, given by Refs. [22, 23] and direct the reader to the more rigorous presentation in Ref. [20] if needed. As we go along in the following paragraphs, we refer the reader to Ref. [22] for the full details. The upcoming presentation will serve as the theoretical framework underlying Chapter 5, and have been included here to make this thesis self-contained.

We start by defining the state space denoted $\chi = \bigcup_{k \in \mathcal{K}} (\{k\} \times \chi_k)$ where k labels the state we are in. The k th state has n_k parameters. In our application, if $k = 2$, say, we are in a state of two basis functions and $n_2 = 2$ as well. The various states with $k = 2$ identifies different particular values of the regression coefficients, collected in the vector $\boldsymbol{\theta}_k$, belonging to these basis functions. Since the coefficients will be real-valued, we let $\boldsymbol{\theta}_k \in \chi_k \subset \mathbb{R}^k$.

We are, generally speaking, trying to infer some probability distribution $\pi(k, \boldsymbol{\theta}_k | \mathcal{D})$, in light of an observed data set \mathcal{D} , on the state space. Keeping k

fixed for the moment we can let a computer program move between two states by generating a set of, say, r random numbers u from the known joint density g . The proposed state, say, x' , is then generated via some deterministic function $h(x, u)$ such that $(x', u') = h(x, u)$. The newly introduced u' contains the r' random numbers generated from the known joint density g' required to move from x' back to x (reversing the move) using the inverse function h' of h . Let us accept the move from x to x' with probability $\alpha(x, x')$ and the reverse move with probability $\alpha(x', x)$. Detailed balance requires that:

$$\int_{(x, x') \in A \times B} \pi(x) g(u) \alpha(x, x') dx du = \int_{(x, x') \in A \times B} \pi(x') g'(u') \alpha(x', x) dx' du'.$$

We want to solve this for $\alpha(\cdot, \cdot)$. In order to do so, we assume $h(\cdot, \cdot)$ and its inverse $h'(\cdot, \cdot)$ are differentiable so that the integral over the primed variables can be converted, via a Jacobian, to one over the unprimed variables. We can make the two integrals identical if the integrands are the same over the integration domain. Therefore, we require, after the change of variables, that:

$$\pi(x) g(u) \alpha(x, x') = \pi(x') g'(u') \alpha(x', x) \left| \frac{\partial(x', u')}{\partial(x, u)} \right|. \quad (3.1)$$

One choice (not the only one) for $\alpha(\cdot, \cdot)$ is to let

$$\alpha(x, x') = \min \left\{ 1, \frac{\pi(x') g'(u')}{\pi(x) g(u)} \left| \frac{\partial(x', u')}{\partial(x, u)} \right| \right\}. \quad (3.2)$$

We can show that Eq. (3.2) is indeed a solution to Eq. (3.1) as follows. Multiply Eq. (3.2) through by the denominator in the fraction:

$$\begin{aligned} \pi(x) g(u) \alpha(x, x') &= \min \left\{ \pi(x) g(u), \pi(x') g'(u') \left| \frac{\partial(x', u')}{\partial(x, u)} \right| \right\} \\ &= \min \left\{ \pi(x') g'(u') \left| \frac{\partial(x', u')}{\partial(x, u)} \right|, \pi(x) g(u) \right\} \\ \Rightarrow \frac{\pi(x) g(u) \alpha(x, x')}{\pi(x') g'(u') \left| \frac{\partial(x', u')}{\partial(x, u)} \right|} &= \min \left\{ 1, \frac{\pi(x) g(u)}{\pi(x') g'(u') \left| \frac{\partial(x', u')}{\partial(x, u)} \right|^{-1}} \right\}, \end{aligned} \quad (3.3)$$

but the left hand side is just $\alpha(x', x)$ from Eq. (3.1) so we get:

$$\alpha(x', x) = \min \left\{ 1, \frac{\pi(x)g(u)}{\pi(x')g'(u')} \left| \frac{\partial(x', u')}{\partial(x, u)} \right|^{-1} \right\}. \quad (3.4)$$

Now switch the labels x and x' in the proposed solution Eq. (3.2), it should equal Eq. (3.4), which we see it does. Also notice how Eq. (3.4) and Eq. (3.2) are reciprocals.

Let us pause and explain how the acceptance probability is actually used in the computer program. For a move from x to x'

Recall that we have kept k fixed during the above presentation, but now let that be allowed to change. That is, as we traverse the state space χ , some states will have $k = 2$ (we use two basis functions to describe the data) and other states will have $k = 4$ (we use four basis functions to describe the data). The question is, what changes in the above discussion. We find a pleasing answer: the dimensions of the state parameters associated with x and those associated with the new state x' can be different as long as the transformation from (x, u) to (x', u') remains a diffeomorphism. We must set up the requirement that $n + r = n' + r'$, which is called *dimension-matching*. As noted in Ref. [22] we can have r and/or r' to be zero which is useful, e.g., in cases where we transition towards lower k values (and thus, we do not need to augment the current state).

3.2.1 State-Dependent Move Types

We will be employing the RJMCMC algorithm in Chapter 5 and will see that, typically, in order to traverse the state space χ we need the concept of state-dependent move types. Consider first standard MCMC, where we have some way

of perturbing the current state, say, with a Gaussian random number. We take a state with $\boldsymbol{\theta} = (2.2, 4.1)^T$, say, and perturb it to $\boldsymbol{\theta}' = (2.4, 4.0)$, e.g. (notice that k is fixed in standard MCMC so has been suppressed) This represents a single move type. But how is such a move type going to let us jump from a state with $k = 2$ with, e.g., $\boldsymbol{\theta}_2 = (1, 4)^T$ to a state with $k' = 4$ with, e.g., $\boldsymbol{\theta}'_4 = (1, 4, 2, 8)^T$ when all we can do is perturb the parameter values? We need a move type which can offer this variation as well, i.e., a move type with a specification such as “augment the current state vector with 2 random numbers”.

We can implement any feature in the Markov chain as long as the chain still satisfies detailed balance. If we have two different move types in state x then detailed balance sums over all the possible move types. In order to maintain detailed balance we need to first introduce the probability of performing move type m in state x denoted $j_m(x)$. Next, we need to realize that, as long as detailed balance is satisfied for each separate move type, the overall detailed balance over all move types is satisfied as well.²² The acceptance probability Eq. (3.2) becomes move-type dependent and changes to:

$$\alpha_m(x, x') = \min \left\{ 1, \frac{\pi(x') j_m(x') g'(u')}{\pi(x) j_m(x) g(u)} \left| \frac{\partial(x', u')}{\partial(x, u)} \right| \right\}, \quad (3.5)$$

where the Jacobian is move-type dependent as well although not explicitly stated in the equation.

In the following section we will show a simple example of developing an RJMCMC algorithm for a linear regression problem.

3.2.2 Developing a Reversible Jump Markov Chain Monte Carlo Algorithm for Linear Regression

Consider a linear regression problem where we have only two basis functions, i.e., k can take values in $\{1, 2\}$. We want to learn whether the data can be explained better with a single basis function or with both basis functions.

We start by defining three move types: the birth move which increments k by 1, the update move, which updates the parameters for a given fixed k , and the death move, which decrements k by 1. If we have a full model ($k = 2$) we will not allow the birth move since there is no valid model with $k = 3$. This affects the move-type probabilities. In particular, if $k = 2$ we have that $j_{\text{birth}}(x) = 0$, $j_{\text{death}}(x) = 1/2$, $j_{\text{update}}(x) = 1/2$. If $k = 1$, we cannot make a death move since $k = 0$ is not a valid model, so $j_{\text{birth}}(x) = 1/2$, $j_{\text{death}}(x) = 0$, $j_{\text{update}}(x) = 1/2$.

Now let us derive the acceptance probability satisfying detailed balance. Consider the update move first since this is the easiest case of the three.

3.2.3 Acceptance Probability of an Update Move in Regression

The update move proceeds as follows. We find ourselves in a state $x = (k, \theta_k)$ and propose a move to $x' = (k, \theta')$, both states having the same k . This is a standard Metropolis-Hastings step. Let us make sense of Eq. (3.5) in this scenario. The way we perturb the current state parameter vector θ is, e.g., by generating $r = k$ random numbers and putting them in u . The deterministic map is thus $\theta' =$

$h(\theta, u) = \theta + u$ giving the Jacobian factor $|\partial\theta'/\partial\theta| = 1$. The proposal density g is simply the Gaussian density evaluated at u . Now consider the reverse move from the proposed state x' back to x . The map is now $\theta = h(\theta', u') = \theta' + u' = \theta' - u$. The last equality follows from the requirement that we must end up in state x (and we originally *added* u to θ). We have that g' is also Gaussian because this is the distribution we would use to perturb the state x' to x , so $g'(u') = g(-u)$, but since the Gaussian is symmetric, we have that $g(u) = g'(u')$ and these factors cancel in the acceptance probability. Finally, the move-type probabilities cancel as well since $j_{\text{update}}(x) = j_{\text{update}}(x')$. The Metropolis-Hastings algorithm for symmetric distributions is actually just the Metropolis algorithm and this is what we have here. The final answer is that, for the update move, we should accept moves with probability (we can ignore k since it is the same for both states, so use directly the parameter θ instead of the more general x):

$$\alpha_{\text{update}}(\theta, \theta') = \min \left\{ 1, \frac{\pi(\theta')}{\pi(\theta)} \right\}. \quad (3.6)$$

To evaluate this probability, we simply need to compute the posterior distribution, which replaces the general distribution $\pi(\cdot)$ in a given Bayesian inference problem, in the two states. Notice that the expression involves the *ratio* of the posterior distribution in the two states which means that the normalization constant cancels, which is certainly one of the biggest advantages of MCMC as explained in the introduction to this chapter.

Next, let us consider the acceptance probability of a birth move.

3.2.4 Acceptance Probability of a Birth Move in Regression

We are in the state $x = (1, \theta)$ and propose a move to state $x' = (2, \theta')$. Since θ' has two dimensions we need to augment θ with an extra random variable that we will draw from a Gaussian (thus, $r = 1$ here leaving u as a single normally distributed variable). So now $\theta' = (\theta'_1, \theta'_2) = h(\theta, u) = (\theta, u)$. Notice that, more generally, there is also the option of generating u as a Gaussian (or anything else we please) and then pass it to another function, say, $k(u)$, leaving us with $(\theta, k(u))$ which would lead to more complicated Jacobians, and could potentially improve the state space traversal, but in this case we simply obtain:

$$\begin{vmatrix} \frac{\partial \theta'_1}{\partial \theta_1} & \frac{\partial \theta'_1}{\partial u} \\ \frac{\partial \theta'_2}{\partial \theta_1} & \frac{\partial \theta'_2}{\partial u} \end{vmatrix} = \begin{vmatrix} 1 & 0 \\ 0 & 1 \end{vmatrix} = 1.$$

The jump probability of a birth move is $j_{\text{birth}}(x) = 1/2$ and since we generated the proposed new state from a Gaussian we have that $g(u)$ is the Gaussian density evaluated at u . This finishes the forward move. Next, let us discuss the reverse move going from x' to x which involves now a death move (from $k = 2$ to $k = 1$). First, the jump probability is $j_{\text{death}}(x') = 1/2$. In order to make this reverse move, we do not need to generate any additional random numbers, so $r' = 0$. We simply need to set $(\theta, u) = h'(\theta') = (\theta'_1, \theta'_2)$. This means that $g'(u')$ should not appear in the detailed balance equation from the beginning and, effectively, we can set it to one here, leaving us with:

$$\alpha_{\text{birth}}(x, x') = \min \left\{ 1, \frac{\pi(2, \theta')}{\pi(1, \theta)} \mathcal{N}(u, \sigma^2)^{-1} \right\}, \quad (3.7)$$

3.2.5 Acceptance Probability of a Death Move in Regression

Let us now derive the death move acceptance probability. As per our remark following Eq. (3.4), we should find that this probability is the inverse of the birth move acceptance probability found in Section 3.2.4.

We start in state $x = (2, \theta)$ and propose to jump to $x' = (1, \theta')$. Since x' is of lower dimensionality, we do this simply by setting $(\theta', u') = h(\theta_1, \theta_2) = (\theta_1, \theta_2)$. The jump probability is $j_{\text{death}}(x) = 1/2$ and since we do not generate any random numbers in this move ($r = 0$), we have that $g(u)$ is effectively one. The Jacobian is:

$$\begin{vmatrix} \frac{\partial \theta'_1}{\partial \theta_1} & \frac{\partial \theta'_1}{\partial u} \\ \frac{\partial \theta'_2}{\partial \theta_1} & \frac{\partial \theta'_2}{\partial u} \end{vmatrix} = \begin{vmatrix} 1 & 0 \\ 0 & 1 \end{vmatrix} = 1.$$

Consider now the reverse move, i.e., we are in state $x' = (1, \theta')$ and propose a birth move to state $x = (2, \theta)$ via a specific set of parameters: $(\theta', u') = (\theta_1, \theta_2)$ given in the previous paragraph. That is to say, we are not interested in a jump from x' to *any* other state, but only interested in the single jump from x' which leads exactly back to x .

The jump probability is $j_{\text{birth}}(x') = 1/2$, and thus cancels with $j_{\text{death}}(\cdot)$, and to perform this reverse move we have to set $(\theta_1, \theta_2) = h'(\theta', u') = (\theta', u' = \theta_2)$ so $g'(u')$ is a Gaussian evaluated at θ_2 . We are ready to write down the acceptance probability:

$$\alpha_{\text{death}}(x, x') = \min \left\{ 1, \frac{\pi(1, \theta')}{\pi(2, \theta)} \mathcal{N}(\theta_2, \sigma^2) \right\},$$

and indeed this is the inverse of Eq. (3.7) when we change variables (think of u in

Eq. (3.7) as being the “second parameter of θ , i.e., θ_2 ”).

CHAPTER 4

IMPROVING THERMODYNAMIC CHARACTERIZATION OF ALLOYS USING INFORMATION THEORY

Construction of binary alloy phase diagrams are important to thermodynamic characterization of materials and generally relies on the cluster expansion discussed theoretically in Chapter 2. In the present chapter, the alloy property will be the *ab initio* quantum mechanical energy, called simply energy, and thus, to offer a more concrete treatment, instead of $q(\cdot)$, we denote the energy $E(\cdot)$. We remind ourselves that, given the clusters with associated ECI, typically obtained from a least squares or a *least absolute shrinkage and selection operator* (LASSO) fit,²⁴ the latter to be discussed in Chapter 5, on a training set \mathcal{D} , the energy of any alloy configuration can be computed, and subsequently used for thermodynamic simulations.^{9,25–30}

When employing the above mentioned fitting techniques to learn the ECI, all structures in \mathcal{D} are weighted equally, but we can provide an intuitive argument as to why this weighting scheme does not necessarily lead to an optimal description of the thermodynamics of the system. Consider the case of a canonical ensemble. The Boltzmann factor, in equilibrium, dictates that more energetic states are exponentially less likely to be observed. So, at low temperatures, the partition function is mostly influenced by the few low energy states. The Boltzmann factor for all other states is essentially zero. Therefore, the *thermodynamic importance* of a state is quantified by its Boltzmann factor. We conclude that, as the temperature is increased, so is the importance of any state. Furthermore, at infinite temperature all states are equally important.

The previous discussion motivated the work in this chapter. It is desirable to investigate techniques which obtain the ECI based on thermodynamic arguments. All the necessary information is encapsulated in the probability distribution over states (PDS). The idea is to bring the PDS induced by the cluster expansion (candidate PDS) as close as possible to the true one. We propose to measure this distance in terms of *relative entropy* (also known as the *Kullback–Leibler divergence*).³¹ Thus, we obtain a *variational* problem, namely, the minimization of the relative entropy functional with respect to the candidate PDS. Though theoretically sound this problem is computationally intractable. To cope with this, we show that the relative entropy functional can be approximated by the variance (with respect to the true PDS) of the difference between the true and the candidate cluster expansion energy. Restricting this approximation on the observed data leads to a weighted least squares problem making the proposed approach computationally attractive.

We test the performance of our method in a study of canonical phase transformations in Si-Ge (two-phase coexistence to disorder) and Mg-Li (order to disorder) alloys at various compositions. A comparison is presented between the relative entropy results and least squares LOOCV where we, for Mg-Li, observe noticeable differences in the transition temperatures. Our results are found to be in better agreement with guiding experimental data.

4.1 Chapter Outline

First, the theoretical framework of relative entropy is discussed in Section 4.2. In Section 4.3, we explain how the *ab initio* observations were made along with details

on the thermodynamic calculations. Then, we present numerical results starting with a toy model in Section 4.3.1. This is followed by a study of a real Si-Ge diamond alloy (Section 4.3.2) predicting the transition from two-phase coexistence to disorder at fixed 50 % composition. The same transition has been simulated in Refs. [32–34] and also by using cluster expansions with least squares LOOCV in Refs. [9, 35]. As a final example, we then turn to an Mg-Li alloy on a body-centered cubic (bcc) lattice (Section 4.3.3) which has recently³⁶ been simulated at low temperature predicting order to disorder phase transformations at 33 %, 50 %, and 66 % Mg. Finally we compare the relative entropy results with both Ref. [36] and guiding experimental data in Ref. [37].

