Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Solid State Chemistry 176 (2003) 587–593
JOURNAL OF SOLID STATE CHEMISTRY
http://elsevier.com/locate/jssc

Application of machine-learning methods to solid-state chemistry: ferromagnetism in transition metal alloys

Gregory A. Landrum* and Hugh Genin[1]

Rational Discovery LLC, 555 Bryant St. #467, Palo Alto, CA 94301, USA

Received 21 January 2003; received in revised form 5 June 2003; accepted 10 June 2003

Abstract

Machine-learning methods are a collection of techniques for building predictive models from experimental data. The algorithms are problem-independent: the chemistry and physics of the problem being studied are contained in the descriptors used to represent the known data. The application of a variety of machine-learning methods to the prediction of ferromagnetism in ordered and disordered transition metal alloys is presented. Applying a decision tree algorithm to build a predictive model for ordered phases results in a model that is 100% accurate. The same algorithm achieves 99% accuracy when trained on a data set containing both ordered and disordered phases. Details of the descriptor sets for both applications are also presented.

Keywords: Ferromagnetism; Transition metal alloys; Disorder; Machine learning; Data mining; Cluster analysis; Decision trees; Quantitative structure–activity relationships

1. Introduction

Machine learning is rapidly becoming an important predictive tool in chemistry [1,2]. Most useful for problems with abundant experimental data, machine-learning techniques have been successfully applied to chemical problems running the gamut from catalyst modeling [3] to prediction of the toxicity of organic molecules [4]. In this contribution we present the results of applying some machine-learning methods to the prediction of ferromagnetism in the transition metals and their binary alloys.

The phrase "machine-learning" subsumes a diverse set of algorithms, such as neural networks, decision trees, and support vector machines. Probably the most widely used machine-learning algorithms (though they are rarely recognized as such) are the regression methods that form the basis of most quantitative structure–activity relationship (QSAR) studies [5]. All of these algorithms have one goal in common: the automated discovery of patterns in data in order to build predictive models [6,7]. Machine-learning algorithms are independent of the particular problem being investigated, be it predicting ferromagnetism or stock prices; they rely solely on the data. An advantage resulting from this characteristic is that the underlying problem domain—the physics or chemistry—need not be fully understood. The trade-off is that machine learning generally requires a large amount of data. In order for the methods to be successful, these data should be internally consistent; i.e., either collected under similar conditions or, at a bare minimum, normalizable.

Solid-state physics offers several problems, such as high-T_c superconductivity and giant magnetoresistance, which are potential candidates for machine-learning approaches. These are industrially relevant problems that are incompletely understood but where accurate experimental data are available. It is not clear whether, at this time, a *sufficient* quantity of consistently measured data is available for successful application of machine learning to these areas. In order to test the effectiveness of machine learning on solid-state problems with limited data, we have studied the problem of predicting a complex materials property from a small set of highly accurate data. We chose ferromagnetism as the materials property because it is both a complex and a well-understood phenomenon, providing an excellent test case [8–12]. In addition, while current

*Corresponding author.
E-mail address: landrum@rationaldiscovery.com (G.A. Landrum).
[1] Brenton Court, Mountain View, CA 94043, USA.

0022-4596/$ - see front matter © 2003 Elsevier Inc. All rights reserved.
doi:10.1016/S0022-4596(03)00343-8

electronic-structure calculations are capable of very accurate predictions of ferromagnetism, they are very computationally expensive and not universally applicable.

For the machine-learning option, we chose to use a simple classification algorithm, a decision tree, and to attempt only binary classification. In other words, the algorithm employed herein attempts to assign the data points—the elements or alloys—to one of two classes: those that display ferromagnetism, and those that do not. A second learning method, hierarchical clustering, is also investigated for its ability to uncover patterns in the data. Note that the scope of this study is limited to predicting merely the existence of a ferromagnetic state, not the magnitude of magnetic moments.

