Prediction of the Atomization Energy of Molecules Using Coulomb Matrix and Atomic Composition in a Bayesian Regularized Neural Networks

Alain B. Tchagang^(✉) and Julio J. Valdés

Digital Technologies Research Centre, National Research Council Canada,
M-50, 1200 Montréal Road, Ottawa, ON K4A 0S2, Canada
{alain.tchagang, julio.valdes}@nrc-cnrc.gc.ca

Abstract. Exact calculation of electronic properties of molecules is a fundamental step for intelligent and rational compounds and materials design. The intrinsically graph-like and non-vectorial nature of molecular data generates a unique and challenging machine learning problem. In this paper we embrace a learning from scratch approach where the quantum mechanical electronic properties of molecules are predicted directly from the raw molecular geometry, similar to some recent works. But, unlike these previous endeavors, our study suggests a benefit from combining molecular geometry embedded in the Coulomb matrix with the atomic composition of molecules. Using the new combined features in a Bayesian regularized neural networks, our results improve well-known results from the literature on the QM7 dataset from a mean absolute error of 3.51 kcal/mol down to 3.0 kcal/mol.

Keywords: Atomization energy · Atomic composition · Bayesian regularization · Coulomb matrix · Electronic properties · Molecules · Neural networks

1 Introduction

Finding new molecules, compounds or materials with desired properties is strategic to the innovation and progress of many chemical, agrochemical and pharmaceutical industries. One of the major challenges consists of making quantitative estimates in the chemical compound space at moderate computational cost (milliseconds per compound or faster). Currently only high level quantum-chemistry calculations, which can take days per molecule depending on property and system, yield the desired chemical accuracy of 1 kcal/mol required for computational molecular and material design [1].

Recent technological advances have shown that data-to-knowledge approaches are beginning to show enormous promise within materials science. Intelligent exploration and exploitation of the vast materials property space has the potential to alleviate the cost, risks, and time involved in trial-by-error approach experiment cycles used by current techniques to identify useful compounds [2]. For example, obtaining atomization energies from the Schrödinger equation solver is computationally

expensive and, as a consequence, only a fraction of the molecules in the chemical compound space can be labeled. By training a machine learning algorithm on the few label ones, the trained quantum mechanics machine learning (QM/ML) model can be used to generalize from these few data points to unseen molecules [3]. One of the central questions in QM/ML is how to represent molecules in a way that makes prediction of molecular properties feasible and accurate [4]. This question has already been extensively discussed in the cheminformatics literature, and many so-called molecular descriptors exist [5]. Unfortunately, they often require a substantial amount of domain knowledge and engineering. Furthermore, they are not necessarily transferable across the whole chemical compound space [1].

In this paper, we follow a more direct approach introduced in [6], and adopted by several other authors. We learn the mapping between the molecule and its atomization energy from scratch using the Coulomb matrix as a low-level molecular descriptor [6, 7]. Coulomb matrix is invariant to translation and rotation but not to permutations or re-indexing of the atoms. Methods to tackle this issue have been proposed. Examples include Coulomb sorted eigenspectrum [1], Coulomb sorted L2 norm of the matrix's columns [7], Coulomb bag of bonds [8], and random Coulomb matrices [1, 3]. Our study extends the work of [1, 3, 6, 7]. Unlike these previous authors, we show that by combining the molecular geometry embedded in the Coulomb matrix with atomicity or atomic composition of molecules (i.e. atom counts of each type in a molecule), the outcome of the QM/ML models can be significantly improved.

To test this new hypothesis, five representations are constructed: (1) sorted Coulomb matrix, (2) atomic composition of molecules, (3) Coulomb eigenspectrum, (4) the combination of the Coulomb eigenspectrum and the atomic composition of molecules, and (5) the combination of the sorted Coulomb matrix and the atomic composition of molecules. Each one is used as input to a well-defined multilayer Bayesian regularized neural networks [9–13]. Results obtained using the combination of either the sorted Coulomb matrix or the Coulomb eigenspectrum with the atomic composition showed better predictions by a difference of more than 1.5 kcal/mol compared to when the sorted Coulomb matrix or the Coulomb eigenspectrum is used solely. More interestingly, the mean absolute error (MAE) = 3.0 kcal/mol obtained in this study is lower than the 3.51 kcal/mol well-known results obtained in [1, 3]. These results confirm the efficacy of using the atomic composition of molecules in a QM/ML model for their electronic properties predictions. Furthermore, the Bayesian regularized neural network is shown to be a suitable candidate for the modeling of molecular data.