4.2 Relative Entropy

Let the binary alloy, as introduced in Chapter 2, be at fixed composition and in a heat bath at temperature T . The PDS is

$$p(\boldsymbol{\sigma}|\beta) = \frac{\exp(-\beta E(\boldsymbol{\sigma}))}{Z}, \quad (4.1)$$

where $\beta = 1/k_{\text{B}}T$, with k_{B} the Boltzman factor, $E(\boldsymbol{\sigma})$ is the *ab initio* energy of alloy state $\boldsymbol{\sigma}$, and $Z = \sum_{\boldsymbol{\sigma}} \exp(-\beta E(\boldsymbol{\sigma}))$ is the partition function. The candidate PDS, given some ECI that we, in this chapter, will denote $\boldsymbol{\gamma}$, is

$$p(\boldsymbol{\sigma}|\boldsymbol{\gamma}, \beta) = \frac{\exp(-\beta E(\boldsymbol{\sigma}|\boldsymbol{\gamma}))}{Z(\boldsymbol{\gamma})}. \quad (4.2)$$

Please refer to Chapter 2 and Eq. (2.11) for the meaning of $E(\boldsymbol{\sigma}|\boldsymbol{\gamma})$. As discussed in the chapter introduction, we seek to bring the candidate PDS (Eq. (4.2)) as close as possible to the true PDS (Eq. (4.1)). The measure of “distance” we choose is the relative entropy, defined by

$$\mathcal{S}_{\text{rel}}[\boldsymbol{\gamma}] := \sum_{\boldsymbol{\sigma}} p(\boldsymbol{\sigma}|\beta) \ln \left[\frac{p(\boldsymbol{\sigma}|\beta)}{p(\boldsymbol{\sigma}|\boldsymbol{\gamma}, \beta)} \right]. \quad (4.3)$$

The relative entropy quantifies the information loss when using the candidate PDS instead of the true one. It has been used in numerous applications ranging from machine learning (variational methods),³⁸ coarse graining,^{39–42} to free energy calculations.⁴³ We postulate that the optimal ECI should be selected by minimizing Eq. (4.3).

Notice that Eq. (4.3) involves a sum over the entire phase space rendering its minimization practically impossible. Several approximations are essential in order to bring the problem to a soluble form. Towards this goal, let us first rewrite Eq. (4.3) as

$$\begin{aligned}\mathcal{S}_{\text{rel}}[\gamma] &= \beta \langle E(\boldsymbol{\sigma}|\gamma) - E(\boldsymbol{\sigma}) \rangle + \ln(Z(\gamma)) + \ln(Z) \\ &= -\langle \Delta \rangle + \ln \langle \exp(\Delta) \rangle \\ &= \ln \langle \exp[\Delta - \langle \Delta \rangle] \rangle,\end{aligned}$$

where $\langle \cdot \rangle$ denotes the expectation with respect to the true PDS and $\Delta := \beta(E(\boldsymbol{\sigma}) - E(\boldsymbol{\sigma}|\gamma))$ can be thought of as the error of the approximation, in units of the thermal energy. We now derive an approximation to the relative entropy whose minimization is equivalent to the solution of a linear system. To leading order in Δ we have

$$\begin{aligned}\mathcal{S}_{\text{rel}}[\gamma] &= \ln \left\langle 1 + (\Delta - \langle \Delta \rangle) + \frac{1}{2}(\Delta - \langle \Delta \rangle)^2 + \dots \right\rangle \\ &\approx \ln \left(1 + \frac{1}{2} \text{Var}[\Delta] \right) \approx \frac{1}{2} \text{Var}[\Delta] \\ &= \frac{1}{2} \left(\sum_{\boldsymbol{\sigma}} p(\boldsymbol{\sigma}|\beta) \Delta^2 - \left(\sum_{\boldsymbol{\sigma}} p(\boldsymbol{\sigma}|\beta) \Delta \right)^2 \right),\end{aligned}\tag{4.4}$$

where, in the last equality, we recall that Δ depends on $\boldsymbol{\sigma}$. Equating the derivative of Eq. (4.4) with respect to γ to zero,

$$\begin{aligned}0 = \frac{\partial \mathcal{S}_{\text{rel}}[\gamma]}{\partial \gamma_i} &= -\beta (\langle \phi_i(\boldsymbol{\sigma}) \Delta \rangle - \langle \Delta \rangle \langle \phi_i(\boldsymbol{\sigma}) \rangle) \\ &= -\beta^2 \text{cov} \left[(E(\boldsymbol{\sigma}) - \hat{E}(\boldsymbol{\sigma}|\gamma)), \phi_i(\boldsymbol{\sigma}) \right]\end{aligned}$$

and solving for γ , yields the following linear system:

$$\text{cov}[\phi(\sigma), \phi(\sigma)] \gamma = \text{cov}[\phi(\sigma), E(\sigma)], \quad (4.5)$$

where

$$\text{cov}[\mathbf{X}, \mathbf{Y}] := \langle (\mathbf{X} - \langle \mathbf{X} \rangle)(\mathbf{Y} - \langle \mathbf{Y} \rangle)^t \rangle$$

with \mathbf{X}^t being the transpose of \mathbf{X} , denotes the covariance of the true PDS between two vector functions \mathbf{X} and \mathbf{Y} of configuration space. However, Eq. (4.5) is still unmanageable. The next step is to approximate it using the observed data.

Assume we have collected N observations $\left\{ \left(\sigma^{(i)}, E(\sigma^{(i)}) \right) \right\}_{i=1}^N$, what we called \mathcal{D} in Chapter 2, where $\sigma^{(i)}$ is the configuration and $E(\sigma^{(i)})$ is the corresponding *ab initio* energy. To approximate Eq. (4.5), we replace all averages with respect to the true PDS with averages over the observed data. In other words, we approximate the average of any quantity \mathbf{X} by

$$\langle \mathbf{X} \rangle \approx \sum_{i=1}^N p_{N,i} \mathbf{X}(\sigma^{(i)}),$$

where

$$\mathbf{p}_N := \left(\frac{\exp[-\beta E(\sigma^{(1)})]}{Z_N}, \dots, \frac{\exp[-\beta E(\sigma^{(N)})]}{Z_N} \right),$$

with $Z_N := \sum_{i=1}^N \exp[-\beta E(\sigma^{(i)})]$. Thus, Eq. (4.5) becomes a weighted least squares problem:

$$\Phi^t \mathbf{W} \Phi \gamma = \Phi^t \mathbf{W} \mathbf{E}, \quad (4.6)$$

where \mathbf{E} is a vector of length N containing the observed energies, and

$$\mathbf{W} := (\mathbf{I}_N - \mathbf{p}_N \mathbf{1}_N) \text{diag}(\mathbf{p}_N) (\mathbf{I}_N - \mathbf{p}_N \mathbf{1}_N)^t,$$

where \mathbf{I}_N is the $(N \times N)$ identity matrix, $\mathbf{1}_N$ a $(1 \times N)$ vector of all elements equal to 1, and $\text{diag}(\mathbf{p}_N)$ an $(N \times N)$ matrix with the vector \mathbf{p}_N as its diagonal.

Eq. (4.6) is what makes the relative entropy method computationally tractable. The extend to which it is valid, requires further mathematical analysis that goes beyond the scope of this work. For the cases considered in Section 4.3, it was observed that the ECI obtained by solving Eq. (4.6) converged as the number of observations N increased in a way similar to least squares. Notice that the ECI are, generally, T -dependent through the weight matrix \mathbf{W} in Eq. (4.6).

Before concluding this section, it is instructive to compare Eq. (4.4) with the least squares loss function

$$\mathcal{L}[\gamma] = \sum_{\sigma} (E(\sigma) - E(\sigma|\gamma))^2 = \beta^{-2} \sum_{\sigma} \Delta^2. \quad (4.7)$$

Notice that the first summand in Eq. (4.4) is similar to that of Eq. (4.7), albeit weighted by the true PDS. The second summand does not have an analogue in least squares. The important observation is that learning efforts are concentrated on the relevant states mostly, at a given temperature. On the contrary, least squares always considers all states equally important. It is therefore anticipated that the relative entropy ECI would perform better when used in thermodynamic calculations.

4.3 Numerical Results

In this section, we explain how the *ab initio* data for the two alloys were obtained. Then, we briefly discuss the details of the underlying Monte Carlo scheme used to extract thermodynamic information needed to detect the phase transitions.

The first principles configurational energies were generated using VASP. For the Mg-Li bcc system we used Perdew–Burke–Ernzerhof (PBE)^{44,45} projector-

augmented-wave⁴⁶ pseudo potentials. We treated the $2p$ and $3s$ orbitals of Mg and the $1s$ orbital of Li as valences. Convergence to within 1 meV/atom was achieved around a wave cutoff of 340 eV, in agreement with Ref. [36] using the same technique, and was checked for both the pure structures and for a couple of randomly chosen mixed structures. Spin polarization was not included. We used a dense gamma-centered Monkhorst–Pack⁴⁷ k -grid for the Brillouin zone integration with just over 7500 k -points per reciprocal atom and scaled the grid according to the size of each supercell being computed. In this way, a training set of 82 structures was generated. Then, a cluster expansions with up until and including five point clusters was constructed using ATAT.^{9,11,13,48,49} The least squares LOOCV ECI were obtained using ATAT. We obtained a LOOCV score of 4.2 meV. The same training set was used to compute the ECI for the relative entropy method at each temperature using Eq. (4.6).

For the Si-Ge diamond system, we also used PBE pseudopotentials and the system converged to within 1 meV/atom around a wave cutoff of 340 eV using the same k -point grid type, density, and scaling as for the Mg-Li system. A training set of 38 structures was constructed and a cluster expansion containing up until and including three point clusters, as in Ref. [9], computed with ATAT. The least squares LOOCV ECI were obtained with ATAT. A LOOCV score of <1 meV was achieved. The training data was used to compute the ECI for the relative entropy method at each temperature using Eq. (4.6).

The clusters obtained with ATAT are used in the relative entropy method as basis functions, i.e., as columns in Φ in Eq. (4.6). The ECI obtained for both systems using the least squares LOOCV method are shown in Fig. 4.1. The empty cluster is left out as it is just an overall scale of the energy irrelevant when sampling

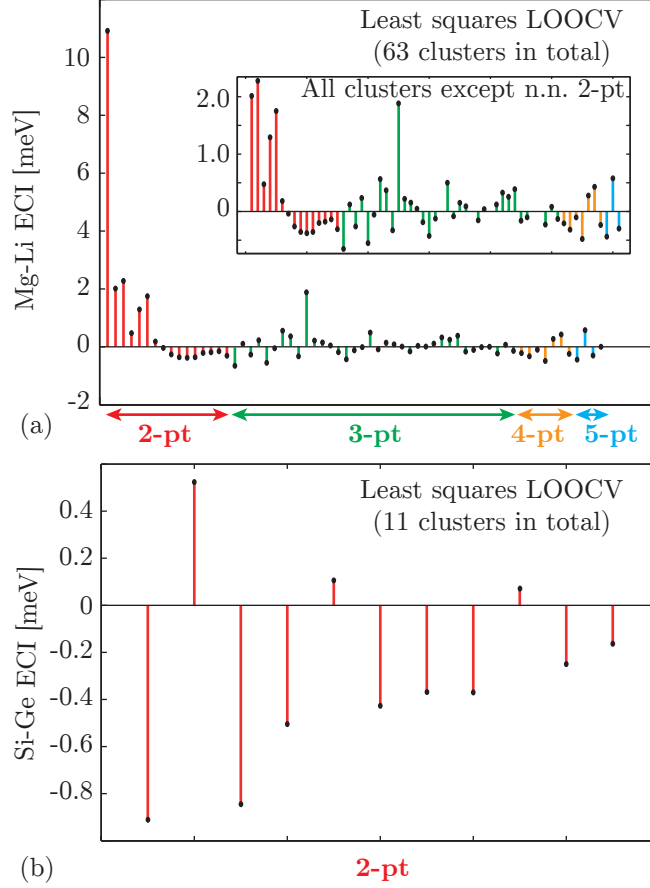


Figure 4.1: (Color) ECI obtained for (a) the Mg-Li and (b) the Si-Ge system using least squares LOOCV as implemented in ATAT. Left-right arrows, with colors and text, group the number of sites in the clusters associated with the ECI. “ i -pt” is shorthand for i point clusters. Within each group, spatially larger clusters are found further to the right. The total number of clusters reported excludes the empty and 1-pt cluster. The inset in (a) shows the ECI for all clusters except the nearest neighbor 2-pt cluster (n.n. 2-pt) for a more detailed variation. Notice the difference in scale in energies between the two systems.

the Boltzmann factor. By a similar argument, valid because the system is in a canonical ensemble, the one-point cluster ECI does not play a role either.

The ground states found for Mg-Li, at the compositions studied in this work, were: B_f , $B2$, and $C11_b$ at 33 %, 50 %, and 66 % Mg, respectively. Ref. [36] finds

the same result except obtaining a C11_b structure at 33 % Mg. We find that only the two pure diamond structures are ground states for the Si-Ge system at zero Kelvin.

Thermodynamic properties were calculated on a $30 \times 30 \times 30$ simulation cell with periodic boundary conditions. Sampling of the PDS was performed using an *Adaptive Sequential Monte Carlo* (ASMC) technique^{50,51} coupled with an underlying Metropolis–Hastings algorithm using double-spin-flip dynamics to conserve composition. ASMC approximates the PDS with a weighted finite set of delta functions, called *particles*, at some convenient initial temperature. The particles are then propagated from this starting point to any other temperature as desired, and the weights are updated accordingly. Steps in temperature are taken adaptively to ensure that the distribution does not undergo large changes, and thermodynamic information is recorded at each step.

The ASMC sampler was initialized with 512 particles, each on its own computational core, at 2000 K. Each particle was initialized randomly. At this high temperature random configurations are distributed (almost) correctly according to the Boltzmann factor offering a very convenient starting point. Nevertheless, each particle was thermalized with 100 sweeps.⁵² A target temperature of 50 K was specified and to improve the chances that we did not miss a transition, we split the interval 50-2000 in bins of size 10 K. ASMC drove the ensemble of particles from the right corner of this interval to the left.

The heat capacity at constant pressure was obtained from the standard statistical relation to the variance in energy

$$C_p = \beta^2 k_B \left(\langle E^2 \rangle - \langle E \rangle^2 \right).$$

For an infinite system a divergence in this quantity signals a phase transition. For

a finite system we should expect a distinct peak.

4.3.1 Magnesium Lithium Toy Model

The purpose of this section is to assess the predictive power of the relative entropy method compared to least squares LOOCV in the context of canonical order to disorder phase transitions. In order to do this we devise an artificial problem in which the phase transition is exactly known. We consider the Mg-Li system and assume that the true energy surface is the one obtained by the set of clusters and ECI determined by ATAT, see Fig. 4.1(a). For this cluster expansion, we compute the phase transition temperature from order to disorder (true transition temperature). We create an artificial data set by evaluating this cluster expansion on the 82 observed configurations and assume that the energies predicted are the true energies. Then, we reduce the cluster expansion to include only the two point clusters (of which there are 16) and train it on the artificial data using both least squares LOOCV and relative entropy. Finally, for each method we compute the order to disorder phase transition temperature for the 66 % Mg composition and compare them with the true one. It must be noted that special attention was paid to ensure that the reduced cluster expansion could predict the same ordered state, which was C11_b, as the artificial model. This was verified by monitoring the pair correlation values versus temperature. The results are shown in Fig. 4.2. The relative entropy method differs by 22 K of the true transition temperature compared to 70 K for least squares LOOCV. This is evidence of the anticipated, enhanced, predictive capabilities of the relative entropy method.

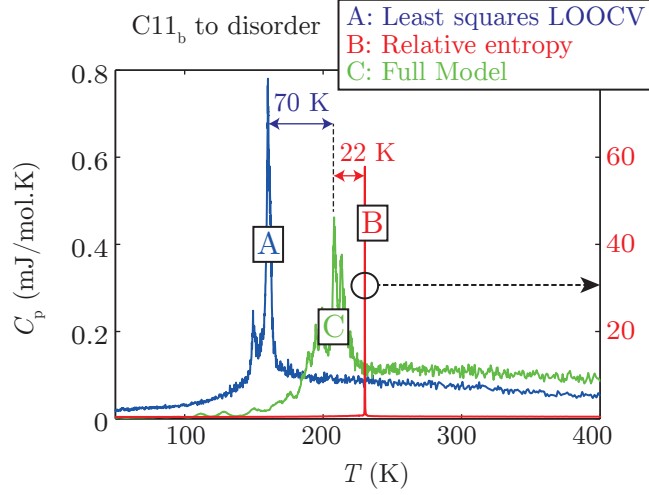


Figure 4.2: (Color) Canonical ordered phase transition from $C11_b$ to disorder in a toy model based on the Mg-Li system at 66 % Mg. The left-right arrows with associated temperatures identify the peak-to-peak distance between predicted transitions of the two considered methods and the full model. Relative entropy is measured on the right scale as indicated by the dashed arrow

4.3.2 Transition from Two-Phase Coexistence to Disorder in Diamond Silicon Germanium

The ECI obtained with relative entropy at the highest temperature 2000 K are shown in Fig. 4.3(a). At high temperatures relative entropy considers all states thermodynamically important. Therefore, the ECI should be comparable to least squares in Fig. 4.1(b) which is seen to be the case. The difference between the methods becomes increasingly pronounced as the temperature is lowered. This claim is supported by the difference in selected ECI versus T as shown in Fig. 4.3(b).