2. Methods

2.1. Descriptors

As is the case in QSAR studies, when using machine-learning methods, substances such as molecules, alloys, and salts are represented using a set of descriptors instead of (or in addition to) their chemical identity. Descriptors can be drawn from experimental properties or computed *ab initio*. They can be as complex as the chemical hardness of a transition metal alloy or as simple as the molecular weight of an organic compound. The process by which a set of descriptors for a problem is selected from the infinity of possibilities varies from modeler to modeler; there is no optimal strategy. The one constant is that, in order to build useful models, the descriptors used must capture the relevant physics and chemistry of the problem. This renders descriptor sets highly problem dependant: the best learning methodology in the world is unlikely to be able to produce a meaningful model for ferromagnetism using a set of descriptors which has been optimized for the prediction of superconductivity. Thus, although machine-learning techniques themselves are problem independent, applying them to solve real problems requires significant domain knowledge in order to be able to identify a suitable set of descriptors.

2.2. Hierarchical clustering

Hierarchical clustering refers to a class of learning algorithms that group data points into clusters based upon the distances between them in descriptor space.[2] The general idea is to start with a group of clusters made up of the individual data points and then to iteratively combine clusters to form a tree (hierarchy) of increasingly large clusters.[3] See Fig. 5 for an example. The resulting cluster tree, which can reveal groupings and regularities in the data set, is often a useful tool for analyzing and understanding the data.

The data for the ferromagnetic alloys discussed below were clustered using a Euclidean distance metric and the group average method [13]. The data were standardized before being clustered.[4]

A more detailed introduction to the theory and practice of hierarchical clustering can be found in Refs. [13,14].

2.3. Decision trees

Decision trees are a class of predictive models for classification that are constructed using *supervised learning* algorithms.[5] A sample tree is shown in Fig. 1. Decision trees classify data points by starting at the top of the tree—the root node—and moving down through the tree by asking a series of "if–then–else" questions of the descriptor values at each branch point until a terminal (leaf) node is hit. For example, the tree in Fig. 1 would start by examining the value of Descriptor 1. If that value is less than v1, the point moves to the left. Continuing on, if the value of Descriptor 3 is less than v3, the point moves to the leaf to the left and is classified as active.

Due to their simplicity and clarity, decision trees tend to lend themselves to interpretation. Unlike many types of machine-learning models, such as neural networks, it is often possible to examine a decision tree and learn something about the physics or chemistry of the problem being modeled.

The decision trees presented in this work were built using a *greedy* algorithm:

1. If all the data fall into the same class, the tree is finished. Otherwise:
2. Select the descriptor which best classifies the data.
3. Split the data into subsets based upon the values of that descriptor.
4. For each subset of the data, proceed to Step 1.

There is no canonical definition of "best classifies the data" to apply in Step 2 of this algorithm. One standard approach is to use concepts introduced by Shannon as part of his development of information theory [15,16]. Central to this method is the concept of the

[2] Clustering is an unsupervised learning algorithm—the known activity values are not used for model building.

[3] The process being described is more accurately termed hierarchical agglomerative clustering.

[4] Standardization translates and scales the values of each descriptor (columns in the data matrix) such that their means are 0 and their deviations are 1. This process helps ensure that all descriptors are treated on equal footing.

[5] Supervised learning uses the activity values of data points when building the model.

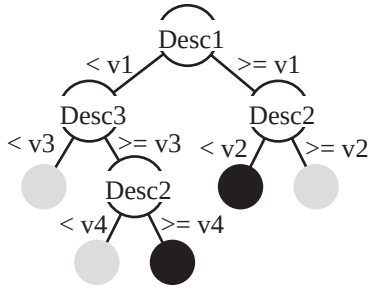

Fig. 1. A sample decision tree. The labels v1, v2, v3 and v4 represent the descriptor boundaries used within each node of the tree (see text). Terminal (leaf) nodes are shaded light/dark if they correspond to active/inactive classifications.

informational entropy of data. The informational entropy of descriptor D_i is defined to be

$$\text{Entropy}(D_i) = \sum_{j=1}^{N_i} -p_{ij}\log_2(p_{ij}), \quad (1)$$

where N_i is the number of values that can be adopted by descriptor i, and p_{ij} is the fraction of the examples in which descriptor i adopts value j.