The rest of this paper is organized as follows. In Sect. 2, the dataset used in this study is described. Section 3 provides a detailed description of the proposed method. Section 4 presents the results and Sect. 5 the conclusions.

2 Materials

The QM7 dataset used in this study is a subset of the GDB-13 dataset [14]. The version used here is the one published in [7] consisting of 7102 small organic molecules and their associated atomization energy. These molecules are composed of a maximum of 23 atoms. Molecules are converted to a suitable Cartesian coordinates representation

using universal forcefield method [15] as implemented in the software OpenBabel [16]. Atomization energies are calculated for each molecule and ranging from -800 to -2000 kcal/mol. Note that all the 7102 molecules are unique and there are no isomers in the set.

3 Methods

Sorted Coulomb matrix, Coulomb eigenspectrum and atomic composition of each molecule are computed using the atomic coordinates and the chemical formulae of each molecule respectively as described in the QM7 dataset. Next, atomic composition, Coulomb eigenspectrum and sorted Coulomb matrix are either combined or used separately as input to a regularized Bayesian neural network for the prediction of the atomization energy.

3.1 Atomicity, Atom Counts or Atomic Composition

Let's define $\Omega = \{\Omega_1, \Omega_2, \dots, \Omega_m, \dots, \Omega_M\}$, the set of possible molecules in the chemical compound space (CCS). By construction, this space is very large. In this study, we will assume that it is bounded by M . Let's define A the set of unique atoms that make Ω . A is bounded by K and it is defined as: $A = \{A^1, A^2, \dots, A^k, \dots, A^K\}$. Let's define a chemical operator “.” that combines atoms among them in a specific numbers α_k^m and according to the laws of chemistry to form a stable molecule Ω_m . The chemical formulae of Ω_m can be written as: $\Omega_m = \alpha_1^m A^1 . \alpha_2^m A^2 . \dots . \alpha_k^m A^k . \dots . \alpha_K^m A^K$, or as in chemical textbook.

$$\Omega_m \equiv A_{\alpha_1^m}^1 A_{\alpha_2^m}^2 \dots A_{\alpha_k^m}^k \dots A_{\alpha_K^m}^K \quad (1)$$

The atomic composition (AC) of molecule Ω_m in the atomic space $[A^1 A^2 \dots A^k \dots A^K]$ is defined as $[\alpha_1^m \alpha_2^m \dots \alpha_k^m \dots \alpha_K^m]$, where α_k^m is a positive integer that represents the number of atom A^k in molecule Ω_m . The AC of the M molecules in the atomic space $[A^1 A^2 \dots A^k \dots A^K]$ can be viewed as an $M \times K$ matrix α , Eq. (2).

$$\alpha = \begin{bmatrix} \Omega_1 \\ \Omega_2 \\ \vdots \\ \Omega_m \\ \vdots \\ \Omega_M \end{bmatrix} = \begin{bmatrix} \alpha_1^1 & \alpha_2^1 & \dots & \alpha_k^1 & \dots & \alpha_K^1 \\ \alpha_1^2 & \alpha_2^2 & \dots & \alpha_k^2 & \dots & \alpha_K^2 \\ \vdots & \vdots & \dots & \vdots & \dots & \vdots \\ \alpha_1^m & \alpha_2^m & \dots & \alpha_k^m & \dots & \alpha_K^m \\ \vdots & \vdots & \dots & \vdots & \dots & \vdots \\ \alpha_1^M & \alpha_2^M & \dots & \alpha_k^M & \dots & \alpha_K^M \end{bmatrix} \quad (2)$$