The prediction of the Si-Ge two-phase coexistence to disorder transition temperature at 50 % composition is seen for ECI obtained using both least squares LOOCV and relative entropy in Fig. 4.4. The least squares result, yielding a tran-




Figure 4.3: (Color) Si-Ge system. (a) ECI obtained with the relative entropy method at 2000 K. Compare to Fig. 4.1(b). (b) shows the difference in ECI versus T between least squares with LOOCV (LS LOOCV) and relative entropy. Circled numbers in (b) identify the clusters in (a) selected for plotting.

sition temperature around 339 K, is in good agreement with the 325 K found in Ref. [9], using the same method but different sets of clusters and Monte Carlo sampling methods.

Relative entropy agrees with least squares LOOCV. This can be explained from the small variation in the ECI between the two methods for this system on the order of 0.01 meV as seen in Fig. 4.3(b).

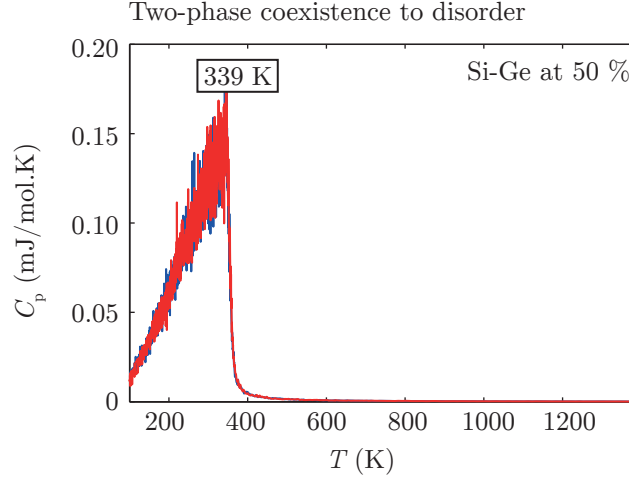


Figure 4.4: (Color) Prediction of two-phase coexistence to disorder phase transition as indicated by a “divergence” in constant pressure heat capacity C_p for Si-Ge on a diamond lattice at 50 % composition. The least squares LOOCV method is compared to relative entropy. Both methods give (almost) identical curves for the heat capacity. The rectangular box near the peak reports the temperature at which the maximum value of the heat capacity was obtained which, to within 1 K, was the same for both methods.

4.3.3 Order to Disorder Phase Transition in bcc Magnesium Lithium

The ECI obtained with relative entropy at the highest temperature 2000 K are shown in Fig. 4.5(a). Similar to Section 4.3.2, at high temperatures the ECI should be comparable to least squares LOOCV, shown in Fig. 4.1(a), which is found to be the case. The difference between selected ECI is shown in Fig. 4.5(b).

In Fig. 4.6 we show the predicted order to disorder transition temperatures for Mg-Li bcc at 33 %, 50 %, and 66 % Mg compositions. Using least squares LOOCV we obtain transition temperatures of 226, 304, and 207 K, respectively. This is in general agreement with Ref. [36] obtaining 190 and 210 K for 33 % and 66 % Mg,

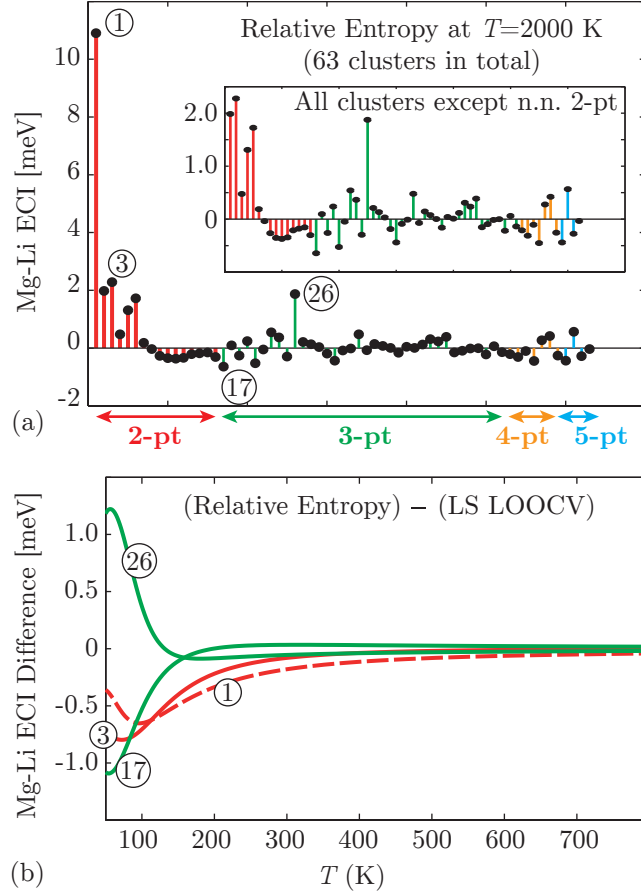


Figure 4.5: (Color) Mg-Li system. (a) ECI obtained with the relative entropy method at 2000 K. Compare with Fig. 4.1(a). (b) shows the difference in ECI versus T between least squares with LOOCV (LS LOOCV) and relative entropy. Circled numbers in (b) identify the clusters in (a) selected for plotting.

respectively. For 50 % composition Ref. [36] reports a transition between 300 and 450 K. Overall, discrepancies are due to different particular fitting methods used and different Monte Carlo techniques.

With the relative entropy method we predict 170, 214, and 240 K for 33 %, 50 %, and 66 % Mg, respectively. The difference between the two methods can be explained from the changing ECI around the transition temperatures as shown in Fig. 4.5(b). The ECI differences are an order of magnitude larger than for the

Si-Ge system.

Interestingly, the relative entropy method is in good agreement with the indicative experimental results in Ref. [37] according to which the order to disorder transitions across all compositions commence between 140 and 200 K. For the highest composition, i.e., 66 % Mg, we note that the experiments had highest error. For this largest composition, the relative entropy predicts a larger transition temperature than least squares LOOCV.

4.4 Conclusion

We have proposed a new paradigm based on a variational principle for obtaining thermodynamically relevant ECI. The principle is build upon the relative entropy which measures the information loss induced by replacing the true PDS with the candidate one. Through a series of suitable approximations we managed to bring the variational problem to a weighted least squares form enabling a practical solution.

We observed differences in predicted order to disorder transition temperatures of the Mg-Li alloy between the proposed method and least squares LOOCV. For the Si-Ge system, where a transition from two-phase coexistence to disorder was studied, we found the two methods to agree.

The main drawback of the method is that it requires a good set of clusters to start with. In this work, we relied on the methodologies implemented in ATAT (namely, least squares LOOCV to capture the ground states). Alternatively, we could have used the state-of-the-art method of compressive sensing. An interesting

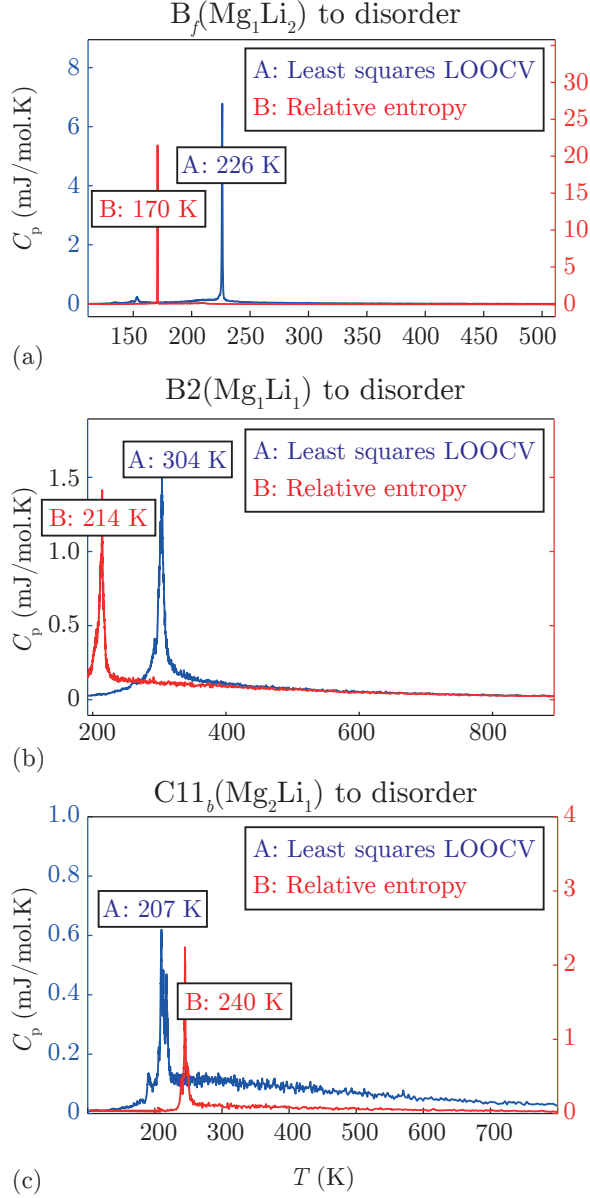


Figure 4.6: (Color) Prediction of the order to disorder phase transitions for Mg-Li bcc as indicated by “divergence” in the constant pressure heat capacity C_p at compositions (a) 33 %, (b) 50 %, and (c) 66 % Mg. In each plot least squares LOOCV (A: blue, associated with the left y -axis) are compared to relative entropy (B: red, associated with the right y -axis). The rectangular boxes on each peak report the temperature at which the maximum value of the heat capacity was obtained.

research question is to investigate the performance of compressive sensing when the error function is replaced by relative entropy. Under the approximations developed in this chapter this corresponds to changing from a Euclidean norm to a weighted norm when accounting for the error term.

CHAPTER 5

BAYESIAN UNCERTAINTY QUANTIFICATION IN USING SURROGATES FOR PREDICTING MATERIALS PROPERTIES

In this chapter we aim to develop a Bayesian framework for quantifying the uncertainty in alloy modeling when using fast parametrized surrogates in place of expensive computer codes as discussed in Chapter 1. We also show how the uncertainty is propagated to quantities predicted from the surrogate. We emphasize that the framework extends beyond cluster expansions (Chapter 2).

Generally, in the most typical setup, a surrogate is learned from some data set, e.g., quantum mechanical energies, and then used to predict some *quantity of interest* (QoI), which could be a ground state line, a phase transition, or some optimal structure (e.g., lowest thermal conductivity in the case where the data are instead thermal conductivities). Of course, the parametrization we choose for the surrogate depends on what data we see from the computer code. Since the code is expensive, we only see a limited amount of data. Furthermore, we require the surrogates to be computationally cheap. This means that, e.g., if the surrogate is represented by a set of basis functions, we are not in liberty to include an arbitrarily large number of such basis functions. The particular surrogate we consider later is such an example. These restrictions on the surrogate mean that, when it is parametrized, we do not know *a priori* the best parametrization. We have to learn it from a set of multiple candidate parametrizations, a pool of candidates, each candidate, from a Bayesian perspective, consistent with the observed limited amount of data. A single value of the QoI is computed from a single surrogate candidate. Since there may be multiple candidates, there may also be multiple values for the same QoI. Our uncertainty about the best surrogate candidate has

thus propagated to the QoI. This is the first source of uncertainty we aim to capture in the present work. From now on, we will simply say parametrization to mean surrogate parametrization/candidate.

Notice also that, the effect of limited data enters implicitly through our belief about the best parametrization pool to choose. For example, upon seeing data set \mathcal{D}_1 it might be that the pool of parametrizations t_1 is better than another pool t_2 . But if we now observe more data, it could very well be that our opinion is reversed, thus choosing t_2 over t_1 . The fewer data points we have, the worse, and generally larger, our pool of parametrizations consistent with the data will be, unless, of course, our prior belief is already sharply tuned to a good pool of parametrizations. However, this is rarely the case, and in by far the most cases we benefit from observing data. From this, it should be clear that the limited data plays a role in our knowledge about the best pool of parametrizations to use. The fact that we only see a limited amount of data therefore introduces a second source of uncertainty (not independent of the first one though) in the QoI, and we will be able to capture this as well.

As it is important, we reiterate that our developed methods are independent of the particular surrogate employed, but we will focus on the cluster expansion introduced in Chapter 2. To recap, the cluster expansion expands the alloy property in basis functions with associated expansion coefficients called ECI, which are obtained by fitting the cluster expansion to a data set. The cluster expansion surrogate is uniquely given once the ECI are specified, so we will consider a surrogate parametrization as being synonymous with the ECI. Although the cluster expansion is exact when untruncated, in practice one needs to make a truncation choice and estimate the ECI from a pool of parametrizations, as discussed in the

beginning of this introduction. We reiterate that this introduces uncertainties in the QoI predicted by the cluster expansion. Although an important question to ask, the sizes of these uncertainties have, strictly speaking, remained unknown until now. We employ a fully Bayesian approach to quantify the uncertainties.

We should mention that non-Bayesian methods have been applied in some works to quantify the uncertainty in QoIs, but we believe that they should be avoided for uncertainty propagation. In particular, there is no rigorous framework for propagating uncertainties through parametrized surrogates, such as the cluster expansion, to the QoI in non-Bayesian frameworks.^{53,54}

In beginning our Bayesian approach, we need to be clear about what we mean by *probability*. We interpret probability as a reasonable degree of belief^{55,56} as opposed to a frequency of some (hypothetical) long-run experiment. The sum and product rules of probability theory then tell us how to manipulate degrees of belief in a rigorous way. From this view of probability, Bayes theorem follows, and can be used to change our knowledge when observing new data in a given problem.^{55,57} This is collectively what is called Bayesian probability theory. We will use a Bayesian approach to introduce a model describing our belief about the best set of ECI with emphasis on sparsity. We include the sparsity feature because alloy properties are expected to be sparsely representable, based on physical arguments.¹⁵ A very successful sparse regression method, from the non-Bayesian literature, is LASSO, which is an L_1 -constrained least squares method. It can be shown that LASSO has a Bayesian interpretation. It corresponds to the posterior mode when the parameters to be learned have independent Laplace distributions as priors.⁵⁸ The above information about LASSO will be used to choose Laplace distributed priors in Section 5.5. The Bayesian posterior distribution (posterior) contains the

information needed to rigorously quantify the QoI uncertainties. In our case, the posterior attains a shape allowing it to be summarized via the 95 % highest posterior density confidence interval (HPD)—the smallest region containing at least 95 % of the posterior mass. We will reduce the effects of other uncertainties as much as possible, and discuss this as we go along.

We employ our framework to two real binary alloy systems. First, we consider bcc magnesium-lithium (Mg-Li) and let the QoI be its ground state line. Then, we turn to diamond silicon-germanium (Si-Ge) and present a computationally more involved example where the QoI is the transition temperature of the disordered to two-phase-coexistence at 50 % composition.

5.1 Chapter Outline

We start out with a general introduction to uncertainty quantification and present our framework in Sections 5.2 and 5.3. In Section 5.5, we present a Bayesian method for describing the ECI with emphasis on sparsity. The resulting posterior from the Bayesian treatment will not be in closed form for its intended use, so we show how samples are drawn from it using Markov chain Monte Carlo (MCMC)⁵⁹ methods in Section 5.6. Having developed the framework, we turn to case studies first discussing uncertainty quantification in the Mg-Li ground state line in Section 5.7, followed by the uncertainty quantification of an Si-Ge phase transition in Section 5.8. Results from these studies are presented in Section 5.10 and a corresponding discussion follows in Section 5.11. The chapter is concluded in Section 5.12.

5.2 Background

In this section we introduce the methods used to quantify the uncertainty in the QoI making no assumption about the form of parametrization of the response surface. Then, in the following section, we show how the cluster expansion makes this parametrization. Independent of the choice of surrogate model we will need data to make the best possible choice of parametrization. Therefore, we first discuss assumptions about the computer code used to obtain the data. Then, we introduce the central element in this work: an operator which acts on the surrogate to produce the QoI, and show how it is used to summarize the present uncertainty quantification task in a single equation, given certain assumptions.

The data acquisition takes place by supplying a set of alloy configurations as input to an expensive computer code, e.g., VASP or the large-scale atomic/molecular massively parallel simulator (LAMMPS)⁶⁰ and obtain a set of corresponding property values as output which collectively form the response surface. This could be quantum mechanical energies per atom or thermal conductivity, respectively. We view the computer code as a function $f(\cdot)$ mapping some input structure, with configuration denoted σ , to a response y . We do not know $f(\cdot)$ and we are most often not interested in it *per se*, but rather some function of this—the QoI. Therefore, we define an operator $I[\cdot]$ taking as input a response surface and returning the QoI which can generally be represented by a set of real numbers. As an example, we can let it return the structure at the global minimum of the surface:

$$I[f(\cdot)] = \arg \min_i f(\sigma_i),$$

or the ground state line of the system. The QoI can also be more complicated such as a transition temperature. The computer code has inherent approximations. For example, VASP approximates the exchange-correlation term in density functional

theory (DFT), implements a particular k -point integration scheme,⁴⁷ and a pseudopotential to approximate the true potential,⁴⁴ etc., all inducing uncertainties in the output. We assume, however, that such uncertainties are small when compared to those arising from not knowing how to choose the surrogate parametrization, and from having observed only a limited amount of data. The presence of code uncertainties means that we do not actually observe the theoretical $f(\boldsymbol{\sigma}_i)$ of structure i . Nevertheless, we will make the assumption that we do observe $f(\cdot)$, but with added Gaussian measurement noise. Call this noisy version of the code y_i . We show in Section 5.5 how the noise can be estimated in the current framework. Incidentally, this does not mean we commit to the noise source necessarily being Gaussian itself, but rather that our estimates depend on the first and second moments of the noise.⁶¹

In its present form, the response surface $f(\cdot)$ is very abstract. To describe it in more concrete terms we replace it with a parametrized surrogate model. The surrogate describes $f(\cdot)$ using a set of parameters θ which can be discrete, continuous, or a mix of both. E.g., if expanding $f(\cdot)$ in a basis, θ identifies the set of basis functions and corresponding expansion coefficients. Without loss of generality, the surrogate is synonymous with θ so we will not differentiate these in the text. We assume that, for some parameter choice, the surrogate captures the response surface *exactly*. If this was not the case we would not be guaranteed that some parameter in the set could, even in principle, capture the surface. This would lead to further uncertainties which would, in general, be very difficult to quantify. The physical implication of our assumption is that the parametrization should account for $f(\cdot)$ being periodic, satisfying certain symmetries, and/or its differentiability, etc. The cluster expansion, to be introduced formally in Section 5.4, does capture the response surface for some parameter and is thus well suited for the present frame-

work. We will assume to be in a situation where the parameter which captures the surface exactly is computationally intractable to obtain. The central question we address in this work is then: of the remaining computationally feasible parameter set, which parameter should be chosen to describe the surface? There is no easy answer and indeed, most of the cluster expansion literature so far has been on how to obtain a “best” set of ECI, but with no rigorous treatment of the uncertainties in the QoI.