Informational entropy values range from $\text{Entropy}(D_i) = 0.0$ for a data set where D_i always takes on the same value to $\text{Entropy}(D_i) = 1.0$ when every possible value is equally likely. Given a set of activity values, A, the *information gain* associated with each descriptor can be calculated

$$\text{Gain}(D_i, A) = \text{Entropy}(D_i) - \sum_{v=1}^{N_A} \frac{|D^v|}{|D|}\text{Entropy}(D_i^v), \quad (2)$$

where N_A is the number of possible activity values, D^v is the subset of the data where $A = v$ and $\text{Entropy}(D_i^v)$ is the informational entropy of descriptor i in that subset.

$\text{Gain}(D_i, A)$, the expected reduction in the informational entropy when the descriptor i is used to divide the data set, is a rigorous measure of how well a given descriptor distinguishes between the activity classes. Application of the information gain to select descriptors in decision trees leads to the ID3 algorithm [6]. The decision trees presented here were generated using a modification of ID3, which allows real-valued descriptors to be used. This approach, similar to the methods described in Ref. [17,18], selects quantization bounds for each descriptor D_i which maximize $\text{Gain}(D_i, A)$.

A good introduction to decision trees and the algorithms used to build them can be found in Ref. [6]. Ref. [19] has a recent survey of advanced techniques and the state of the art.

2.4. Evaluation of model quality

How do we know how good our models are? There are probably as many ways to answer this question as there are modelers. We consider two components of model quality: accuracy and overfitting.

The accuracy of a predictive model can be assessed by means of the *accuracy ratio* (E). We define this ratio as: $E = A_{\text{true}}/A_{\text{guess}}$, where A_{true} denotes the accuracy of a model and A_{guess} denotes the accuracy to be expected by randomly guessing. For example, in a two-class problem (i.e., where we are predicting "active" or "inactive") guessing has a 50% chance of being correct, so $A_{\text{guess}} = 0.5$. Higher values of E indicate higher quality models.

Unlike accuracy, where more is better, *overfitting* is something we want to minimize. Overfitting—memorization of the training data—is caused by using too few data points to train models with too many adjustable parameters. Models that are significantly overfit do not, in general, make accurate predictions when applied to new data. A familiar example of a model with overfitting problems is a cubic equation fit to 3 points: the fit to the training data is perfect, but the model clearly is not going to generalize well. To assess the degree of overfitting exhibited by predictive models, we used *shuffle tests*, which measure the ability of the model to overfit the data. In a shuffle test the activity values of the data set are randomly shuffled (permuted); this operation removes all physical connection between the descriptors and activity.[6] A model is then built using the same parameters used for the true model. The accuracy of this "randomized" model (A_{random}) is used to calculate an *overfitting ratio* (R): $R = A_{\text{random}}/A_{\text{true}}$. This R value is a measure of the degree of overfitting of the predictive model: higher quality models have low R values.

Another standard technique for determining degree of overfitting is the *hold-out test*, where the model is presented with new data (not used to construct it) in order to see how well it generalizes. Because the data sets used here are too small to for hold-out tests to provide meaningful results, we relied on shuffle tests to provide a measure of overfitting.

2.5. Descriptor calculation

The electronic structure calculations were performed using the linear muffin-tin orbital (LMTO) method [20–23] within the local spin-density approximation [24]. All calculations were checked for convergence of energies, orbital moments and magnetic moments with respect to the number of k points used in the reciprocal space integrations [25]. Atomic-sphere radii used in the calculations were chosen using an automated procedure. The program used was TB-LMTO version 4.7 [26]. Where required, chemical hardness [27,28] values for the transition metals and their alloys were calculated from

[6] One hopes that there was such a connection in the first place!

