

# Fast and Accurate Modeling of Molecular Atomization Energies with Machine Learning

Matthias Rupp,<sup>1,2</sup> Alexandre Tkatchenko,<sup>3,2</sup> Klaus-Robert Müller,<sup>1,2</sup> and O. Anatole von Lilienfeld<sup>4,2,\*</sup>

<sup>1</sup>*Machine Learning Group, Technical University of Berlin, Franklinstr 28/29, 10587 Berlin, Germany*

<sup>2</sup>*Institute of Pure and Applied Mathematics, University of California Los Angeles, Los Angeles, CA 90095, USA*

<sup>3</sup>*Fritz-Haber-Institut der Max-Planck-Gesellschaft, 14195 Berlin, Germany*

<sup>4</sup>*Argonne Leadership Computing Facility, Argonne National Laboratory, Argonne, Illinois 60439, USA*

(Dated: September 14, 2011)

We introduce a machine learning model to predict atomization energies of a diverse set of organic molecules, based on nuclear charges and atomic positions only. The problem of solving the molecular Schrödinger equation is mapped onto a non-linear statistical regression problem of reduced complexity. Regression models are trained on and compared to atomization energies computed with hybrid density-functional theory. Cross-validation over more than seven thousand small organic molecules yields a mean absolute error of  $\sim 10$  kcal/mol. Applicability is demonstrated for the prediction of molecular atomization potential energy curves.

Solving the Schrödinger equation (SE),  $H\Psi = E\Psi$ , for assemblies of atoms is a fundamental problem in quantum mechanics. Alas, solutions that are exact up to numerical precision are intractable for all but the smallest systems with very few atoms. Hierarchies of approximations have evolved, usually trading accuracy for computational efficiency [1]. Conventionally, the external potential, defined by a set of nuclear charges  $\{Z_I\}$  and atomic positions  $\{\mathbf{R}_I\}$ , uniquely determines the Hamiltonian  $H$  of any system, and thereby the potential energy by optimizing  $\Psi$ , [2]  $H(\{Z_I, \mathbf{R}_I\}) \xrightarrow{\Psi} E$ . For a diverse set of organic molecules, we show that one can use machine learning (ML) instead,  $\{Z_I, \mathbf{R}_I\} \xrightarrow{\text{ML}} E$ . Thus, we circumvent the task of explicitly solving the SE by training once a machine on a finite subset of known solutions. Since many interesting questions in physics require to repeatedly solve the SE, the highly competitive performance of our ML approach may pave the way to large scale exploration of molecular energies in chemical compound space [3, 4].

ML techniques have recently been used with success to map the problem of solving complex physical differential equations to statistical models. Successful attempts include solving Fokker-Planck stochastic differential equations [5], parameterizing interatomic force fields for fixed chemical composition [6, 7], and the discovery of novel ternary oxides for batteries [8]. Motivated by these, and other related efforts [9–12], we develop a non-linear regression ML model for computing molecular atomization energies in chemical compound space [3]. Our model is based on a measure of distance in compound space that accounts for both stoichiometry and configurational variation. After training, energies are predicted for new (out-of-sample) molecular systems, differing in composition and geometry, at negligible computational cost, *i.e.* milli seconds instead of hours on a conventional CPU. While the model is trained and tested using atomization energies calculated at the hybrid density-functional the-

ory (DFT) level [2, 13, 14], any other training set or level of theory could be used as a starting point for subsequent ML training. Cross-validation on 7165 molecules yields a mean absolute error of 9.9 kcal/mol, which is an order of magnitude more accurate than counting bonds or semi-empirical quantum chemistry.

We use the GDB data base, a library of nearly one billion organic molecules that are stable and synthetically accessible according to organic chemistry rules [15]. While potentially applicable to any stoichiometry, as a proof of principle we restrict ourselves to small organic molecules. Specifically, we define a controlled test-bed consisting of all 7165 organic molecules from the GDB data base with up to seven “heavy” atoms that contain C, N, O, or S, being saturated with hydrogen atoms. Atomization energies range from -800 to -2000 kcal/mol. Structural features include a rich variety of chemistry such as double, and triple-bonds; (hetero)cycles, carboxy, cyanide, amide, alcohol, and epoxy-groups. For each of the many stoichiometries, many constitutional (differing chemical bonds) but no conformational isomers are part of this data base. Based on the string representation of molecules in the data base, we generated Cartesian geometries with OpenBabel [16]. Thereafter, the PBE0 approximation to hybrid DFT [17, 18] in converged numerical basis, as implemented in the FHI-aims code [19] (tight settings/tier2 basis set), was used to compute reference atomization energies for training. Our choice of the PBE0 hybrid functional is motivated by small errors ( $< 5$  kcal/mol) for thermo-chemistry data that includes molecular atomization energies [20].