The focus of this chapter is not so much on the ECI themselves, or equivalently, on the response surface itself, but rather on the QoI predicted from these. Our approach will be to describe the ECI with a Bayesian approach but different than Ref. [62] in the prior specification. We reiterate that the choice of the best parameter will implicitly depend on the data set size. This is because we “pin down” the response surface more as the data set increases. Thus, incompatible surrogates, recognized by the Bayesian method, are removed leaving a smaller set of possible ECI to choose from and thus less variability.

5.3 Bayesian Uncertainty Quantification

We now develop the Bayesian tools needed for the proposed uncertainty quantification of the QoI. Before seeing any data, we have some belief of possible surrogates best describing the QoI. We associate with our belief a prior probability distribution $p(\theta)$ (prior). Since we are ultimately interested in the QoI rather than θ let I denote a realization of the operator (which outputs the QoI) $I[f(\cdot; \theta)]$. We are interested in $p(I)$ which is our prior belief about the QoI. Of course, since I derives directly from θ via the application of the associated operator $I[f(\cdot; \theta)]$, $p(I)$

has to depend on $p(\theta)$. This dependence, however, is not trivial unless we make the assumption that the operator itself does not introduce further uncertainties in I . We shall later see a case where this could happen. With this assumption, consider first a single surrogate θ_1 and its associated predicted QoI I_1 . If we are certain about θ_1 we necessarily have to be certain about I_1 . However, the contrary statement is not true. To see this, note that multiple θ 's might lead to the same value of I_1 . For example, if I_1 is a transition temperature it is conceivable that two different surrogates θ_1 and θ_2 happen to yield the same transition. With the help of the Dirac δ -function the functional dependence between θ_1 and any value of I can then be summarized in the form:

$$p(I) = p(\theta_1) \delta(I[f(\cdot; \theta_1)] - I). \quad (5.1)$$

But since the uncertainty in I can also be affected from other surrogates than θ_1 the general form of Eq. (5.1), accounting for all possible surrogates, becomes:^{63,64}

$$p(I) = \int d\theta p(\theta) \delta(I[f(\cdot; \theta)] - I), \quad (5.2)$$

which is then our prior belief in the values taken on by the QoI. The prior belief can, most typically, be much improved by incorporating actual observations of the response surface into our state of knowledge. Therefore, we now run the computer code for a set of structures and obtain a data set $\mathcal{D} = \{(\boldsymbol{\sigma}_i, y_i)\}_{i=1}^n$. The likelihood, denoted $\mathcal{L}(\mathcal{D}|\theta, \cdot)$, quantifies how compatible a surrogate θ is with \mathcal{D} and contains all information about θ present in \mathcal{D} . Bayes rule provides a clear prescription for how we should update our belief about the best surrogate upon observing \mathcal{D} :

$$p(\theta|\mathcal{D}, \cdot) \propto p(\theta) \mathcal{L}(\mathcal{D}|\theta, \cdot). \quad (5.3)$$

This change of knowledge propagates to the QoI as:

$$p(I|\mathcal{D}, \cdot) = \int d\theta \delta(I[f(\cdot; \theta)] - I) p(\theta|\mathcal{D}, \cdot), \quad (5.4)$$

with the same rationale as that preceding Eq. (5.2). Eq. (5.4) is the posterior on the QoI. Plotting its HPD provides an answer to our opening question of how to quantify the uncertainty in predicted QoIs from surrogates used for alloy modeling. However, solving the equation is not trivial in general.

Note that a very large spread in $p(I|\mathcal{D})$ informs us that we are not capturing the response surface well with our parametrization method and data set. Different applications tolerate different spreads of course, but the purpose in the present work is to get a feel for the sizes of these uncertainties when applying the method to two real alloys using a popular parametrization method to be discussed next.

At this point we can relate our framework to the thorough work of Ref. [65], referred to as KH in what follows. They develop a general Bayesian framework for matching θ to a set of observations. This is known as calibration, and they use Gaussian processes to do so.^{66,67} In our work, we will use a basis function approach,⁶⁸ the cluster expansion, to be discussed in the following section. We consider KH an alternative approach to ours, as far as the calibration aspect is concerned. But we emphasize that our work also provides a rigorous way to propagate the distribution of the calibrated model to the QoI. This allows us to quantify the uncertainty of the QoI. Surely, it would be interesting to compare KH to our approach at the calibration level. For a discussion of, and references to, other Bayesian alternatives please see Ref. [65].

5.4 Parametrizing the Response Surface with the Cluster Expansion

The cluster expansion, theoretically introduced in Chapter 2, provides a possible parametrization method of the response surface $f(\cdot)$ discussed in the previous section by expanding it in correlation functions constructed from cluster orbits. Together with the ECI, these functions play the role of θ in Eq. (5.4), and we recall that, for *some* set of ECI, the cluster expansion captures the true response surface exactly. This makes the expansion well suited for estimating the particular uncertainties in the QoI we aim to capture from truncating the expansion. For notational convenience, we introduce a vector $\boldsymbol{\gamma}$ to carry the information about which clusters have been retained in the truncation. Furthermore, for clearer notation later, we now let $\boldsymbol{\beta}_{\boldsymbol{\gamma}}$ be the ECI. The vector $\boldsymbol{\gamma}$ contains a one/zero in its j th spot if basis function j is included/excluded in the expansion and if $\gamma_j = 0$ then so is $\beta_{\gamma,j}$. Thus, the cluster expansion parametrizes $f(\cdot)$ as $f(\cdot; \boldsymbol{\gamma}, \boldsymbol{\beta}_{\boldsymbol{\gamma}})$.

The difference between the observed data and the cluster expansion parametrization is the error. We assume this error to be Gaussian. In view of later developments we cast this in the notationally simpler form, notationally slightly different from that discussed in Chapter 2, because of the added information $\boldsymbol{\gamma}$:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta}_{\boldsymbol{\gamma}} + \boldsymbol{\epsilon}, \quad (5.5)$$

where we recall that \mathbf{X} is the design matrix. The i th element in $\boldsymbol{\epsilon}$ is the Gaussian error of observation i . We assume the error identical for all observations, with mean zero and constant variance δ^2 . The framework presented here can be used to estimate δ , and hence the assumed noise, from the data.

5.5 Bayesian Treatment of the Response Surface Parametrization

The cluster expansion parametrizes the response surface, but how do we choose the best set of parameters, i.e., the cluster orbits and corresponding ECI, in light of some data set \mathcal{D} ? In this section we introduce a Bayesian model on the ECI from Section 5.4. In particular, we introduce the prior, likelihood, and posterior on $\theta = (\gamma, \beta_\gamma)$ in Eq. (5.5). As discussed in the chapter introduction, we will base the Bayesian model on LASSO through our prior beliefs. The price of this will be a posterior distribution that is not in closed form. However, we show later how to use sampling methods to still solve Eq. (5.4). The prior is the joint distribution of three parameters $p(\beta_\gamma, \delta^2, k|\cdot)$, where k is the number of basis function to be used, that is, it is the number of columns in the design matrix. We introduce this parameter in Eq. (5.7) below. We shall assume prior logical independence between the parameters. In other words, knowledge about β_γ , before seeing any data, does not provide evidence about the value of the data noise or k . The logical independence allows us to separate the joint prior into independent factors $p(\beta_\gamma|\cdot)p(\delta^2)p(k|\cdot)$. We therefore proceed by specifying the prior one parameter at a time. A Laplace distribution is assigned to our prior belief in the magnitude of the ECI:

$$p(\beta_{\gamma,j}|\tau, \gamma) = \frac{1}{2\tau} \exp\left(-\frac{|\beta_{\gamma,j}|}{\tau}\right), \quad (5.6)$$

where $\beta_{\gamma,j}$ is the j th coefficient of β_γ and τ is a *hyperparameter*. It is thought of as a shrinkage tuning parameter for $\beta_{\gamma,j}$ with $j \in \gamma$. $|\cdot|$ is the absolute value. The amount of shrinkage is unknown *a priori* so will be described with a non-

informative hyperprior:

$$p(\tau) \propto \tau^{-1}.$$

The data noise is also unknown leaving us to choose a non-informative hyperprior as well:

$$p(\delta^2) \propto \delta^{-2}.$$

We can enforce sparsity in how many cluster orbits best describe the data. In other words, we can shrink the number of columns needed in the design matrix. A sub-model of size k is a particular set of active cluster orbits. To favor lower values of k , the prior probability of a sub-model follows the right truncated Poisson distribution at k_{\max} :

$$p(k|\lambda) = \frac{\exp(-\lambda) \lambda^k}{C k!}, \quad (5.7)$$

with C a normalization constant, $k = 1, \dots, k_{\max}$, and λ , a hyperparameter, can be thought of as a second shrinkage parameter inducing sub-models with less clusters. Equal probability is assigned to each sub-model having the same value of k . The hyperprior on λ is chosen as non-informative:

$$p(\lambda) \propto \lambda^{-1}.$$

Consider now the likelihood $\mathcal{L}(\mathbf{y}|\boldsymbol{\beta}_\gamma, \gamma, \cdot)$. Given the linear model in Eq. (5.5) and our belief about the error, we write it as:

$$\mathcal{L}(\mathbf{y}|\boldsymbol{\beta}_\gamma, \gamma, \cdot) = (\delta^2)^{-n/2} \exp\left(-\frac{\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}_\gamma\|_2^2}{2\delta^2}\right), \quad (5.8)$$

where $\|\cdot\|_t$ is the t -norm. Incidentally, although Eq. (5.8) contains no explicit correlations this does not mean we cannot capture such effects in the data.⁵⁵ For this problem, the hyperparameters τ , λ , and δ can be integrated out, and thus

implicitly dealt with. This serves us a more robust fitting method. Combining the above equations one can show that the posterior takes the form:

$$p(\boldsymbol{\gamma}, \boldsymbol{\beta}_{\boldsymbol{\gamma}} | D) \propto \Gamma(k) B(k, k_{\max} - k + 1) \|\boldsymbol{\beta}_{\boldsymbol{\gamma}}\|_1^{-k} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}_{\boldsymbol{\gamma}}\|_2^{-n}, \quad (5.9)$$

where k is the number of active components in $\boldsymbol{\gamma}$, $\Gamma(\cdot)$ is the Gamma function, $B(\cdot, \cdot)$ is the Beta function, and the reason for keeping k -dependent prefactors is related to how we are going to draw samples from this which is the topic of the next section. Eq. (5.9) is our posterior belief about the cluster expansion truncation and ECI values best describing the data at hand. For more details on this particular model, please see Ref. [69].

5.6 Sampling from the Posterior: Reversible Jump Markov Chain Monte Carlo

Our ultimate goal is to sample from the posterior on the QoI [Eq. (5.4)] using the posterior on possible surrogates [Eq. (5.9)]. The integral in Eq. (5.4), however, can not be carried out analytically. Therefore, we turn to numerical sampling methods. The idea is to sample from the surrogate posterior in regions where it takes large values (i.e., we sample surrogates very likely to be true). Each surrogate θ is then passed through $I[f(\cdot; \theta)]$ finally providing a sample from the posterior on the QoI $p(I | \mathcal{D})$.

We will use an MCMC sampling technique to sample from the surrogate posterior. If we worked with a fixed model dimension k in Eq. (5.9), standard MCMC would suffice, see Chapter 3. However, we consider the possibility that k cluster orbits might not describe the data better than, say, k' orbits. This requires

a dimension-changing scheme leaving us to choose the RJMCMC method introduced, and discussed theoretically, in Section 3.2. As in Section 3.2, we consider the three move types: update the ECI in the current model, increment the model dimension (birth), i.e., add another correlation function and associated ECI, or to decrement the current model dimension (death). If updating the ECI a zero-mean Gaussian distribution is chosen with variance α^2 to perturb the state. We will consider only the models $k + 1$ (birth move) and $k - 1$ (death move) as valid candidates if currently at k active basis functions. The reason for not choosing, say, $k + 2$ is because we assume that $k + 1$ is always the smallest perturbation of the state possible (although we have no notion of distance in this more complicated state space) hence leading to higher acceptance rates. Here, we will simply provide the jump probabilities for our case, but more information is available in Ref. [69].

A proposal step size is drawn, called u , from a Gaussian. If we just want to update the ECI vector in the current model k from β_γ to β'_γ we select a random element, say the j th, from β_γ and propose the update $\beta'_{\gamma,j} = \beta_{\gamma,j} + u$. We accept this change with probability (see Section 3.2.3)

$$\min \left\{ \left(\frac{\|\beta'_\gamma\|_1}{\|\beta_\gamma\|_1} \right)^{-k} \left(\frac{\|\mathbf{y} - \mathbf{X}\beta'\|_2}{\|\mathbf{y} - \mathbf{X}\beta\|_2} \right)^{-n}, 1 \right\},$$

which we recognize to be a standard Metropolis acceptance probability which does not explicitly contain the proposal distribution (Gaussian) because this distribution is symmetric.¹⁸ If we instead propose a birth move to increment the model dimension we select an inactive correlation function and propose to revive it simultaneously setting the associated ECI value equal to the proposal u . This move is accepted with probability (see Section 3.2.4)

$$\min \left\{ \frac{k^2}{k_{\max} - k} \frac{\|\beta'\|_1^{-(k+1)} \|\mathbf{y} - \mathbf{X}\beta'\|_2^{-n}}{\|\beta_1\|^{-k} \|\mathbf{y} - \mathbf{X}\beta\|_2^{-n}} \frac{p(\gamma' \rightarrow \gamma)}{p(\gamma \rightarrow \gamma')} \mathcal{N}(u; 0, \alpha^2)^{-1}, 1 \right\}, \quad (5.10)$$

where $\mathcal{N}(u; 0, \alpha^2)$ is the Gaussian proposal distribution evaluated at u . Notice how

now the proposal distribution enters explicitly in the third factor this time. The quantity $p(a \rightarrow b)$ is the probability of ending in state b given that you started in a , and is provided in Ref. [69]. A death move is just the inverse of Eq. (5.10), see Section 3.2.5. Note that the only adjustable parameter in the RJMCMC method is α . A value too small will not allow good exploration of the space and a value too large leads to many rejected steps. We use a value leading to an approximately 30 % total acceptance rate. This is close to the widely accepted optimal acceptance probability of 23.4 %.⁷⁰ Interestingly, this parameter can be estimated from the data themselves leaving *no* adjustable parameters.

Having now coupled the Bayesian LASSO (BL) method in Section 5.5 with the RJMCMC sampling scheme in the present section we will refer to the method as BL-RJMCMC in the rest of the chapter. We are now in place to sample from Eq. (5.4). In the next two sections we discuss two particular alloy quantities of interest and how to apply Eq. (5.4).

5.7 Uncertainty Quantification of the Magnesium-Lithium Ground State Line

In this section we give an example of using Eq. (5.4) when I represents the ground state line of formation energies in the bcc Mg-Li system when using the cluster expansion as a parametrization so $\theta = (\gamma, \sigma_\gamma)$. This is equivalent to quantifying the uncertainty in the $T = 0$ K phase diagram. The response surface is internal quantum mechanical energies of the alloy. The operator $I[f(\cdot; \theta)]$ then takes a set of cluster orbits and associated ECI, i.e., a surrogate, and produces the ground state line as output.

We can describe in more detail what $I[\cdot]$ does in the present problem. First, change notation from the general response y_i in Section 5.2 to the more expressive $E(\boldsymbol{\sigma}_i)$ with $\boldsymbol{\sigma}_i$ being the configuration of structure i . Let $\Delta E(\boldsymbol{\sigma}_i)$ be the formation energy of structure i defined below. Take $E(\boldsymbol{\sigma}_{\text{Mg}})$ and $E(\boldsymbol{\sigma}_{\text{Li}})$ to be the energies of the structures containing only Mg and only Li atoms, respectively. If structure i contains a fraction of Mg atoms $x_{\text{Mg}}^{(i)}$, then:

$$\Delta E(\boldsymbol{\sigma}_i) \equiv E(\boldsymbol{\sigma}_i) - \left(E(\boldsymbol{\sigma}_{\text{Mg}})x_{\text{Mg}}^{(i)} + E(\boldsymbol{\sigma}_{\text{Li}})(1 - x_{\text{Mg}}^{(i)}) \right), \quad (5.11)$$

and $I[\cdot]$ then forms the ground state line from the formation energies of all structures, see, e.g., Ref. [71].