Table 1
Descriptors used to characterize the ordered transition metals and alloys

	Name	Description
1	has3d?	Nonzero if any of the metals present has valence 3*d* electrons.
2	has4d?	Nonzero if any of the metals present has valence 4*d* electrons.
3	has5d?	Nonzero if any of the metals present has valence 5*d* electrons.
4	ElConc	Valence electron concentration (number of valence electrons in unit cell/volume of the unit cell).
5	AtVol	Available atomic volume (volume of unit cell/number of atoms in unit cell).
6	Hardness	Calculated chemical hardness of the metal or alloy.

the LMTO density of states using a previously reported procedure [12].

The calculated atomic energy levels used for the Max ΔE_d descriptor were taken from the "Atomic Reference Data for Electronic Structure Calculations" database provided by the National Institute of Standards and Technology [29]. The energy levels used were calculated using the local spin-density approximation.

3. Results and discussion

3.1. Predicting ferromagnetism in ordered transition metal alloys

The data set used to build predictive models for ordered transition metals and their alloys consists of 61 elements and binary alloys. Structural information and data on the existence of ferromagnetism in these phases were obtained from standard sources [30,31]. The compounds were characterized using a set of six descriptors Table 1. With the exception of *Hardness*, the calculation of which is discussed above, the values of these descriptors were generated arithmetically using values looked up in a database.

As we mentioned above, there is no optimal strategy for selecting descriptors. We arrived at the set in Table 1 using domain knowledge and Occam's Razor as well as intuition. Much is known about the electronic basis of ferromagnetism, and we used that knowledge to guide us towards descriptors that were relevant to the underlying physics. In addition, we have limited computational resources, so we further biased our pool by choosing descriptors that were easy to calculate. Physical relevance trumped ease of computation in some cases; for example, chemical hardness is not trivial to calculate but is known to be strongly related to ferromagnetism, so we included it. Yet another tool that we used was the method of structure maps, described below. In the end, though, it was ultimately our chemical intuition that determined exactly which descriptors we employed.

The first indication that Table 1 has a good set of descriptors for modeling this data set can be seen by plotting calculated hardness versus available atomic volume and coloring the points by whether or not they are ferromagnetic, Fig. 2. Property maps of this sort are analogous to structure maps [32] and quantum structural diagrams, and can sometimes be used to identify important descriptors [33,34]. Fig. 2 shows a clearly delineated region of descriptor space that contains only ferromagnetic compounds.

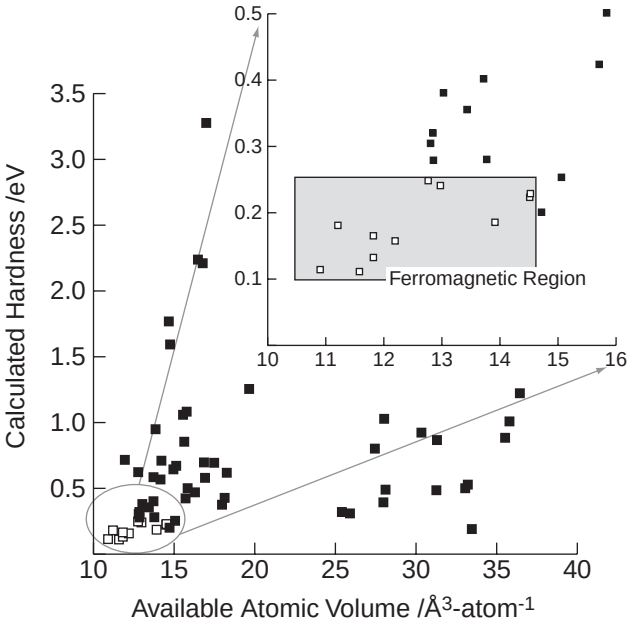

Fig. 2. Property map for ferromagnetism in ordered binary transition metal alloys. Dark-/light-colored points correspond to nonmagnetic/ferromagnetic phases, respectively.