Row $\alpha(m,:)$ of α corresponds to the AC of the m^{th} molecule (Ω_m). Column $\alpha(:,k)$ corresponds to the number of atom A^k in each molecule of Ω . K is an integer and correspond to the number of unique atoms that makes Ω . For example, given a set of

seven molecules: $\Omega = \{\text{CH}_4, \text{C}_2\text{H}_2, \text{C}_3\text{H}_6, \text{C}_2\text{NH}_3, \text{OC}_2\text{H}_2, \text{ONC}_3\text{H}_3, \text{SC}_3\text{NH}_3\}$. The set of unique atoms that makes Ω is $A = \{\text{C}, \text{H}, \text{N}, \text{O}, \text{S}\}$. The matrix α is then:

$$\alpha = \begin{bmatrix} \text{CH}_4 \\ \text{C}_2\text{H}_2 \\ \text{C}_3\text{H}_6 \\ \text{C}_2\text{NH}_3 \\ \text{OC}_2\text{H}_2 \\ \text{ONC}_3\text{H}_3 \\ \text{SC}_3\text{NH}_3 \end{bmatrix} = \begin{bmatrix} 1 & 4 & 0 & 0 & 0 \\ 2 & 2 & 0 & 0 & 0 \\ 3 & 6 & 0 & 0 & 0 \\ 2 & 3 & 1 & 0 & 0 \\ 2 & 2 & 0 & 1 & 0 \\ 3 & 3 & 1 & 1 & 0 \\ 3 & 3 & 1 & 0 & 1 \end{bmatrix} \quad (3)$$

It is obvious that this representation is not unique. That is two molecules with identical atomic composition may have different electronic properties. Isomers are great examples in this case. They are compound with the same molecular formulas but that are structurally different in some way, and they can have different chemical, physical and biological properties [17]. It is also worth to note that such molecular representation had been explored in the past in quantitative structure activity relationship and correspond to a different form of the Atomistic index developed by Burden [13].

3.2 Coulomb Matrix

The Coulomb matrix (CM) has recently been widely used as molecular descriptors in the QM/ML models [1, 3, 6, 7]. Given a molecule its Coulomb matrix $\text{CM} = [c_{ij}]$ is defined by Eq. (4).

$$c_{ij} = \begin{cases} 0.5Z_i^{2.4} & \text{for } i = j \\ \frac{Z_i Z_j}{\|R_i - R_j\|} & \text{for } i \neq j \end{cases} \quad (4)$$

Z_i is the atomic number of atom i , and R_i is its position in atomic units [7]. CM is of size $I \times I$, where I corresponds to the number of atoms in the molecule. It is symmetric and has as many rows and columns as there are atoms in the molecule. As we mentioned earlier, the Coulomb matrix is invariant to rotation, translation but not to permutation of its atoms. Several techniques to tackle this issue have been explored in the literature. Examples include sorted Coulomb matrix and Coulomb eigenspectrum.

Sorted Coulomb Matrix (SCM). This approach sorts the CMs by descending order with respect to the norm-2 of their columns and simultaneously permuting their rows and columns accordingly. After the ordering step and given the symmetry of these matrices, it is customary to only consider their lower triangular part [6, 7], and to unfold them row-wise in a 1-dimensional (1D) vector of length $L = \sum_{i=0}^I (I - i)$, where I here corresponds to the number of atoms of the largest molecule. In this study, the 1D vector is called the SCM signal $x(m,:) = x_m[l]$, with $l = 1$ to L and m

corresponds to molecule Ω_m . For a set of M molecules, their 1D SCM signals can be organized in an $M \times L$ matrix x :

$$x = \begin{bmatrix} x_{11} & x_{12} & \dots & x_{1l} & \dots & x_{1L} \\ x_{21} & x_{22} & \dots & x_{2l} & \dots & x_{2L} \\ \vdots & \vdots & \dots & \vdots & \dots & \vdots \\ x_{m1} & x_{m2} & \dots & x_{ml} & \dots & x_{mL} \\ \vdots & \dots & \dots & \vdots & \dots & \vdots \\ x_{M1} & x_{M2} & \dots & x_{Ml} & \dots & x_{ML} \end{bmatrix} \quad (5)$$

The m^{th} row of x represents the 1D SCM signal of the m^{th} molecule. Given that molecules have different number of atoms, the short ones are padded with zeros so that all the 1D SCM signals have the same length L .