Our goal is to answer how certain we are about the predicted ground states of Mg-Li and implicitly how well the cluster expansion captures the energies. Practically speaking what we need is to draw ground state line samples from Eq. (5.4) and plot the HPD. This can be done as follows. Assume we are given an equilibrated RJMCMC chain on the posterior. Then, we draw samples from the posterior, which corresponds to drawing different possible parametrizations of the response surface. For each parametrization we compute the formation energies of the data and plot the ground state line. This is repeated for as many samples from the chain as possible and the HPD is computed. In Section 5.10 we show the result of this.

5.8 Uncertainty Quantification of Disorder to Two-Phase-Coexistence Phase Transition

Next, consider another QoI: A disorder to two-phase-coexistence phase transition in the diamond Si-Ge alloy at 50 % composition. Computationally, the phase transition is found by starting the alloy at a high temperature which for Si-Ge could be 2000 K. The temperature has to be large enough for the entropic energy in the free energy to dominate. Then, the temperature is gradually lowered while monitoring the specific heat at constant pressure:

$$C_p(T) = \beta^2 k_B \langle (E - \langle E \rangle)^2 \rangle, \quad (5.12)$$

where β is the inverse temperature defined in terms of the temperature T as $1/k_B T$, k_B is the Boltzmann constant, $\langle \cdot \rangle$ denotes an average, and E is the internal energy of the alloy. Lowering the temperature decreases the entropic energy allowing the configurational energy to become comparatively stronger finally demanding a certain ordering of the atoms. In a two-phase-coexistence phase the atoms do not want to mix on the lattice but stays separated into pure forms. A peak (divergence, in the limit of an infinite lattice) in the specific heat signals the phase transition.

The procedure for quantifying the uncertainty is similar to that discussed in Section 5.7 but now with I representing a phase transition temperature. This demonstrates the generality of Eq. (5.4). Parametrizations of the response surface are again sampled from the RJMCMC chain and for each sample the phase transition is computed. This, in turn, provides a sample from Eq. (5.4). Plotting the HPD of multiple samples quantifies the uncertainty in the phase transition.

Let us now discuss the numerical method used to compute the phase transition. We used an adaptive sequential Monte Carlo (ASMC) method coupled via in-house

code to the MCMC library of ATAT. This is not to be confused with the other MCMC method already used to sample the posterior. The ASMC code works as follows. A set of 64 so-called particles is started at a high temperature of 2000 K. Each particle is an initial state of the system: a $30 \times 30 \times 30$ Monte Carlo cell with periodic boundary conditions. Using double-spin-flip dynamics they are independently evolved via the ASMC method to 50 K. The step size is adaptive, but we maximally allow a 10 K jump to ensure we do not miss a transition. At each step (including the initialization) 100 flips per lattice site were performed. For more details please see Ref. [72]. There are uncertainties associated with the number of particles and the seed used for the internal pseudo-random-number generator. Thus, the present situation is a case where $I[\cdot]$ could introduce further uncertainties. If we allow this we are not only capturing the uncertainties from the truncated cluster expansion and limited data but also from the ASMC method. To avoid the ASMC uncertainties, we use the same number of particles with each sample and the same seed.

5.9 Data

We used the quantum mechanical energies of both alloys computed via VASP in Ref. [72]. This provided a set of responses. ATAT was then used to generate the design matrix \mathbf{X} . We must make a choice about what cluster orbits (basis functions) can be considered by the Bayesian model. The RJMCMC method will then select a subset of these. For bcc Mg-Li we generated 16 two-point (2-pt) and 4 3-pt cluster orbits besides the empty and 1-pt cluster orbits contained in both systems. The maximum cluster orbit spatial extent was 20 Å. Generating more cluster orbits, especially for this system, leads to better predictions. But in

the present work we want to demonstrate features of the uncertainty in the QoI induced by the cluster orbit and ECI choice, and not predict some alloy property with high precision. Therefore, it is essentially irrelevant what exact choice of cluster orbits we make. Of course, if the uncertainty is too large it can be reduced by including larger cluster orbits but this was not needed here. For the Si-Ge diamond system we generated 9 2-pt and 6 3-pt cluster orbits with a maximum spatial extent of 8 Å.

In Fig. 6.6 we show the two data sets. Fig. 6.6(a) contains the formation energies of bcc Mg-Li versus Mg concentration. The Mg-Li data set is in general agreement with a similar one obtained in Ref. [36]. Discrepancies are due to different VASP settings but more importantly the difference in structures considered. In particular, ATAT was used in our case to generate the data set. In Fig. 6.6(b) the formation energies of diamond Si-Ge versus Ge concentration. The formation energy is a measure of how well the pure alloys mix. A negative value indicates affinity towards mixing on the lattice, this is the case for Mg-Li, and the opposite is true for positive values, seen for the Si-Ge system.

5.10 Results

For both systems we needed to set up an RJMCMC chain on Eq. (5.4) and equilibrate it in order to draw posterior samples. In Fig. 5.2 we provide details of these simulations versus MCMC step number for both systems. For Mg-Li(Si-Ge) we ran the chain for 9(8) million steps burning the first 4 million samples. This took 6 hours to run on a single core (a parallelization of the MCMC chain is possible).⁷³ Fig. 5.2(a) shows the model complexity during the run. The sparse nature of the

BL-RJCMC is apparent for both systems in that most time is spent exploring surrogates with around a third of the complexity maximally possible, which is 22 for Mg-Li and 16 for Si-Ge. This is because of our prior choices in Eqs. (5.6) and (5.7). In comparison, least squares and LASSO with 10-fold cross-validation (LASSO-CV) have high complexities of 22 and 20, respectively, for Mg-Li and 16 and 10, respectively, for Si-Ge. The zero-point cluster orbit was not used in any fitting process (as the noise is assumed independent and identically distributed). Therefore, 21 coefficients are shown for Mg-Li and 16 for Si-Ge with the understanding that the ECI not shown is common to all methods and equates to the mean of the data. Fig. 5.2(b) estimates the data noise δ entering the model through ϵ in Eq. (5.5). The noise agrees, for both systems, with the order of magnitude expected from DFT calculations which is meV's per atom.⁷⁴ For the bcc Mg-Li system we use the approach discussed in Section 5.7 to quantify the uncertainty in the ground state line from the cluster expansion truncation and the fact that we have limited data. Please find this result in Fig. 5.3(a) where black dots are the VASP formation energies along with the black line being the VASP ground state line. The median of the BL-RJCMC posterior samples is shown as the full blue line and the shaded area encloses the HPD.

For the diamond Si-Ge system we use the sampled ECI to estimate the uncertainty in the disorder to two-phase-coexistence transition temperature as discussed in Section 5.8. Please refer to Fig. 5.3(b) for the results. Five coefficients were sampled from the RJCMC chain and the transition temperature run from each set of coefficients, based on identical initial conditions, are shown in the main figure. The inset summarizes this by plotting the HPD of Eq. (5.4). The median is plotted as a vertical black line.

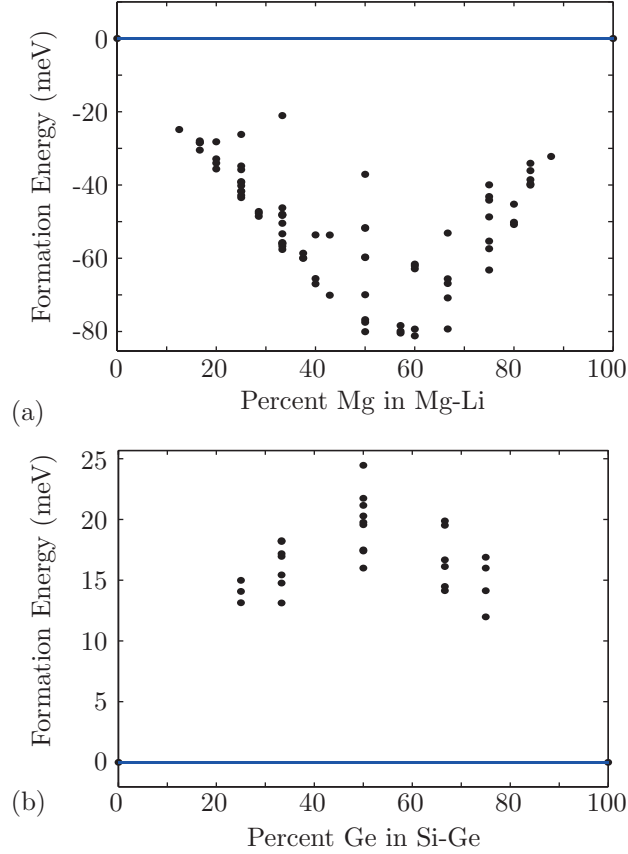


Figure 5.1: Data sets obtained via VASP. (a) Shows the quantum mechanical formation energies of 82 Mg-Li structures and (b) for 38 Si-Ge structures. The line of zero formation energy is shown as well.

5.11 Discussion

Most work so far in the field has focused on obtaining the best surrogate surface itself. Therefore, we provide a comparison of the most likely surrogates predicted by BL-RJMCMC to the (point) surrogates predicted by least squares and LASSO-CV. See Fig. 5.5. The BL-RJMCMC chain is started with just one cluster orbit active and an ECI value of 1 eV. From this, it finds an almost identical result as LASSO-CV, but it is even sparser. Interestingly, in Fig. 5.5(a) BL-RJMCMC predicts a strong signal from the nearest neighbor 2-pt cluster orbit but nothing from

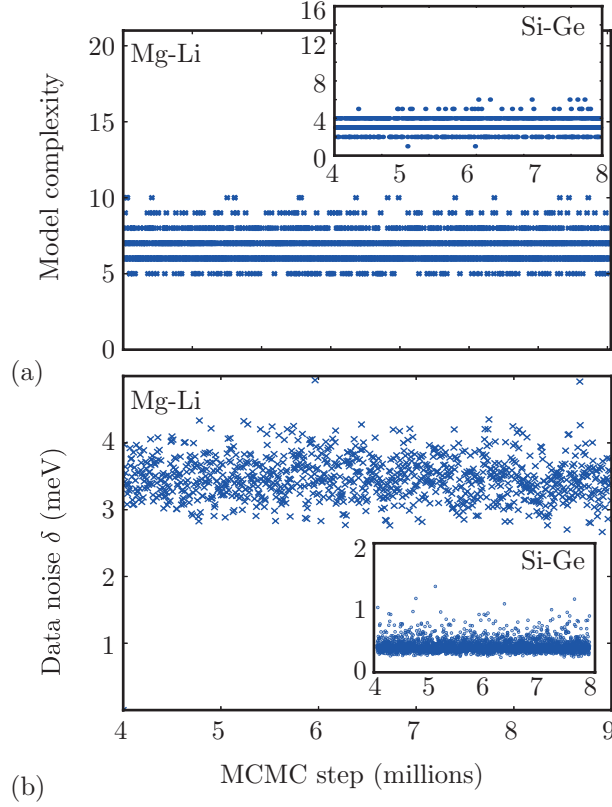


Figure 5.2: BL-RJMCMC simulation details for the Mg-Li system. In (a) we monitor the model complexity, i.e., k in Eq. (5.9), where the maximum complexity k_{\max} is 21 (not counting the empty cluster orbit). In (b) we show the estimated data noise ϵ in Eq. (5.5). Insets show similar results for the diamond Si-Ge system where, in (a), the maximum complexity is 16 (not counting the empty cluster orbit).

the largest 2-pt cluster orbits. It then picks up a signal again at the two shortest 3-pt cluster orbits, in agreement with LASSO-CV, but predicts no signal from the largest 3-pt cluster orbit as opposed to LASSO-CV. With the smaller cluster orbits carrying most of the signal, the BL-RJMCMC provides a more physically satisfying result compared to both least squares and LASSO-CV.¹⁰ In Fig. 5.5(b) BL-RJMCMC picks up some strong 2-pt cluster orbit signals but no 3-pt cluster orbit signals, among the cluster orbits chosen. In the latter case both least squares




Figure 5.3: (Color) Uncertainty in the ground state line of bcc Mg-Li. Black dots are the formation energies obtained from VASP and their ground state line is the black line. Red dash-dot and green dash lines are LASSO-CV and least squares, respectively. The BL-RJMCMC result is plotted as a blue full line (median of the posterior) with the shaded areas representing the HPD.




Figure 5.4: (Color) Uncertainties in the disorder to two-phase-coexistence of Si-Ge at 50 % composition identified by a peak in Eq. (5.12). Five different runs are shown each in a different color. The inset summarizes the main figure via the posterior median as a black vertical line surrounded by the HPD as the shaded area.

and LASSO-CV predict signals for these large cluster orbits. An interesting observation can be made for the third ECI where least squares and LASSO-CV disagree

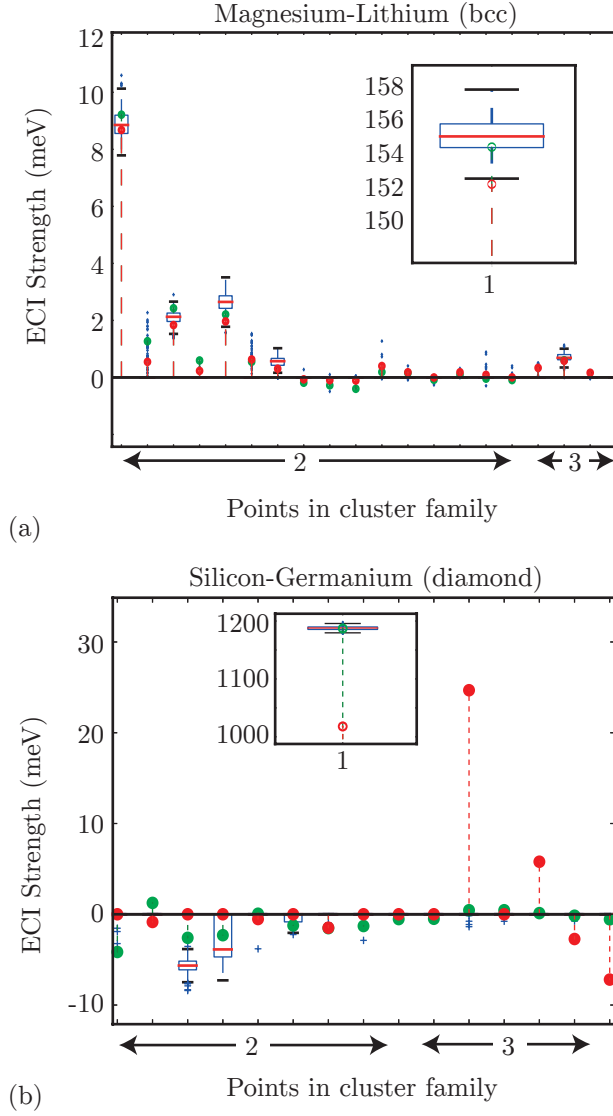


Figure 5.5: (Color) ECI obtained via the BL-RJMCMC method (box plots), least squares (red stems), and LASSO-CV (green stems) plotted for the (a) Mg-Li bcc and (b) Si-Ge diamond system. Insets show the single site cluster orbit for increased clarity.

on its sign. The ECI sign is indirectly coupled to the underlying physics: A negative value indicates that this 2-pt cluster orbit prefers the atoms in the cluster to be of the same type. The opposite is true for a positive ECI value. In other words therefore, whenever least squares and LASSO-CV disagree on the underly-

ing *physics* about preferred atomic ordering, BL-RJMCMC predicts that the data support no evidence of drawing a conclusion about this.

The posterior samples from the BL-RJMCMC chain are used in Eq. (5.4) to draw samples from $p(I|\mathcal{D})$. In the first alloy, bcc Mg-Li, the samples are ground state lines. In Fig. 5.3 we summarize these samples using the HPD. Uncertainties on the order of 10 meV are seen around the median in the vicinity of 50 % composition. The fact that the relative uncertainties are around 12 % provides evidence that the cluster expansion is capturing the response surface. Whether this is tolerable or not cannot be answered here. For pure comparison we have shown the point estimated ground state lines of LASSO-CV (green dash) and least squares (red dot-dash) as well. Notice that the black line from our limited data set is not necessarily the true ground state line. Strictly speaking, the true line is computed only from an *infinite* number of structures. Therefore, the HPD is not required to contain this line.

In the second alloy the BL-RJMCMC samples were phase transitions. The HPD of the two-phase coexistence transitions in Si-Ge varies by 19.7 K based on five posterior samples. The median transition is at 333 K in good agreement with the result in Ref. [9]. Discrepancies are due to a different choice of cluster orbit set and MCMC method used. We predict a relative uncertainty in the transition around 6 % which, at first, seem very small. However, the quantum mechanical energy of Si-Ge is typically well described by smaller cluster orbits of two and three points compared to, e.g., Mg-Li typically requiring up to five point cluster orbits. Indeed, the difference in uncertainties between the two systems here are in line with what we found in Ref. [72]: there was a small difference in using least squares and relative entropy on Si-Ge but a large difference for Mg-Li. Therefore,

we attribute the smaller uncertainty in Si-Ge to the fact that we have captured the response surface better than that of Mg-Li.