Decision trees classify data by dividing descriptor space into discrete regions (each node in the tree can be viewed as introducing one or more hyperplanes which segment descriptor space), so it is no surprise that the tree built using this data set, shown in Fig. 3, performs well.

The decision tree of Fig. 3 classifies 100% of the data points correctly. Because random guessing would provide 50% accuracy, this corresponds to $E = 2.0$. Note that the tree-building algorithm has automatically selected the two most important descriptors (*Hardness* and *AtVol*) from the pool of six available. Because the decision tree reproduces the ferromagnetic region seen in Fig. 2 and uses only two descriptors for the 61 data points, we are not concerned about this model overfitting the data set, and so no shuffle test was performed.

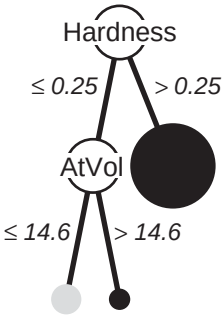

Fig. 3. Decision tree for prediction of ferromagnetism in ordered binary alloys. Leaf nodes are colored dark/light to indicate that they classify alloys as nonmagnetic/ferromagnetic. The relative sizes of the leaves are determined by the number of alloys that they contain.

3.2. Predicting ferromagnetism in disordered transition metal alloys

Many industrially interesting transition metal alloys are not ordered, so it would be useful to have a predictive model that is capable of making predictions for disordered alloys as well as ordered ones. To build such a model we added 42 binary alloys to the previous data set, yielding a new set of 103 alloys. The alloys added are all disordered, but their structure types and unit cell parameters have been measured. The structural data were taken from the same source used above [30].

The descriptor set used for the previous model requires electronic structure calculations in order to calculate the chemical hardness. Therefore, this model cannot be applied to disordered alloys.[7] The other descriptors can all be easily generated from unit cell parameters, so only *Hardness* needs to be replaced. Given the importance of quantum-mechanical exchange in ferromagnetism [10–12,35] a measure of the strength of the exchange interaction was selected. This new descriptor, Max ΔE_d, is calculated using the maximum valence d orbital exchange splitting in the free atoms of the elements that make up the alloy. For example, consider FeCo: the 3d exchange splittings in atomic Fe and Co are 3.53 and 2.76 eV, respectively [29]; Max ΔE_d for FeCo is 3.53 eV. Note that Max ΔE_d is not intended to be an approximation to *Hardness*. It is a distinct quantity which is included here due to its relevance to the underlying physics of ferromagnetism and the fact that it is efficiently calculable for disordered alloys.

A property map for the 103 transition metal alloys using Max ΔE_d and atomic volume as the axes is shown in Fig. 4.

[7] This is not strictly true; it is possible to approximate the electronic structure of disordered systems by using supercell techniques. However, these calculations are very computationally expensive; we would like to find something faster.