One of the most important ingredients for ML is the choice of an appropriate data representation that reflects prior knowledge of the application domain, *i.e.* a model of the underlying physics. A variety of such “descriptors” are used by statistical methods for chem- and bio-informatics applications [21–23]. For modeling atomization energies, we use the same molecular information that enters the Hamiltonian for an electronic structure calcu-

lation, namely the set of Cartesian coordinates,  $\{\mathbf{R}_I\}$ , and nuclear charges,  $\{Z_I\}$ . Our representation consists of atomic energies, and the inter-nuclear Coulomb repulsion operator. Specifically, we represent *any* molecule by a ‘‘Coulomb’’ matrix  $\mathbf{M}$ ,

$$M_{IJ} = \begin{cases} 0.5Z_I^{2.4} & \forall I = J, \\ \frac{Z_I Z_J}{|\mathbf{R}_I - \mathbf{R}_J|} & \forall I \neq J. \end{cases} \quad (1)$$

Here, off-diagonal elements correspond to the Coulomb repulsion between atoms  $I$  and  $J$ , while diagonal elements encode a polynomial fit of atomic energies to nuclear charge.

Using ML we attempt to construct a non-linear map between molecular characteristics and atomization energies. This requires a measure of molecular (dis)similarity that is invariant with respect to translations, rotations, and the index ordering of atoms. To this end, we measure the distance between two molecules by the Euclidean norm of their diagonalized Coulomb matrices:  $d(\mathbf{M}, \mathbf{M}') = d(\boldsymbol{\epsilon}, \boldsymbol{\epsilon}') = \sqrt{\sum_I |\epsilon_I - \epsilon'_I|^2}$ , where  $\boldsymbol{\epsilon}$  are the eigenvalues of  $\mathbf{M}$  in order of decreasing absolute value. For matrices that differ in dimensionality,  $\boldsymbol{\epsilon}$  of the smaller system is extended by zeros. Note that by representing chemical compound space in this way, (i) any system is uniquely encoded because stoichiometry as well as atomic configuration are explicitly accounted for, (ii) symmetrically equivalent atoms contribute equally, (iii) the diagonalized  $\mathbf{M}$  is invariant with respect to atomic permutations, translations, and rotations, and (iv) the distance is continuous with respect to small variations in inter-atomic distances or nuclear charges. As discussed in Ref. [24], these are all crucial criteria for representing atomistic systems within statistical models.

In Fig. 1, relative atomization energies, as a function of  $d(\mathbf{M}, \mathbf{M}')$ , and a histogram of distances are shown for all pairs of molecules in our data set. The inset exemplifies the distances between three molecular species, pyrrol, thiophene, and ethanol: Within our measure of similarity the nitrogen containing aromatic heterocycle pyrrol is  $\sim 10$  times farther away from its sulfur containing analogue, thiophene, than from ethanol. This is due to the large difference in nuclear charges between atoms from different rows in the periodic table.

Within our ML model [25–27], the energy of a molecule  $\mathbf{M}$  is a sum over weighted Gaussians,

$$E^{est}(\mathbf{M}) = \sum_{i=1}^N \alpha_i \exp\left[-\frac{1}{2\sigma^2} d(\mathbf{M}, \mathbf{M}_i)^2\right], \quad (2)$$

where  $i$  runs over all molecules  $\mathbf{M}_i$  in the training set. Regression coefficients  $\{\alpha_i\}$  and length-scale parameter  $\sigma$  are obtained from training on  $\{\mathbf{M}_i, E_i^{ref}\}$ . Note that each training molecule  $i$  contributes to the energy not only according to its distance, but also according to its specific weight  $\alpha_i$ . The  $\{E_i^{ref}\}$  were computed at PBE0 DFT level of theory.