If one wishes to reduce the uncertainty further in either system, more cluster orbits can be included. Also, more data should be obtained. This would move the posterior of θ closer to the true response surface and allow us to capture the response surface better. Improving the Bayesian fitting method is another alternative. In particular, if one has further prior knowledge about the system (e.g., one knows the data noise to be some particular value, or that the data are correlated in some way) this should be included in the prior specification and will lead to a better fit of the response surface. This can readily be incorporated into the present framework.

5.12 Conclusion

We have presented a rigorous Bayesian framework for quantifying the uncertainty in predicted quantities of interests from the cluster expansion arising from two sources: lack of knowledge about the best truncation choice and corresponding ECI, and also from having observed a limited amount of data. We presented a framework for carrying out this quantification in general but considered two particular quantities of interest in two different binary alloys: bcc Mg-Li and diamond Si-Ge. The main idea was to set up a Bayesian posterior on the possible truncations and associated ECI consistent with the observed data. Samples were then drawn, via an RJMCMC sampling scheme, from the posterior and propagated through an operator $I[\cdot]$ which provided the quantity of interest as output when given a particular truncation of the cluster expansion, i.e., a particular surrogate.

Essentially, the operator converted uncertainties in the cluster expansion to uncertainties in the quantity of interest. We applied our framework to two alloys. First, we demonstrated the uncertainty in the predicted ground state line of formation energies in bcc Mg-Li, and concluded that there was up to 12 % relative uncertainty when plotting the 95 % HPD. In the second alloy considered, diamond Si-Ge, we provided uncertainties in the phase transition temperature at 50 % composition from disorder to two-phase-coexistence. We found a 6 % relative uncertainty when plotting the 95 % HPD.

Finally, we address a potential issue with the BL-RJCMC method. An important question is how to choose the starting guess. The BL-RJCMC method *per se* is not built for finding global extrema. In this work, however, we found that a great variety of starting guesses lead to the same conclusions. In other words, the response surfaces were adequately smooth for our purpose. This might not be true in general. An interesting research topic is to couple the BL-RJCMC presented here with some annealing scheme⁷⁵ or other sampling scheme^{73,76} to get closer to a globally relevant part of the posterior. This would make the method much more robust to the shape of the response surface.

CHAPTER 6

USING SURROGATES FOR MATERIALS DESIGN

We discussed the motivation for surrogate models in Chapter 1 and introduced a particular surrogate model, the cluster expansion, useful for alloys, in Chapter 2. To recap, the cluster expansion expands the response surface in basis functions with associated expansion coefficients called ECI. So far in this thesis we have not discussed the computation of the clusters themselves. We have tacitly assumed that the clusters were computable in order to construct the design matrix. For most bulk systems this is the case. In fact, generally speaking, for bulk systems, we do not have to worry about the computational cost of obtaining the clusters. This is unfortunately not so when faced with low-dimensional systems. Of course, predicting structures possessing lower dimensional geometries is just as, if not more, interesting as predicting bulk structures, so we need surrogates which can successfully describe such systems as well. This motivated the work in this chapter where we set out to modify the traditional cluster expansion to work more fluently for structures of arbitrary geometries.

As we have seen in the previous chapters that to perform a cluster expansion the open source software package ATAT can be employed. Non-bulk systems such as surfaces can be cluster expanded with ATAT,⁷⁷ but consider now using ATAT to cluster expand *any* geometry—even one fully void of periodicity and/or large in size. The computation of the clusters in ATAT and of the associated correlation functions for arbitrary structure shapes easily becomes prohibitively expensive to compute for large unit cells. One of the main reasons is that the number of clusters to compute, as will be shown in Section 6.2, grows extremely fast with the unit cell size.

In this chapter, we introduce the *cluster expansion ghost lattice method* (CE-GLM). The CE-GLM is an idea, which enables, via minor modifications, a cluster expansion software package to cluster expand structures of *arbitrary* shapes with the same computational effort as the corresponding bulk system. Its main strength is in the ease of implementation in existing cluster expansion code. What makes the method different than a regular cluster expansion with vacuum atom types, is in the way it handles large low-symmetry structures by using a practically convenient *ad hoc* grouping of the clusters. The work of Ref. [78] provided evidence that *some ad hoc* grouping of the clusters should be expected to be necessary for low-symmetry systems since the number of unknowns otherwise becomes insurmountable. We propose to make the *ad hoc* choice unique and transferable.

In order to demonstrate that our choice of using bulk clusters in non-bulk low-symmetry systems can be useful, we set out to reproduce the work of Ref. [78], but by using the CE-GLM implemented in ATAT. If successful, this will provide a strong case for the usefulness of the CE-GLM method.

6.1 Chapter Outline

We start with a discussion of how the CE-GLM extends the cluster expansion of Chapter 2 in Section 6.2. Then, in Section 6.3.1, we provide an overview of the nanowire problem considered, followed by details on how the nanowire data set was obtained in Section 6.3.2. In Section 6.3.3 we discuss how the thermal conductivities of the nanowires were computed. The results of this chapter are provided in Section 6.4 before concluding in Section 6.5.

6.2 The Cluster Expansion Ghost Lattice Method

The CE-GLM removes the computational overhead of cluster expanding non-bulk possibly low-symmetry systems compared to bulk structures. The cluster expansion of 2D sheets, nanowires, spheres, etc., takes the same time as cluster expanding a bulk structure equivalent in size—the equivalence quantified, e.g., by using the bulk structure which just exactly encloses the non-bulk structure. This is accomplished by introducing a new type of site in the system called a *ghost* site. The cluster function $\Gamma(\sigma)$ of any cluster, which includes a ghost site, is zeroed.

In addition, for low-symmetry systems, the CE-GLM chooses to use the bulk clusters and group them by the bulk symmetries. We discuss the implication of this at the end of this section, but before that, we present how one employs the CE-GLM. A large box is created containing lattice sites on the same underlying lattice as the non-bulk structure in question. Then, a subset of lattice sites inside this box is selected to represent the non-bulk structure. The remaining sites are the *ghost sites*. We note that this is not a new idea in the cluster expansion framework; the novelty is in how the ghost sites are kept in the system and thus enables the use of bulk clusters grouped by their bulk symmetries. In Fig. 6.1 we consider an example of a CE-GLM implementation for a hypothetical system, namely a single-layer Si-Ge slab on a simple cubic (sc) lattice. The large box in Fig. 6.1(a) is on an sc lattice as well, containing white ghost sites. As a side note, we can think of this entire box as the equivalent bulk structure. The non-bulk structure is not visible inside this box until the cross-section, identified by the black surrounding square in both Fig. 6.1(a) and Fig. 6.1(b), is taken. This reveals the structure in Fig. 6.1(b). A black ellipse delineates a 3-pt cluster originating from an atom inside the structure, but which includes ghost sites. The atoms within the cluster

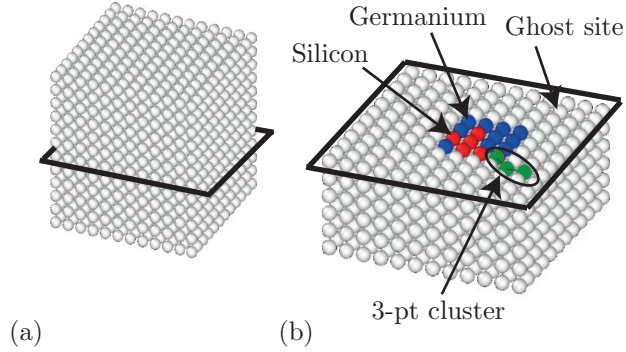


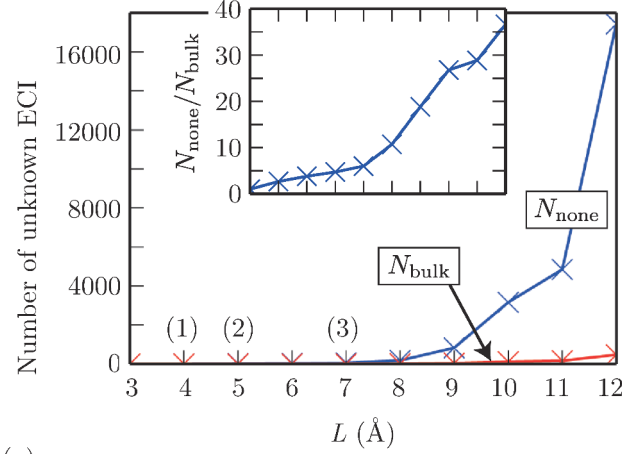
Figure 6.1: (Color) Demonstration of the CE-GLM for some hypothetical single-layer slab of Si (Silicon; red) and Ge (Germanium; blue) on a simple cubic lattice. Ghost sites are white. The black square surrounding the boxes in (a) and (b) show where a cross-section is taken in (a) to reveal the slab in (b). A black ellipse in (b) delineates a 3-pt cluster containing ghost sites. The sites in the cluster are colored green, purely as a visual construct, to associate them with the cluster.

are colored green, which is purely a visual construct. Since the cluster includes at least one ghost site, its cluster function is zero, and hence does not contribute to the cluster expanded property.


The CE-GLM employs the bulk clusters and associated bulk symmetries in the cluster expansion even for non-bulk geometries. To further understand the importance of symmetries in cluster expansions, and why an *ad hoc* grouping is expected to be necessary in most low-symmetry cluster expansions, consider the following hypothetical scenario. We want to cluster expand some function of configuration on an fcc lattice with a 5 Å lattice constant, and clusters with up till 4 points are to be included. This is a very common cluster expansion problem.^{13,15,79,80} The unknown ECI are learned by training the cluster expansion on some data set containing the property values of, say, 80-200 alloy structures.

Assume, for the sake of the argument, that the more clusters that are included, the better the property is captured. For quantum mechanical energies, e.g., this assumption seems to hold well.¹³ Therefore, it is of interest to include large sized clusters. Assume that the hypothetical property follows the size-hierarchy rules in Ref. [10]. In particular, we will include all 2-pt clusters of a given maximum spatial extent L . Then, we also include all 3-pt clusters within some lesser extent, chosen here as $L - 1$, and finally all 4-pt clusters with spatial extents less than $L - 2$ are included as well. In this sense, all clusters are bound by the length L quantifying the maximum spatial extent of any included cluster. As L increases, the number of clusters, generally speaking, increases as well. This, in turn, increases the number of unknown ECI.

In Fig. 6.2(a), we plot the number of unknown ECI versus L , first if we do not exploit the fcc lattice space group symmetries (N_{none}), and then together with the number of unknowns if we do use all the 48 space group symmetries of the bulk (N_{bulk}). The inset shows the ratio of the two curves from the main plot. Notice how this ratio increases rapidly with L . The fcc lattice is illustrated in Fig. 6.2(b) where red lines illustrate the largest 2-pt cluster for three different values of L . The parenthesized numbers can be matched with those above the abscissa in Fig. 6.2(a) to find the corresponding value of L . Notice that L is not necessarily the length of the largest 2-pt cluster. Furthermore, no other clusters than the largest 2-pt cluster are illustrated in Fig. 6.2(b) for any given L . From Fig. 6.2(a) we learn that, e.g., when the 2-pt cluster has maximum spatial extent $L = 10 \text{ \AA}$ we have $N_{\text{bulk}} = 118$ and $N_{\text{none}} = 3160$. Since these cluster sizes are very common it implies that, exploiting symmetries in the cluster expansion critically lowers the number of unknowns. In Ref. [81], a similar plot for N_{bulk} is given for the bcc lattice.



(a)



(b)

Figure 6.2: (Color) (a) Number of unknown ECI to be determined for an fcc lattice, representing some hypothetical alloy, versus the maximum spatial extent L of any included cluster in the cluster expansion, when not using any symmetries in the system (N_{none}) and when using all 48 bulk space group symmetries (N_{bulk}). The inset shows the ratio of the main plot curves versus the same L . The parenthesized numbers above the abscissa at L equal to 4, 5, and 7 Å are matched with those in (b). (b) The fcc lattice considered with a 5 Å lattice constant and its lattice sites represented as green balls. The cubic cell structure is delineated with black lines. The red lines with parenthesized numbers show the maximum allowed length of any 2-pt cluster included in the cluster expansion when choosing L at three different values along the abscissa in (a).

The choice of grouping the clusters by the bulk symmetries of the underlying lattice in the CE-GLM makes the clusters unique in the sense of being user-independent. Also, if one develops a cluster grouping, particular to a problem at

hand, this will, almost never, be useful for other structure shapes thus requiring new groupings to be developed for each new property and/or geometry. This issue is dealt away with in the CE-GLM since the choice is to use the bulk clusters. It is in this sense that the CE-GLM is transferable. Also, the choice presented in this work has the added feature that, if no ghost sites are present, the system is truly bulk, in which case the CE-GLM naturally merges into the standard cluster expansion without being a special case. Finally, the choice of cluster-grouping-scheme for low-symmetry systems does not increase computational efforts when compared to the bulk because of the ghost sites.

Finally, it is important to point out the fact that we do not pretend that the choice of bulk clusters will always work equally well on any geometry and configurational property. If it does not perform well, one may have to consider using a more elaborate grouping scheme. But the CE-GLM could provide a starting point for cluster expanding exotic geometries and the method might turn out to perform sufficiently well. In the context of the work presented in Refs. [82, 83], the CE-GLM is a computationally efficient way to obtain the bulk contribution to the energy, which is the first term in Eqs. (13) and (14) in Ref. [82], but without the added surface term perturbation. With the help of the ghost site concept, the CE-GLM can, in this light, be seen as a computationally very efficient algorithm—essentially without overhead in going from the bulk to the non-bulk—to compute this bulk contribution for arbitrary geometry shapes and sizes.

6.3 Application: Predicting Lattice Thermal Conductivity of Si-Ge Nanowires

6.3.1 Nanowires as Thermoelectric Devices

We consider Si-Ge nanowires used as thermoelectric devices.^{84,85} The efficiency of a thermoelectric device is quantified by the dimensionless figure of merit:⁸⁶

$$ZT = \frac{S^2 \sigma_{\text{el}}}{\kappa} T_{\text{avg}}, \quad (6.1)$$

where S is the Seebeck coefficient quantifying the induced thermoelectric voltage in response to a temperature gradient across the device, σ_{el} is the electrical conductivity of the device, T_{avg} its average temperature, and κ its thermal conductivity. The thermal conductivity κ can be decomposed into contributions from the electrons (el) and the phonons/lattice (lat) writing $\kappa = \kappa_{\text{el}} + \kappa_{\text{lat}}$, but typically, and this will be assumed here as well, $\kappa_{\text{el}} \ll \kappa_{\text{lat}}$.^{86,87} Therefore, from now on, we let $\kappa \approx \kappa_{\text{lat}}$. Thus we see from Eq. (6.1) that, keeping everything else constant, a minimal κ leads to a maximal ZT . The best thermoelectric devices currently have $ZT \approx 2$ at room temperature,^{86,88,89} and with $ZT > 3$, we can begin replacing compressor-based refrigeration units.^{90–92}

Ref. [93] discovered that Si nanowires possessed two orders of magnitude lower thermal conductivity compared to the bulk value of 150 W/m.K.⁹⁴ This brought a lot of interest into studying nanowires for thermoelectric applications. We are interested in the heat transferred along the nanowire axis, so we now let κ mean the axial lattice thermal conductivity. The thermal conductivity has been shown to depend on the configuration of the Si and Ge atoms,⁹⁵ and is thus well suited to be described by the cluster expansion.

In the present work, 50 nm long Si nanowires are considered, with roughened surfaces—roughened to help scattering phonons—which currently achieve a ZT around 1. A factor of, at least, 10 reduction in κ is, however, still needed. The authors in Ref. [78] used a clever meta cluster expansion (MCE) to capture κ . In essence, this is a cluster expansion technique which uses an *ad hoc* grouping of the clusters based on physical intuition. Then, in conjunction with the learned MCE, a genetic algorithm was employed to find the configuration with lowest κ . They found that, a silicon nanowire with pure planes of germanium (PPG)—the planes spaced about 1 nm apart—was the lowest thermal conductivity configuration. While the method was very successful, extending it to other properties and/or system geometries requires a new set of cluster groupings. This is in general a very complex task

6.3.2 Creating the Nanowire Data Set

This section details how we obtained the nanowire data set \mathcal{D} . First, computing the correlation functions requires a set of cluster orbits. Given these cluster orbits, the data set generation is a two-step process. First, the nanowire must be created, and its thermal conductivity computed. Then, we must compute the correlation functions in Eq. (2.8) using the orbits. Carrying this out for all wires gives $\boldsymbol{\kappa}$ and \mathbf{X} , respectively, in Eq. (2.12) (the general property vector \mathbf{q} is now a vector of thermal conductivities that we denote $\boldsymbol{\kappa}$). As discussed in the following section, we used molecular dynamics (MD), as implemented in LAMMPS, to compute κ .

Interestingly, it was shown in Ref. [78] that, instead of modeling long (~ 50 nm) nanowires with roughened surfaces, as we intend to do here, we can instead model short (~ 2 nm) ones with perfect surfaces. This is computationally easier, as we

need not worry about modeling the surface roughness, and we have fewer atoms overall. Therefore, we adopt this approach here as well.