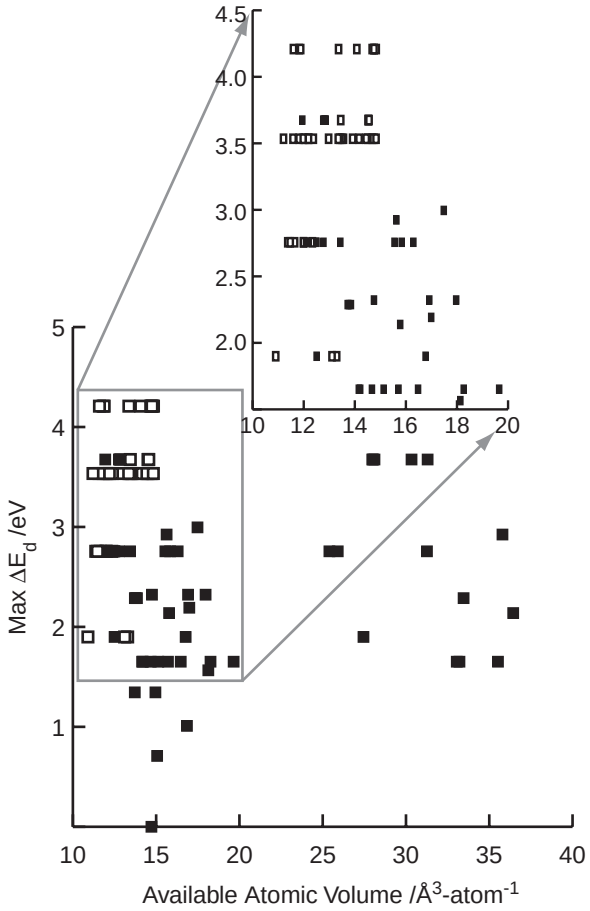

Fig. 4. Property map for ferromagnetism in ordered and disordered binary transition metal alloys. Dark/light points correspond to nonmagnetic/ferromagnetic phases.

Though the ferromagnetic alloys are grouped together to some extent in Fig. 4, they do not show the clean segmentation that was possible with the previous data set (Fig. 2). Two descriptors alone are no longer sufficient to build a good predictive model. In order to find patterns in the space defined by our six descriptors, we employed hierarchical clustering. The results of the clustering are promising, Fig. 5.

One of the first things to leap out from the cluster tree of Fig. 5 is the large grouping of nonmagnetic (red) compounds at the right side of the tree. Tracing up through the diagram, it can be seen that this grouping is the result of two clusters that are differentiated at a very early step from the rest of the data. These two clusters, containing more than 30 nonmagnetic phases, define large volumes of descriptor space where ferromagnetic compounds are *not* found. In the zoomed region at the bottom of Fig. 5, which shows a medium sized cluster containing both magnetic and nonmagnetic phases, one can see that the compounds grouped together in clusters tend to be chemically similar. This suggests that the descriptor set chosen is chemically reasonable.

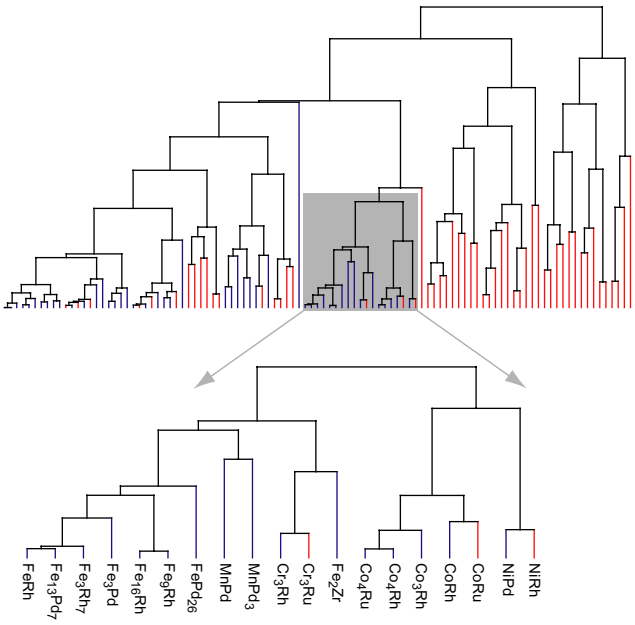

Fig. 5. Top: cluster tree for the binary transition metal alloys. Terminal (leaf) clusters are colored red/blue correspond to nonmagnetic/magnetic alloys. The shaded area is shown in greater detail at the bottom of the Figure, where the leaves have been labeled. The vertical scale of the bottom tree has been altered to clarify the presentation.