Coulomb Eigenspectrum (CES). The Coulomb eigenspectrum [6, 7] is obtained by solving the eigen value problem $Cv = \lambda v$, under constraint $\lambda_i \geq \lambda_{i+1}$ where $\lambda_i > 0$. The spectrum $(\lambda_1, \dots, \lambda_l)$ is used as the representation and it corresponds to a 1D signal: $z(m,:) = z_m[n]$, with $n = 1$ to N and m corresponds to molecule Ω_m . For a set of M molecules, their 1D CES signals can be organized in an $M \times N$ matrix z :

$$z = \begin{bmatrix} z_{11} & z_{12} & \dots & z_{1n} & \dots & z_{1N} \\ z_{21} & z_{22} & \dots & z_{2n} & \dots & z_{2N} \\ \vdots & \vdots & \dots & \vdots & \dots & \vdots \\ z_{m1} & z_{m2} & \dots & z_{mn} & \dots & z_{mN} \\ \vdots & \dots & \dots & \vdots & \dots & \vdots \\ z_{M1} & z_{M2} & \dots & z_{Mn} & \dots & z_{MN} \end{bmatrix} \quad (6)$$

The m^{th} row of z represents the 1D CES signal of the m^{th} molecule. Given that molecules have different number of atoms, the short ones are padded with zeros so that all the 1D CES signals have the same length N .

3.3 Input of the QM/ML Model

Let's define X as the input to the neural network defined below. In order to test the usefulness of the AC in the prediction of the electronic properties of molecules, we have considered five different inputs and compared them against each other. The five inputs are: $X = \alpha$ (only the AC is used), $X = z$ (only the CES is used), $X = x$ (only the SCM is used), $X = [\alpha \ z]$ (AC and CES are combined and used as inputs), and finally $X = [\alpha \ x]$ (AC and SCM are combined and used as inputs). By combining AC, CES and SCM, taking the Z-scores of X prior to its utilization as input to the ML model becomes an obvious choice. The Z-score of X will return a matrix of same size X' , where each column of X' has mean 0 and a standard deviation of 1 [18].

3.4 Output – Atomization Energy of Molecules

The output to the QM/ML is the atomization energy E . It quantifies the potential energy stored in all chemical bonds. As such, it is defined as the difference between the potential energy of a molecule and the sum of potential energies of its composing isolated atoms. The potential energy of a molecule is the solution to the electronic Schrödinger equation $H\Phi = E\Phi$, where H is the Hamiltonian of the molecule and Φ is the state of the system. The atomization energy of molecules are organized in an $M \times 1$ column vector $y = [y_1 \ y_2 \ \dots \ y_m \ \dots \ y_M]^T$. The superscript T indicates the transpose operator. The entry y_m is a real number that corresponds to the atomization energy of the m^{th} molecule.

3.5 Bayesian Regularized Neural Networks

Neural networks (NN) are universal function approximators that can be applied to a wide range of problems such as classification and model building. It is already a mature field within machine learning and there are many different NN paradigms. Multilayer feed-forward networks are the most popular and a large number of training algorithms have been proposed. Compared to other non-linear techniques, in multilayer NNs, the measure of similarity is learned essentially from data and implicitly given by the mapping onto increasingly many layers. In general, NNs are more flexible and make fewer assumptions about the data. However, it comes at the cost of being more difficult to train and regularize [3]. In this paper, we used the Bayesian regularization method to train our NNs [9–12].

Bayesian methods are optimal methods for solving learning problems. Any other method not approximating them should not perform as well on average. They are very useful for comparison of data models as they automatically and quantitatively embody “Occam’s Razor” [19]. Complex models are automatically self-penalizing under Bayes’ Rule. Bayesian methods are complementary to NNs as they overcome the tendency of an over flexible network to discover nonexistent, or overly complex, data models.