The CE-GLM requires us to create two representations of the same nanowire structure. One representation for use in ATAT-GLM to compute the correlation functions. We call this the *ATAT representation*. The other representation was generated for directly computing κ in LAMMPS without periodic images and ghost sites. We call this representation the *LAMMPS representation*. Creating both representations started out at a common origin as follows: a large cubic box containing a diamond lattice, with a lattice constant $a_0 = 5.431 \text{ \AA}$, was created. A cylindrical region was selected along the $[111]$ direction representing the wire. OVITO⁹⁶ was used along the way to verify the nanowire geometry. All sites inside the box, but not inside the nanowire or any of the periodic images along its axis, were selected as ghost sites. The non-ghost atoms, including images of the wire along its axis, to be shown shortly, were chosen as Si or Ge depending on the particular configuration of the wire. This is the ATAT representation of the nanowire. Notice that the number of periodic images needed here depends on the cluster with largest spatial extent. One periodic image was enough along each axial direction of the nanowire in our case, since the largest cluster orbit included was 1.2 nm and the wire had a length of roughly 2 nm. Further details about exactly which cluster orbits were included are presented later.

Fig. 6.3 shows the ATAT representation of the nanowire. At first, it can be difficult to see the atoms in the nanowire. This is because all the ghost sites are included—shown as smaller black atoms—and they engulf the nanowire, just as in Fig. 6.1(a). Notice the repetition of the wire by one image along its axis in both directions. The image atoms are colored slightly lighter than the atoms in

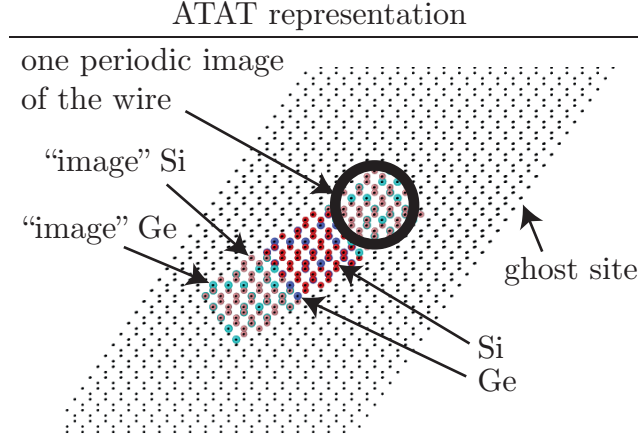


Figure 6.3: (Color) The ATAT representation of an arbitrary Si (red atoms) nanowire from the data set \mathcal{D} alloyed with 6.4 % Ge (blue atoms). Ghost sites are black atoms, made small so the wire becomes visible. The atoms in the two images of the wire along its axis are colored with a slightly lighter color to distinguish them from the main wire. The thick black circle identifies one of the two nanowire images. Ge atoms are made slightly larger than Si atoms for visual clarity.

the main wire. Of course, if we were not modeling the wire as being infinitely periodic along its axis, the images should not be there. Also, to be on the safe side, the distance from the wire to any point on the surface should be larger than the longest length of any cluster included for the same reasons, which can easily be achieved by making the cell large enough; there is little increased overhead in computing the correlation functions with an increase in the number of ghost sites. The entire geometry is bulk, so it is straightforward to parse with ATAT-GLM.

Next, to create the LAMMPS representation of the nanowire, all periodic images of the ATAT representation, and all the ghost sites, were removed. The nanowire axis was rotated to lie along the $[100]$ direction in the box. The box dimension along the wire axis was changed to the periodic length of the wire closest to 2 nm, which is $2\sqrt{3}a_0 \approx 1.88$ nm along $[111]$ in a diamond lattice before thermalization,

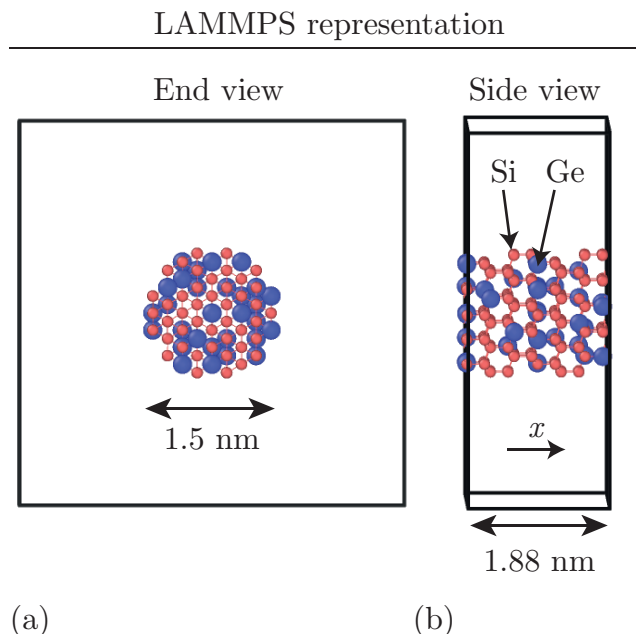


Figure 6.4: (Color) The corresponding LAMMPS representation of the nanowire in Fig. 6.3. The black surrounding box is the LAMMPS simulation cell. (a) Shows the nanowire looking down its axis, and (b) illustrates the same wire from the side. The horizontal arrow inside the cell in (b) identifies the x -axis. Double headed arrows report sizes of the wire, and the smaller arrows in (b), over the wire, identify the atom types. Ge atoms are made slightly larger than Si atoms for visual clarity.

to be discussed in the next section. The nanowire diameter was made 1.5 nm, also before thermalization, leaving it with 220 atoms. Please consult Fig. 6.4 for a visualization of the LAMMPS representation of the nanowire. Fig. 6.4(a) and (b) show the end and side view, respectively, with the black box in each view being the LAMMPS simulation cell. In Fig. 6.4(b), notice how the wire becomes periodic along its axis when applying periodic boundary conditions along x in LAMMPS.

6.3.3 Computing the Nanowire Thermal Conductivities and the Design Matrix

To compute the thermal conductivities κ of the LAMMPS representations of the nanowires, the Green-Kubo method⁹⁷ was employed. Ref. [98] provides a thorough study, and detailed discussion, of this. We will refer to the LAMMPS representations of the nanowires as simply (nano)wires in this section. The wire axis is aligned along x with periodic boundary conditions, and fixed boundary conditions along dimensions transverse to the axis, so

$$\kappa = \frac{V_{\text{wire}}}{k_{\text{B}}T^2} \lim_{\tau_m \rightarrow \infty} \int_0^{\tau_m} \langle J_x(\tau) J_x(0) \rangle d\tau, \quad (6.2)$$

where V_{wire} is the nanowire volume, T is its temperature, k_{B} is the Boltzmann constant, $J_x(\tau)$ is the x component of the heat flux vector \mathbf{J} at time τ and $\langle J_x(\tau) J_x(0) \rangle$ is the heat flux autocorrelation function (HCACF) at lag τ . The heat flux vector \mathbf{J} must be collected while the system is in a constant particle number N , system volume V , and energy E (NVE) ensemble following a thermalization according to the constant N , V , and temperature T (NVT) ensemble.

The MD time step was set to 1 fs and the heat flux vector was defined as:

$$\mathbf{J} = \frac{1}{V_{\text{wire}}} \left[\sum_{i=1}^N e_i \mathbf{v}_i + \frac{1}{2} \sum_{i=1}^N \sum_{\substack{j=1 \\ i \neq j}}^N (\mathbf{F}_{ij} \cdot \mathbf{v}_i) \mathbf{r}_{ij} + \frac{1}{6} \sum_{i,j,k=1}^N (\mathbf{r}_{ij} + \mathbf{r}_{ik}) (\mathbf{F}_{ijk} \cdot \mathbf{v}_i) \right],$$

where the sums are over all N atoms in the nanowire, \mathbf{r}_i is the position of atom i , $\mathbf{r}_{ij} \equiv \mathbf{r}_i - \mathbf{r}_j$, e_i and \mathbf{v}_i are the total per-atom energy (kinetic plus potential) and velocity of atom i , respectively, \mathbf{F}_{ij} is the pair potential force on atom i due to its neighbor j , and \mathbf{F}_{ijk} , in the third term, is the 3-body force term. The forces are obtained from the potential energy landscape which is modeled with the Tersoff scheme.⁹⁹ Please refer to Ref. [98] for more details, and for a discussion of

different, but equally reasonable ways, of defining this heat flux vector. Note that $\langle \cdot \rangle$ in Eq. (6.2) is an ensemble average, which, in an MD simulation, becomes a time average, assuming, as we will, ergodicity is satisfied. The time τ_m should be much smaller than the total simulation time, but larger than the time required for the HCACF to decay to zero.⁹⁸

To collect \mathbf{J} , the system needs to be thermalized in an NVT ensemble at $T = 300$ K. In order to do this, it is important to thermalize the nanowire surface, especially for wires of this small size having large surface to volume ratios. We found the following procedure successful. The atomic coordinates in the wire were first adjusted using a conjugate gradient method until finding a (possibly local) minimum in the potential energy. We then began an annealing process to thermalize the surface: an initial set of velocities were chosen for each atom from a mean zero Gaussian distribution with standard deviation scaled to 1000 K, and the system was run for 500 ps. Then, the temperature was gradually lowered 100 K at a time over 10 ps. At each temperature, we ran the system for 100 ps before decreasing by another 100 K, continuing in this way until reaching 300 K. After obtaining a room temperature NVT ensemble, the nanowire axis was pressurized to 1 bar using a constant N , pressure P , and $T = 300$ K ensemble. This was necessary to reduce strains from the size mismatch between the Si and Ge atoms. Indeed, we observed the pressure to be, typically, around 500 bar before pressurizing to 1 bar. After the NPT ensemble run, the system was switched back to the room temperature NVT ensemble and run for 1 ns.

Following this, the system was finally switched to an NVE ensemble, and run for 16 ns. The total energy stayed constant to within numerical accuracy expected with the Verlet integrator, and the temperature stayed constant, on average, by

oscillating around 300 K with a 40 K amplitude. The axial heat flux J_x was output every fourth time step (for memory reasons), and the HCACF computed from this. The HCACF was then integrated. Determining when to stop the integration, i.e., choosing the best τ_m in Eq. (6.2), called τ_m^* , was done as follows. Forty moving averages of various window sizes ranging from 50 ps to 200 ps were computed. The time point where the standard deviation of all window sizes obtained a minimum was chosen as τ_m^* .

To verify our overall implementation, and to show explicitly how the choice of τ_m was carried out, please refer to Fig. 6.5. We ran a pure Si wire, Fig. 6.5(a), and the PPG wire, Fig. 6.5(b). Our approach gives 4.1 ± 0.4 W/m.K for the pure Si wire and 0.12 ± 0.03 W/m.K for the PPG wire. The errors are determined from re-running the wires with different initial random velocity distributions. In comparison, Ref. [78] obtains 4.1 ± 0.3 W/m.K for the pure Si wire and 0.23 ± 0.05 W/m.K for the PPG wire with the same method of determining the standard deviations. Discrepancies between these numbers are due to different MD software used, a different surface reconstruction technique, and different thermalization and pressurization timings and methods.

Finally, this is how the design matrix \mathbf{X} was obtained. The i th row of the design matrix contains the cluster-orbit-averaged correlation functions of nanowire i . It is *a priori* unclear which cluster orbits are needed to describe thermal conductivity. In this work we used up till 5-pt cluster orbits with maximum spatial extents of the cluster orbits chosen as 12, 8, 6, and 5 Å for the 2, 3, 4, and 5-pt cluster orbits, respectively. Since the largest cluster orbit is 1.2 nm, one periodic nanowire image, which is 1.88 nm, is enough along the axis. For each structure in the data set, the ATAT representation of the nanowire was parsed by ATAT-GLM providing the




Figure 6.5: (Color) Demonstration of the method used to obtain the best τ_m in Eq. (6.2) called τ_m^* . All ordinates are in units of W/m.K. (a) Shows the pure Si wire and (b) the PPG wire. Top plots are κ versus τ_m in Eq. (6.2). In the bottom plots the μ/σ graph is the mean/standard deviation of 40 different unweighted running averages of window sizes ranging from 50 to 200 ps applied to the top graphs. The μ/σ graph is measured on the right/left ordinate. The red circles with vertical lines crossing the abscissa mark the minima of the σ graphs and hence the times τ_m^* .

nanowire correlation functions as output.

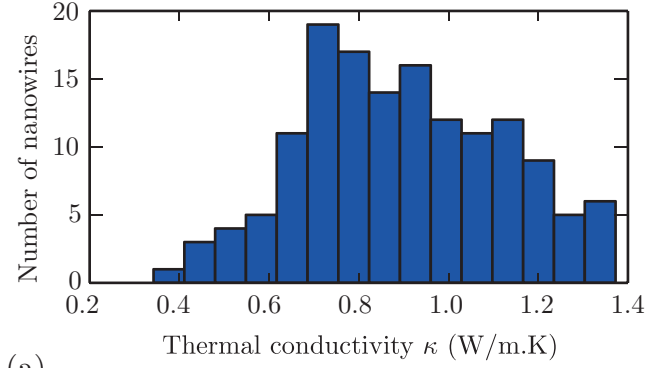
6.4 Results

In Fig. 6.6(a) we show a histogram of \mathcal{D} containing 145 wires each with a random Si-Ge configuration. The concentration of Ge was restricted to the range 3 % to 22 % as in Ref. [78], which could be due to, e.g., weight constraints, and the wires


had an average thermal conductivity of 0.90 W/m.K compared to 1.1 W/m.K in Ref. [78]. We will refer to a wire from \mathcal{D} as a random wire (RW). Please see the top left of Fig. 6.6(b) for an illustration of this type of wire. For the purpose of learning β we split the RW data set into a training set (RW train), of 140 wires, and a test set (RW test) containing 5 wires. The ECI were learned on RW train, using least squares with repeated random sub-sampling validation using 10 splits,¹⁰⁰ and then initial predictions were made on RW test. The predictions on RW test serves as a preliminary test of the CE-GLM in the following sense. If the prediction errors are much larger than the training errors at this stage, the CE-GLM is not capturing the underlying physics and we should not trust further predictions. In Fig. 6.7(a) we first show the least squares fit on RW train. The average *root-mean-square* (RMS) training error of all random splits was 0.15 W/m.K.

In Fig. 6.7(b) we show the initial test of the RW train fit on RW test, and as can be seen, the predictions made on RW test have errors comparable to those on RW train. Indeed, the test error is 0.12 W/m.K. Note that the errors in general are relatively large. This was the same observation made in Ref. [78], and is due to inherent sizable noise in equilibrium MD simulations.

The purpose of this work is to show whether the CE-GLM method would have lead to the same conclusions as in Ref. [78]. To this end, a data set of 18 nanowires was created, each wire being a perturbed configuration from the PPG wire. We refer to these as the similar-to-PPG (SPPG) wires. As an example, one perturbation is obtained by substituting a random Ge atom from one of the pure Ge planes in the PPG wire with a random Si atom in the non-planes region of the wire. Each SPPG wire has its own number of atoms, ranging from 1 to 8, randomly substituted in this way. See Fig. 6.6(b) for an example SPPG wire with



(a)



(b)

Figure 6.6: (Color) (a) Nanowire data set of random Si-Ge configurations plotted in a histogram with the number of elements versus their thermal conductivity κ . (b) Each type of wire considered in the fitting and predictions of thermal conductivity: wires with random Si-Ge configurations (RW), the pure planes of Ge wire (PPG), and wires that are similar to the PPG wire (SPPG). Red/blue atoms are Si/Ge. For the SPPG wire, black circles show where two random atoms, one atom from one of the pure Ge planes, and one atom from the non-plane region, were substituted to perturb the PPG wire.

1 atom substituted.

Notice that, importantly, the RW test wires, the SPPG wires, and the single

PPG wire, are not used in training the least squares model. Only RW train is employed for this task. If the CE-GLM is useful here, we will predict that a large majority (if not all) of the SPPG wires have lower thermal conductivity than the RWs. In Fig. 6.7(b) we see that this is also the case. Furthermore, the RMS prediction error is 0.25 W/m.K, which is similar in magnitude to the RW data sets. Also, a larger error for the, more or less, ordered SPPG wires is to be expected since the predictions are made based on a fit to the *random* data set. In principle, this should not matter since we assume that the thermal conductivity is captured well in all of configuration space. However, in practice, since we only observe a limited amount of data, it will typically be the case that we enjoy an increased predictive capability in regions closer to the structures employed in the training process.

As a final additional test, we see whether CE-GLM predicts the PPG wire to have overall lowest thermal conductivity of all wires considered above. This is an important test. If it (almost) has the lowest conductivity, we would have been able to provide the same conclusion as Ref. [78] if coupling the cluster expanded thermal conductivity with an optimization routine such as the genetic algorithm or an adaptive sequential Monte Carlo method.⁷³ Remarkably, Fig. 6.7(b) shows that we do indeed make this important prediction. Note that, whether this is indeed the global minimum configuration is out the scope of this work to determine. These results provide evidence that the CE-GLM does capture the thermal conductivity of Si-Ge nanowires to, at least, the same degree as the MCE, and is therefore a useful method to apply in this problem.




Figure 6.7: (Color) Predictions of thermal conductivity from the cluster expansion ghost lattice method (CE-GLM) versus the results from molecular dynamics. In (a) the least squares fit on RW train is shown. Only one of the 10 train/test splits are shown (green/white circles). In (b) the CE-GLM is tested on three different test sets: RW test, the SPPG wires, and the PPG wire. Note that in (b), RW train (green circles) is the same as in (a). Error bars on the SPPG wires and the PPG wire stem from 10 different random splits of the RW data into train and test. For more information on the various nanowire data sets see, e.g., Fig. 6.6(b).