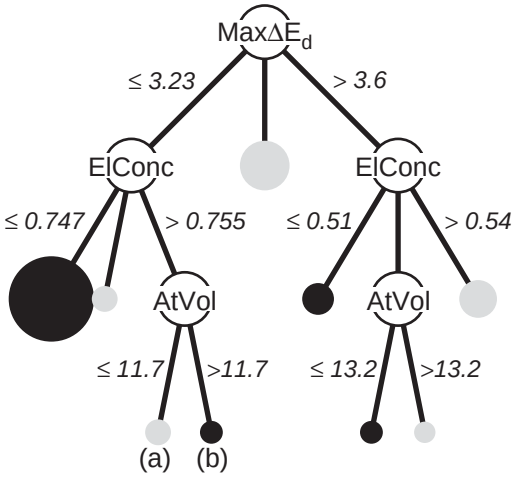

Fig. 6. Decision tree for prediction of ferromagnetism in ordered and disordered binary alloys. Leaf nodes are colored gray/black to indicate that they classify alloys as ferromagnetic/nonmagnetic. The relative sizes of the leaves are determined by the number of alloys that they contain.

Confirmation of the quality and suitability of this descriptor set is provided by the decision tree built using it, Fig. 6.

This relatively simple tree, which uses only three of the six available descriptors, correctly classifies 99% of the alloys in the data set ($E = 1.98$). The misclassified alloy, NiPt, ends up in the nonmagnetic leaf labeled (b) in Fig. 6 instead of the ferromagnetic leaf (a). The magnetic behavior of NiPt changes with ordering (disordered NiPt is ferromagnetic, while the ordered phase is nonmagnetic) and our simple model does not have sufficient resolution to reproduce this phenomenon.

Because the data set is not large enough to allow us to perform a reasonable hold-out test, we must test for overfitting using shuffle tests. To this end, ten shuffle tests were performed and the results averaged. The mean accuracy of these randomized models was 51% with a sample standard deviation of 9%, giving a mean overfitting ratio $R_{avg} = 0.52$. The randomized models display an accuracy which is statistically indistinguishable from that expected by guessing (50%). Overfitting is unlikely to be a problem in the decision tree shown in Fig. 6.

4. Conclusions

The implications of this investigation for the applicability of machine-learning techniques in materials problems are very promising. Using data sets that are tiny on the scale of most machine-learning problems—61 data points for the ordered-phase study, and 103 for the combined ordered–disordered study—we have built models with high predictive accuracy and little overfitting. In addition, we have shown that, using just a handful of descriptors, the very simple and rapid decision tree algorithm can, at least in outline, manage the complexity of ferromagnetism. The details, such as magnetic moments, remain to be addressed.

Can we extrapolate these results to other areas of solid-state physics, such as high-T_c superconductivity? The short answer is: possibly very soon. As mentioned above, the descriptor sets for different problems are going to be different, so the first task will be to find a set of suitable descriptors for the problem. Because ferromagnetism is well understood, it is relatively straightforward to select a good descriptor set. High-T_c superconductivity is not so well understood. Additionally, while there is much experimental data available, there is not yet enough consistently measured data. But this situation is constantly improving. Preliminary studies have shown great promise in classifying known data for superconductors [36]. It is very likely that machine-learning based predictions will become useful in the near future, as more data become available on which to build models, and as more descriptor sets are tested.

Acknowledgments

We would like to thank Professors Roald Hoffmann, Stephen Lee and Frank DiSalvo at Cornell University for inspirational discussions. G.A. Landrum gratefully

acknowledges Prof. Dr. Richard Dronskoswki at the RWTH-Aachen for his support of Dr. Landrum's postdoctoral stay in Aachen and for sparking his interest in ferromagnetism.