Unlike a standard back-propagation NN training method where a single set of parameters (weights, biases, etc.) are used, the Bayesian approach to NN modeling considers all possible values of network parameters weighted by the probability of each set of weights. Bayesian inference is used to determine the posterior probability distribution of weights and related properties from a prior probability distribution according to updates provided by the training set D using the Bayesian regularized NN model, H_i . Where orthodox statistics provide several models with several different criteria for deciding which model is best, Bayesian statistics only offer one answer to a well-posed problem.

$$P(w|D, H_i) = \frac{P(D|w, H_i)P(w|H_i)}{P(D|H_i)} \quad (7)$$

Bayesian methods can simultaneously optimize the regularization constants in NNs, a process which is very laborious using cross-validation [9].

4 Results and Discussions

As we mentioned earlier, the version of the QM7 dataset used in this study is the one published in [7] and it is composed of $M = 7102$ molecules and contains up to five types of atoms: Carbon (C), Hydrogen (H), Oxygen (O), Nitrogen (N), and Sulfur (S). Therefore the set of unique atoms is $A = \{C, H, N, O, S\}$. The matrix α is of size $M \times K = 7102 \times 5$. The largest molecule is made of $I = 23$ atoms. Thus the CES matrix z is of size $M \times N = 7102 \times 23$, and the SCM matrix x is of size $M \times L = 7102 \times 276$, because $L = \sum_{i=0}^{23} (23 - i) = 276$. The column vector of atomization energy y is of size 7102×1 . The QM7 dataset is randomly divided into 80% training and 20% testing sets. Performance is measured using the root mean square error (RMSE), Eq. (8), the mean absolute error (MAE), Eq. (9), and the Pearson correlation coefficient r_{ppe} , Eq. (10).

$$RMSE = \sqrt{\frac{1}{M} \sum_{m=1}^M (y_m - y_m^e)^2} \quad (8)$$

$$MAE = \frac{1}{M} \sum_{m=1}^M |y_m - y_m^e| \quad (9)$$

$$r_{ppe} = \frac{\sum_{m=1}^M (y_m - \bar{y})(y_m^e - \bar{y}^e)}{\sqrt{\sum_{m=1}^M (y_m - \bar{y})^2} \sqrt{\sum_{m=1}^M (y_m^e - \bar{y}^e)^2}} \quad (10)$$

4.1 Results

We used the Matlab implementation of the regularized Bayesian network to model the relationship between the inputs (X) and the output (y). Figure 1 for example shows the Matlab architecture of one of the networks used.

Fig. 1. Matlab representation of one hidden layer neural networks with 17 neurons when the SCM is used as input (276 inputs and 1 output).

In the literature, there is no clear and rational approach on how to select the number of neurons and the number of hidden layers of a NN. The middle ground is usually to select an architecture that will neither under-fit nor over-fit the model. In this study we tested several architecture based on some empirical observations also coming from the literature with the goal for avoiding under-fitting and overfitting of the model. Table 1 shows the results obtained using different NN architectures with the data partitioned into 90% training and 10% for validation and testing.

Table 1. Statistics of the results using three different network architectures, each trained using Bayesian regularization. MAE and RMSE are in kcal/mol.

Network architecture	Input*	Statistics		
		MAE	RMSE	r_{ppe}
[17]	α	13.82	18.05	0.9967
	z	9.40	12.29	0.9985
	x	5.02	6.72	0.9995
	$[\alpha \text{ z}]$	8.38	11.09	0.9988
	$[\alpha \text{ x}]$	3.70	5.0	0.9997
[16 $\times$ 8 $\times$ 4]	α	13.82	18.05	0.9967
	z	8.70	11.35	0.9987
	x	4.83	6.58	0.9996
	$[\alpha \text{ z}]$	7.59	10.06	0.9990
	$[\alpha \text{ x}]$	3.42	4.73	0.9998
[18 $\times$ 9 $\times$ 3]	α	13.80	18.04	0.9967
	z	8.57	11.16	0.9987
	x	4.40	5.95	0.9996
	$[\alpha \text{ z}]$	7.38	9.80	0.9990
	$[\alpha \text{ x}]$	3.0	4.22	0.9998