6.5 Conclusion

In this chapter we introduced a modified cluster expansion method, called the cluster expansion ghost lattice method (CE-GLM), which is distinct to the regular

ternary cluster expansion with vacuum atom types. The CE-GLM uses the bulk clusters and bulk symmetries for any geometry. The ghost site—zeroing cluster functions of clusters containing it—turns any structure into an effective bulk geometry and makes the CE-GLM computationally efficient when implemented in existing cluster expansion software packages—no matter the shape and/or size of the structure in question when compared to an equivalent bulk system. We discussed how to implement the method in the cluster expansion software package ATAT.

To test the usefulness of the CE-GLM the method was employed to predict the thermal conductivity of Si-Ge nanowires. When comparing to a similar work carried out recently we find great agreement: the lowest thermal-conductivity-nanowire configuration is that of an Si wire with planes of pure Ge spaced 1 nm apart. This establishes that the CE-GLM can be useful for cluster expanding non-bulk low-symmetry systems with minor modification to existing cluster expansion software.

The main drawback of the CE-GLM is that it does not provide a way to systematically check whether increasingly precise surface terms matter greatly in the problem at hand or not.

In future work it will be interesting to employ the method on other even more exotic geometries and preferably compare to other cluster grouping methods. Furthermore, since the CE-GLM computes the unperturbed, i.e., the bulk, energy of any structure, a study, for non-bulk structures, of how the surface contribution plays a role would be useful.

CHAPTER 7

CONCLUSION

Surrogates used in place of expensive computer code are essential to optimization and characterization of materials *in silico*. The surrogate attempts to approximate the materials response surface of any given property as well as possible from a data set containing samples of the true response surface. The frameworks developed in the thesis are general, but for our applications, we focused mainly on the important subset of materials called alloys. In the case of alloys, a particular surrogate exists called the cluster expansion.

The main topics presented in this thesis has been the improvement of surrogates for materials characterization, optimization, and design and the quantification of the uncertainty in predicted materials properties and quantities of interest (QoIs) from using surrogates.

Thermodynamic characterization of materials is extremely important since it tells us if the material will be stable or not if created in the real world. In Chapter 4, we showed that the concept of relative entropy employed to measure the distance between a candidate and the true Boltzmann distribution provided significant improvements to thermodynamic predictions of the material in much closer agreement with experiments. This should provide evidence that materials science can benefit greatly from information theory. This merge is indeed the current direction of the computational materials science field.

Predicting materials QoIs are important to both materials characterization, optimization, and to design. However, without knowing *how well* we know the QoI, our predictions are meaningless. This motivated the work of Chapter 5, where

we developed a rigorous framework for propagate and quantify the uncertainty in predicted QoIs. We demonstrated that the predictive variance in QoIs should be expected to be around 5-10 %, which can be improved for the application at hand, when the materials property is quantum mechanical energy, in agreement with what has been assumed in the field so far. This uncertainty propagation is essential in order to approach future materials design in an informed and reliable way and our work is the first to attempt such a rigorous propagation of uncertainty from the surrogate to the QoI in the computational materials science field. There is, however, still much more work needed in this area.

In Chapter 6, we observed that the traditional cluster expansion method, developed for bulk systems, did not perform well for low-symmetry systems. We considered a materials design problem involving the minimization of thermal conductivity in Si-Ge nanowires but were, at first, unable to proceed using traditional methods. Therefore, we set out to develop a fast cluster expansion method for such systems and found a successful approach called the cluster expansion ghost lattice method capable of capturing the thermal conductivity and predict the lowest conductivity structure. Our method is approximate, but had huge benefits in terms of computational implementation and speed when applied to any system of arbitrary shape and proved useful in the particular application of predicting nanowire thermal conductivity.

As final words, we are headed towards the realization of virtual materials laboratories, but there are still many developments and discoveries to be made before we reach this goal. We sincerely hope that our contributions and ideas help reach this goal sooner rather than later. Rest assured that some day in the future these laboratories will see the light of day.

BIBLIOGRAPHY

- [1] G. Kresse and J. Furthmüller, Phys. Rev. B **54**, 11169 (1996).
- [2] G. Kresse and J. Furthmüller, Comput. Mat. Sci. **6**, 15 (1996).
- [3] J. Sanchez, Physical review B **48**, 14013 (1993).
- [4] N. Ashcroft and N. Mermin, *Solid State Physics* (Saunders College, Philadelphia, 1976).
- [5] J. M. Sanchez, F. Ducastelle, and D. Gratias, Physica A: Statistical Mechanics and its Applications **128**, 334 (1984).
- [6] T. Mueller and G. Ceder, Physical Review B **82**, 184107 (2010).
- [7] S. G. Brush, Reviews of modern physics **39**, 883 (1967).
- [8] G. Ceder, Computational Materials Science **1**, 144 (1993), URL <http://www.sciencedirect.com/science/article/pii/0927025693900058>.
- [9] A. van de Walle and G. Ceder, Journal of Phase Equilibria **23**, 348 (2002).
- [10] N. A. Zarkevich and D. D. Johnson, Phys. Rev. Lett. **92**, 255702 (2004), URL <http://link.aps.org/doi/10.1103/PhysRevLett.92.255702>.
- [11] A. van de Walle, Calphad **33**, 266 (2009), URL <http://www.sciencedirect.com/science/article/pii/S0364591608001314>.
- [12] C. L. Lawson and R. J. Hanson, *Solving least squares problems*, vol. 161 (SIAM, 1974).
- [13] A. Van de Walle, M. Asta, and G. Ceder, Calphad **26**, 539 (2002).

- [14] V. Blum, G. L. Hart, M. J. Walorski, and A. Zunger, Physical Review B **72**, 165113 (2005).
- [15] L. J. Nelson, G. L. Hart, F. Zhou, and V. Ozoliņš, Physical Review B **87**, 035125 (2013).
- [16] J. R. Norris, *Markov chains*, 2 (Cambridge university press, 1998).
- [17] J. Sethna, *Statistical mechanics: entropy, order parameters, and complexity*, vol. 14 (Oxford University Press, 2006).
- [18] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, The journal of chemical physics **21**, 1087 (1953).
- [19] W. K. Hastings, Biometrika **57**, 97 (1970).
- [20] P. J. Green, Biometrika **82**, 711 (1995).
- [21] C. Geyer, Handbook of Markov Chain Monte Carlo pp. 3–48 (2011).
- [22] P. J. Green and D. I. Hastie, Genetics **155**, 1391 (2009).
- [23] P. J. Green, Oxford Statistical Science Series pp. 179–198 (2003).
- [24] R. Tibshirani, Journal of the Royal Statistical Society. Series B (Methodological) pp. 267–288 (1996).
- [25] A. Kohan, P. Tepeš, G. Ceder, and C. Wolverton, Computational Materials Science **9**, 389 (1998), URL <http://www.sciencedirect.com/science/article/pii/S0927025697001687>.
- [26] P. R. Alonso, P. H. Gargano, and G. H. Rubiolo, Calphad **38**, 117 (2012), URL <http://www.sciencedirect.com/science/article/pii/S0364591612000429>.

- [27] R. Arróyave and M. Radovic, Phys. Rev. B **84**, 134112 (2011), URL <http://link.aps.org/doi/10.1103/PhysRevB.84.134112>.
- [28] C. Ravi, B. K. Panigrahi, M. C. Valsakumar, and A. van de Walle, Phys. Rev. B **85**, 054202 (2012), URL <http://link.aps.org/doi/10.1103/PhysRevB.85.054202>.
- [29] Z. W. Lu, D. B. Laks, S.-H. Wei, and A. Zunger, Phys. Rev. B **50**, 6642 (1994), URL <http://link.aps.org/doi/10.1103/PhysRevB.50.6642>.
- [30] A. S. Dalton, A. A. Belak, and A. Van der Ven, Chemistry of Materials **24**, 1568 (2012), URL <http://pubs.acs.org/doi/abs/10.1021/cm203283v>.
- [31] S. Kullback and R. A. Leibler, The Annals of Mathematical Statistics **22**, 79 (1951), URL <http://projecteuclid.org/euclid.aoms/1177729694>.
- [32] P. C. Kelires and J. Tersoff, Phys. Rev. Lett. **63**, 1164 (1989), URL <http://link.aps.org/doi/10.1103/PhysRevLett.63.1164>.
- [33] S. de Gironcoli, P. Giannozzi, and S. Baroni, Phys. Rev. Lett. **66**, 2116 (1991), URL <http://link.aps.org/doi/10.1103/PhysRevLett.66.2116>.
- [34] M. Laradji, D. P. Landau, and B. Dünweg, Phys. Rev. B **51**, 4894 (1995), URL <http://link.aps.org/doi/10.1103/PhysRevB.51.4894>.
- [35] A. Qteish and R. Resta, Phys. Rev. B **37**, 6983 (1988), URL <http://link.aps.org/doi/10.1103/PhysRevB.37.6983>.
- [36] R. H. Taylor, S. Curtarolo, and G. L. Hart, Physical Review B **81**, 024112 (2010).
- [37] C. Barrett and O. R. Trautz, Transactions AIME **175**, 579 (1948), URL <http://www.onemine.org/search/index.cfm?fullText=trautz&start=10>.

- [38] C. M. Bishop et al., *Pattern recognition and machine learning*, vol. 4 (springer New York, 2006).
- [39] M. S. Shell, The Journal of Chemical Physics **129**, 144108 (pages 7) (2008), URL <http://link.aip.org/link/?JCP/129/144108/1>.
- [40] A. Chaimovich and M. S. Shell, Phys. Rev. E **81**, 060104 (2010), URL <http://link.aps.org/doi/10.1103/PhysRevE.81.060104>.
- [41] A. Chaimovich and M. S. Shell, The Journal of Chemical Physics **134**, 094112 (pages 15) (2011), URL <http://link.aip.org/link/?JCP/134/094112/1>.
- [42] M. S. Shell, The Journal of Chemical Physics **137**, 084503 (pages 13) (2012), URL <http://link.aip.org/link/?JCP/137/084503/1>.
- [43] I. Bilionis and N. Zabaras, The Journal of Chemical Physics **138**, 044313 (pages 12) (2013), URL <http://link.aip.org/link/?JCP/138/044313/1>.
- [44] J. P. Perdew, K. Burke, and M. Ernzerhof, Phys. Rev. Lett. **77**, 3865 (1996), URL <http://link.aps.org/doi/10.1103/PhysRevLett.77.3865>.
- [45] J. P. Perdew, K. Burke, and M. Ernzerhof, Phys. Rev. Lett. **78**, 1396 (1997).
- [46] G. Kresse and D. Joubert, Phys. Rev. B **59**, 1758 (1999), URL <http://link.aps.org/doi/10.1103/PhysRevB.59.1758>.
- [47] H. J. Monkhorst and J. D. Pack, Physical Review B **13**, 5188 (1976).
- [48] A. van de Walle and M. Asta, Modelling and Simulation in Materials Science and Engineering **10**, 521 (2002), URL <http://stacks.iop.org/0965-0393/10/i=5/a=304>.
- [49] G. L. W. Hart and R. W. Forcade, Phys. Rev. B **77**, 224115 (2008).

- [50] P. Del Moral, A. Doucet, and A. Jasra, Journal of the Royal Statistical Society: Series B (Statistical Methodology) **68**, 411 (2006), ISSN 1467-9868, URL <http://dx.doi.org/10.1111/j.1467-9868.2006.00553.x>.
- [51] I. Bilionis and P. Koutsourelakis, Journal of Computational Physics **231**, 3849 (2012), ISSN 0021-9991, URL <http://www.sciencedirect.com/science/article/pii/S0021999112000630>.
- [52] M. E. Newman and G. T. Barkema, *Monte Carlo methods in statistical physics*, vol. 13 (Clarendon Press Oxford, 1999).
- [53] Z. R. Kenz, H. Banks, and R. C. Smith (2013).
- [54] R. C. Smith, *Uncertainty Quantification: Theory, Implementation, and Applications*, vol. 12 (SIAM, 2013).
- [55] E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge University Press, 2003).
- [56] P. S. d. Laplace, *Analytical Theory of Probability* (1812).
- [57] H. Jeffreys, *Theory of Probability* (1939).
- [58] T. Park and G. Casella, Journal of the American Statistical Association **103**, 681 (2008).
- [59] W. R. Gilks, *Markov chain monte carlo* (Wiley Online Library, 2005).
- [60] S. Plimpton, Journal of computational physics **117**, 1 (1995).
- [61] G. L. Bretthorst, in *Maximum entropy and Bayesian methods* (Springer, 1996), pp. 1–42.
- [62] T. Mueller and G. Ceder, Physical Review B **80**, 024103 (2009).

- [63] I. Bilonis and N. Zabaras, *Inverse Problems* **30**, 015004 (2014).
- [64] D. S. Sivia, *Data analysis: A Bayesian Tutorial* (Oxford University Press, 1996).
- [65] M. C. Kennedy and A. O'Hagan, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **63**, 425 (2001).
- [66] G. S. Kimeldorf and G. Wahba, *The Annals of Mathematical Statistics* pp. 495–502 (1970).
- [67] A. O'Hagan and J. Kingman, *Journal of the Royal Statistical Society. Series B ...* **40**, 1 (1978), URL <http://www.jstor.org/stable/2984861>.
- [68] P. S. Craig, M. Goldstein, J. C. Rougier, and A. H. Seheult, *Journal of the American Statistical Association* **96**, 717 (2001).
- [69] X. Chen, Z. Jane Wang, and M. J. McKeown, *Signal Processing* **91**, 1920 (2011).
- [70] G. O. Roberts, *Annals of Applied Probability* **7**, 110 (1997).
- [71] A. Zunger, L. G. Wang, G. L. W. Hart, and M. Sanati, *Modelling and Simulation in Materials Science and Engineering* **10**, 685 (2002).
- [72] J. Kristensen, I. Bilonis, and N. Zabaras, *Physical Review B* **87**, 174112 (2013).
- [73] P. Del Moral, A. Doucet, and A. Jasra, *Statistics and Computing* **22**, 1009 (2012).
- [74] W. Xie, W. Xiong, C. A. Marianetti, and D. Morgan, *Physical Review B* **88**, 235128 (2013).

- [75] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, *Science* **220**, 671 (1983).
- [76] J. Skilling, *Bayesian inference and maximum entropy methods in science and engineering* **735**, 395 (2004).
- [77] B. Han, A. Van der Ven, G. Ceder, and B.-J. Hwang, *Physical Review B* **72**, 205409 (2005).
- [78] M. Chan, J. Reed, D. Donadio, T. Mueller, Y. Meng, G. Galli, and G. Ceder, *Physical Review B* **81**, 174303 (2010).
- [79] G. D. Garbulsky and G. Ceder, *Phys. Rev. B* **51**, 67 (1995), URL <http://link.aps.org/doi/10.1103/PhysRevB.51.67>.
- [80] G. Ceder, G. D. Garbulsky, D. Avis, and K. Fukuda, *Phys. Rev. B* **49**, 1 (1994), URL <http://link.aps.org/doi/10.1103/PhysRevB.49.1>.
- [81] G. L. Hart, V. Blum, M. J. Walorski, and A. Zunger, *Nature materials* **4**, 391 (2005).
- [82] D. Lerch, O. Wieckhorst, G. L. W. Hart, R. W. Forcade, and S. Müller, *Modelling and Simulation in Materials Science and Engineering* **17**, 055003 (2009), URL <http://stacks.iop.org/0965-0393/17/i=5/a=055003>.
- [83] R. Drautz, H. Reichert, M. Fähnle, H. Dosch, and J. Sanchez, *Physical review letters* **1**, 1 (2001), URL <http://journals.aps.org/prl/abstract/10.1103/PhysRevLett.87.236102>.
- [84] M. Amato, M. Palummo, R. Rurali, and S. Ossicini, *Chemical Reviews* **114**, 1371 (2014), URL <http://pubs.acs.org/doi/abs/10.1021/cr400261y>.
- [85] M. Shelley and A. A. Mostofi, *EPL (Europhysics Letters)* **94**, 67001 (2011).

- [86] D. Donadio and G. Galli, Physical review letters **102**, 195901 (2009).
- [87] T. T. Vo, A. J. Williamson, V. Lordi, and G. Galli, Nano letters **8**, 1111 (2008).
- [88] B. Qiu, H. Bao, G. Zhang, Y. Wu, and X. Ruan, Computational Materials Science **53**, 278 (2012).
- [89] C. C. Yang and S. Li, ChemPhysChem **12**, 3614 (2011).
- [90] D. M. Rowe, *Materials, Preparation, and Characterization in Thermoelectrics*, vol. 1 (CRC press, 2012).
- [91] T. Harman, P. Taylor, M. Walsh, and B. LaForge, Science **297**, 2229 (2002).
- [92] A. Majumdar, Science **303**, 777 (2004).
- [93] S. G. Volz and G. Chen, Applied Physics Letters **75**, 2056 (1999).
- [94] C. Glassbrenner and G. A. Slack, Physical Review **134**, A1058 (1964).
- [95] B. Abeles, Phys. Rev. **131**, 1906 (1963), URL <http://link.aps.org/doi/10.1103/PhysRev.131.1906>.
- [96] A. Stukowski, Modelling and Simulation in Materials Science and Engineering **18**, 015012 (2010), URL <http://ovito.org/>.
- [97] D. Frenkel and B. Smit, *Understanding molecular simulation: from algorithms to applications*, vol. 1 (Academic press, 2001).
- [98] P. K. Schelling, S. R. Phillpot, and P. Keblinski, Physical Review B **65**, 144306 (2002).
- [99] J. Tersoff, Phys. Rev. B **39**, 5566 (1989).
- [100] S. M. Weiss and C. A. Kulikowski (1990).