References

[1] P. Villars, K. Brandenburg, M. Berndt, et al., Eng. Appl. Artif. Intell. 13 (2000) 497–505.
[2] A. Persidis, Nat. Biotechnol. 18 (2000) 237–238.
[3] T.R. Cundari, J. Deng, H.F. Pop, et al., J. Chem. Inf. Comput. Sci. 40 (2000) 1052–1061.
[4] G. Espinosa, A. Arenas, F. Giralt, J. Chem. Inf. Comput. Sci. 42 (2002) 343–359.
[5] C. Hansch, A. Leo, D.H. Hoekman, et al. (Eds.), Exploring QSAR: Fundamentals and Applications in Chemistry and Biology, American Chemical Society, Washington DC, 1995.
[6] T. Mitchell, Machine Learning, McGraw-Hill, New York, 1997.
[7] V. Cherkassky, F. Mulier, Learning from Data: Concepts, Theory, and Methods, Wiley, New York, 1998.
[8] A. Aharoni, Introduction to the Theory of Ferromagnetism, Clarendon Press, Oxford, 1996.
[9] D. Vollhardt, N. Blümer, K. Held, et al., Adv. Solid State Phys. 38 (1999) 383.
[10] S. Chikazumi, Physics of Ferromagnetism, 2nd Edition, Clarendon Press, Oxford, 1997.
[11] G.A. Landrum, R. Dronskowski, Angew. Chem. Int. Ed. Engl. 38 (1999) 1389–1393.
[12] G.A. Landrum, R. Dronskowski, Angew. Chem. Int. Ed. Engl. 39 (2000) 1560–1585.
[13] H.C. Romesburg, Cluster Analysis for Researchers, Krieger, Malabar, 1990.
[14] V. Kumar, An introduction to cluster analysis for data mining, 1995, http://www-users.cs.umn.edu/~han/dmclass/cluster_survey_10_02_00.pdf.
[15] C.E. Shannon, Bell System Tech. J. 27 (1948) 379–423.
[16] C.E. Shannon, Bell System Tech. J. 27 (1948) 623–656.
[17] U.M. Fayyad, K.B. Irani, Mach. Learn. 8 (1992) 87–102.
[18] U.M. Fayyad, K.B. Irani, in: R. Bajcsy (Ed.), Proceedings of the 13th International Joint Conference on Artificial Intelligence, Morgan-Kaufmann, Los Altos, CA, 1993, pp. 1022–1027.
[19] S.K. Murthy, Data Mining Knowledge Discovery 2 (1998) 345–389.
[20] O.K. Andersen, Phys. Rev. B 12 (1975) 3060.
[21] O.K. Andersen, C. Arcangeli, R.W. Tank, et al., Tight-Binding Approach to Computational Materials Science, MRS, Pittsburgh, 1998.
[22] O.K. Andersen, O. Jepsen, Phys. Rev. Lett. 53 (1984) 2571.
[23] H.L. Skriver, The LMTO Method, Springer, Berlin, 1984.
[24] U.V. Barth, L. Hedin, J. Phys. C 5 (1972) 1629.
[25] P.E. Blöchl, O. Jepsen, O.K. Andersen, Phys. Rev. B 49 (1994) 16223.
[26] G. Krier, O. Jepsen, A. Burkhardt, et al., The TB-LMTO-ASA program, v4.7.
[27] R.G. Pearson, Chemical Hardness: Applications from Molecules to Solids, Wiley-VCH, Weinheim, 1997.
[28] W. Yang, R.F. Parr, Proc. Natl. Acad. Sci. USA 82 (1985) 6723–6726.
[29] S. Kotochigova, Z. Levine, E. Shirley, et al., National Institute of Standards and Technology, Atomic Reference Data for Electronic Structure Calculations, v2000.
[30] P. Villars, L.D. Calvert, Pearson's Handbook of Crystallographic Data for Intermetallic Phases, 2nd Edition, ASM International, Ohio, 1991.
[31] K. Adachi, D. Bonnenberg, J.J.M. Franse, et al., Landolt Börnstein Group III/19: Magnetic Properties of Metals, Springer, Berlin, 1986.
[32] J.K. Burdett, J. Alloys, Compounds 197 (1993) 281–289.
[33] K.M. Rabe, J.C. Phillips, P. Villars, et al., Phys. Rev. B 45 (1992) 7650–7676.
[34] K.M. Rabe, MRS Bull. 18 (1993) 31–37.
[35] H. Ibach, H. Lüth, Solid-State Physics, 2nd Edition, Springer, Berlin, 1995.
[36] G.A. Landrum, H. Genin, Mater. Res. Soc. Symp. Proc. 700 (2002), http://www.mrs.org/members/proceedings/fall2001/s/S7_6.pdf.