* α = Atomic Composition (AC), z = Coulomb Eigen Spectrum (CES), x = Sorted Coulomb Matrix (SCM)

The results obtained show that the association of the AC with either CES or SCM significantly improved the prediction accuracy. For example, with the three hidden layer network [18 $\times$ 9 $\times$ 3], the MAE goes from 13.80 kcal/mol when only the AC is used, to 4.40 kcal/mol when only the SCM is used down to 3.0 kcal/mol when the AC is combined with SCM. Similar observation is made when CES is combined with the AC (see Table 1). These results suggest that AC represents an interesting feature for the predictions of the electronic properties of molecules. Furthermore, the MAE = 3.0 kcal/mol obtained is lower than the MAE of 9.9 kcal/mol [6, 7] using kernel ridge regression and MAE of 3.51 kcal/mol obtained in [1, 3] using a multilayer NN associated with random coulomb matrices and a binarization scheme to augment the data in order to use a more complex multilayer neural network than the one used in this study. Clearly, the QM7 dataset result of [1, 3] is improved by a factor difference of 0.5 kcal/mol in this study. Our result is close to the acceptable 1 kcal/mol chemical accuracy.

4.2 Discussions

Predicting molecular energies quickly and accurately across the CCS is an important problem as the QM calculations take days and do not scale well to more complex molecules. ML is a good candidate for solving this problem as it encourages the framework to focus on solving the problem of interest rather than solving the more general Schrödinger equations. In this paper, we have developed further the learning-from-scratch approach initiated in [6] and provided a new ingredient for learning a successful mapping between raw molecular geometries and atomization energies. Our results suggest important discoveries and open new venues for future research.

Atomicity, Atom Counts or Atomic Composition Represents an Interesting Feature for QM/ML Models. Atomic composition (AC), i.e. atom counts of each type in a molecule is a representation that does not contain any molecular structural information. But our analysis suggests a correlation between the AC representation and the atomization energy. The combination of AC with CES or SCM yields a new molecular representation which inherits all the properties of either CES or SCM representation respectively. Even though the AC representation is not unique (case of isomers as we mentioned earlier), by combining it with the SCM for example, the pair [AC SCM] inherit all the properties of SCM and becomes a representation that is uniquely defined, invariant to rotation, translation and re-indexing of the atoms, given that the SCM had already been sorted in decreasing order to tackle the non-invariance to atom re-indexing. Similar observation can be made with the CES.

Bayesian Regularized Neural Networks are Suited for Molecular Data. The Bayesian regularization approach used in this study seems to fit molecular data very well. Similar observation was made in [12] when developing quantitative structure activity relationship (QSAR) model of compounds active at the benzodiazepine and muscarinic receptors. The results obtained here further prove the point that Bayesian regularized neural networks possess several properties useful for the analysis of molecular data. One advantage of the Bayesian regularized neural networks is that the number of effective parameters used in the model is less than the number of weights, as some weights do not contribute to the models. This minimizes the likelihood of overfitting. The concerns about overfitting and overtraining are also removed by this method so that the production of a definitive and reproducible model is attained [9–12].

5 Conclusions

In this study, we show that by combining the atomic composition of molecules with their Coulomb matrix representation, the output of the quantum mechanics machine learning model can be significantly improved. Using the QM7 dataset as a test case, our results show a decrease by a difference of 1.5 kcal/mol when the sorted Coulomb matrix representation is combined with the atomic composition compared to when the sorted Coulomb matrix is used alone. Furthermore, our results improve well-known results from the literature on the QM7 dataset from a mean absolute error of 3.51 kcal/mol down to 3.0 kcal/mol. These results suggest that the atomic composition

of molecules contain interesting information useful for quantum mechanics machine learning model and should not be neglected.

References

1. Montavon, G., et al.: Learning invariant representations of molecules for atomization energy prediction. In: Pereira, F., Burges, C.J.C., Bottou, L., Weinberger, K.Q. (eds.) *Proceedings of the 25th International Conference on Neural Information Processing Systems - (NIPS 2012)*, Curran Associates Inc., USA, vol. 1, pp. 440–448 (2012). doi:2999134.2999184
2. Xue, D., Balachandran, P.V., Hogden, J., Theiler, J., Xue, D., Lookman, T.: Accelerated search for materials with targeted properties by adaptive design. *Nat. Commun.* **15**(7), 11241 (2016). <https://doi.org/10.1038/ncomms11241>
3. Montavon, G., et al.: Machine learning of molecular electronic properties in chemical compound space. *New J. Phys.* **15**(9), 095003 (2013). <https://doi.org/10.1088/1367-2630/15/9/095003>
4. Von Lilienfeld, O.A., Tuckerman, M.E.: Molecular grand-canonical ensemble density functional theory and exploration of chemical space. *J. Chem. Phys.* **125**(15), 154104 (2006). <https://doi.org/10.1063/1.2338537>
5. Mauri, A., Consonni, V., Todeschini, R.: Molecular descriptors. In: Leszczynski, J. (ed.) *Handbook of Computational Chemistry*. Springer, Dordrecht (2016). https://doi.org/10.1007/978-94-007-6169-8_51-1
6. Rupp, M., Tkatchenko, A., Müller, K.-R., von Lilienfeld, O.A.: Fast and accurate modeling of molecular atomization energies with machine learning. *Phys. Rev. Lett.* **108**, 058301 (2012). <https://doi.org/10.1103/physrevlett.108.058301>
7. Rupp, M.: Machine learning for quantum mechanics in a nutshell. *Int. J. Quan. Chem.* **115**, 1058–1073 (2015). <https://doi.org/10.1002/qua.24954>
8. Hansen, K., et al.: Machine learning predictions of molecular properties: accurate many-body potentials and nonlocality in chemical space. *J. Phys. Chem. Lett.* **6**, 2326–2331 (2015). <https://doi.org/10.1021/acs.jpclett.5b00831>
9. MacKay, D.J.C.: A practical bayesian framework for backprop networks. *Neural Comput.* **4**, 415–447 (1992). <https://doi.org/10.1162/neco.1992.4.3.448>
10. Mackay, D.J.C.: Probable networks and plausible predictions - a review of practical bayesian methods for supervised neural networks. *Comput. Neural Sys.* **6**, 469–505 (1995). https://doi.org/10.1088/0954-898x_6_3_011
11. Mackay, D.J.C.: Bayesian interpolation. *Neural Comput.* **4**, 415–447 (1992). <https://doi.org/10.1162/neco.1992.4.3.415>
12. Buntine, W.L., Weigend, A.S.: Bayesian back-propagation. *Complex Sys.* **5**, 603–643, (1991). <https://doi.org/10.1007/s00138-012-0450-4>
13. Burden, F.R.: Robust QSAR models using bayesian regularized neural networks. *J. Med. Chem.* **42**, 3183–3187 (1999). <https://doi.org/10.1021/jm980697n>
14. Blum, L.C., Raymond, J.-L.: 970 million druglike small molecules for virtual screening in the chemical universe database GDB-13. *J. Am. Chem. Soc.* **131**, 8732–8733 (2009). <https://doi.org/10.1021/ja902302h>
15. Rappé, A.K., Casewit, C.J., Colwell, K.S., Goddard, W.A., Skiff, W.M.: UFF, a full periodic table force field for molecular mechanics and molecular dynamics simulations. *J. Am. Chem. Soc.* **114**(25), 10024–10035 (1992). <https://doi.org/10.1021/ja00051a040>
16. Guha, R., et al.: The blue obelisk, interoperability in chemical informatics. *J. Chem. Inf. Model.* **46**(3), 991–998 (2006). <https://doi.org/10.1021/ci050400b>

17. Gorzynski, S.J.: General Organic and Biological Chemistry, 2nd edn, p. 450. McGraw-Hill, New York (2010)
18. Aho, K.A.: Foundational and Applied Statistics for Biologists, 1st edn. CRC Press, Boca Raton (2014)
19. Hugh, G., Gauch, H.G.: Scientific Method in Practice. Cambridge University Press, Cambridge (2003)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.