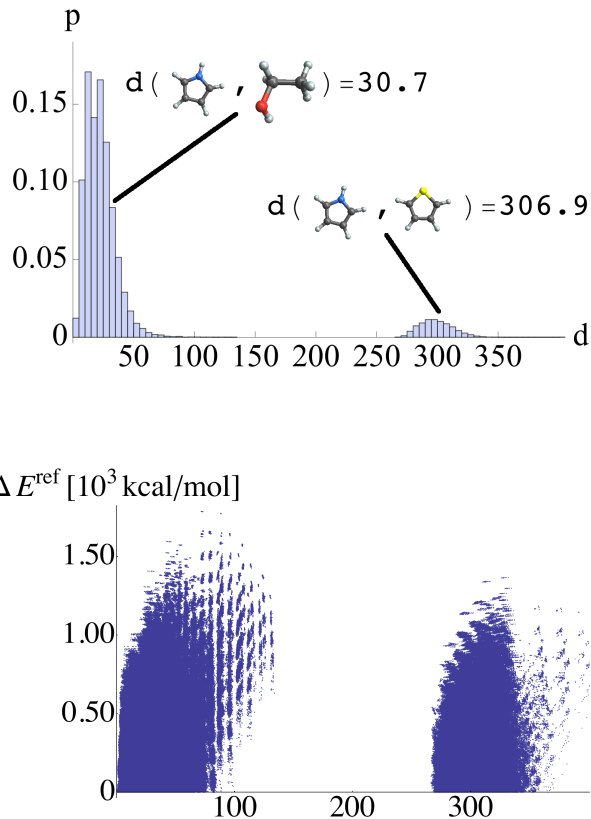


FIG. 1: (Color online) Top: Distribution of distances,  $d(\mathbf{M}, \mathbf{M}')$ , for all molecular pairs occurring in the first 7165 small organic molecules from the GDB data base [15]. The inset exemplifies two distances, pyrrol/ethanol and pyrrol/thiophene (N: Blue, O: Red, S: Yellow, C: Black, H: White). Bottom: Absolute differences in atomization energies between  $\mathbf{M}$  and  $\mathbf{M}'$  as a function of  $d(\mathbf{M}, \mathbf{M}')$ .

To determine  $\{\alpha_i\}$ , we used kernel ridge regression [26]. This regularized model limits the norm of regression coefficients,  $\{\alpha_i\}$ , thereby ensuring the transferability of the model to new compounds. For given length-scale  $\sigma$  and regularization parameter  $\lambda$ , the explicit solution to the minimization problem,

$$\min_{\boldsymbol{\alpha}} \sum_i (E^{est}(\mathbf{M}_i) - E_i^{ref})^2 + \lambda \sum_i \alpha_i^2 \quad (3)$$

is given by  $\boldsymbol{\alpha} = (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{E}^{ref}$ ,  $K_{ij} = \exp[-d(\mathbf{M}_i, \mathbf{M}_j)^2 / (2\sigma^2)]$  being the kernel matrix of all training molecules, and  $\mathbf{I}$  denoting the identity matrix.

We used stratified [28] five-fold cross-validation [26, 27] for model selection and to estimate performance. Parameters  $\lambda$  and  $\sigma$  were determined in an inner loop of five-fold cross-validation using a logarithmically scaling grid. This procedure is routinely applied in machine learning and statistics to avoid over-fitting and overly optimistic error estimates.

The dependence of the cross-validated ML performance on the number of molecules in training set,  $N$ , is illustrated in Fig. 2 (top). When increasing  $N$  from 500 to 7000, the mean absolute error (MAE) falls off from more than 17 kcal/mol to less than 10 kcal/mol. Furthermore, the width  $\sigma$  of the Gaussian kernel decreases from 460 to 25 on the distance scale of Fig. 1. A small  $\sigma$  emphasizes compound pairs for which the distance is small, whereas a larger  $\sigma$  allows for contributions from distant pairs. This is to be expected for increased number of training molecules. Because of the discrete nature of chemical space (nuclear charges can only assume integer values), however, we do not expect continuous coverage for  $N \rightarrow \infty$ , implying that  $\sigma$  will converge to a small but finite value. The regularization hyperparameter  $\lambda$  remains small throughout, consistent with the fact that we model noise-free numerical solutions of the approximated Schrödinger equation. An asymptotic fit of the form  $\sim 1/\sqrt{N}$ , based on statistical theory [26, 29] suggests that the MAE can be lowered to  $\sim 7.6$  kcal/mol for  $N \rightarrow \infty$ . It is remarkable that already for the here presented, relatively small training set sizes, ML achieves errors of roughly one percent on the relevant scale of energies, clearly outperforming bond counting or semi-empirical quantum chemistry methods. The cross-validated performance for a training set size of  $N = 1000$  is displayed in Fig. 2 (bottom). There is good correlation with the DFT data. For comparison, corresponding correlations are shown for bond counting [30], and semi-empirical quantum chemistry (PM6 [31]) computed with MOPAC [32]. While the latter two methods exhibit a systematic shift in slope, the inset highlights that the ML correlation accurately reproduces clustering, and slope of one.

In order to assess transferability and applicability of our model to chemical compound space, we use a ML model trained on  $N = 1000$  molecules (model 1k). The training set of model 1k contains all small molecules with 3 to 5 heavy atoms, and a randomized stratified selection of larger compounds covering the entire energy range. The thousand Coulomb matrices corresponding to the OpenBabel configurations were included as well as four additional Coulomb matrices per molecule. These additional matrices were scaled in order to represent the repulsive wall, the dissociative limit, and the energy minimum at  $f = 1$  [33, 34]. All predictions are made for molecules that were *not* used during training of the model.

For testing the transferability, we applied the 1k model to the remaining 6k molecules. The calculations yield errors that hardly change from the estimated performance in the training with a MAE of 15.2 kcal/mol. For the selected molecular subset of the seven thousand smallest molecules in the GDB database [15], we therefore conclude that training on 15% of the molecules permits predictions of atomization energies for the remaining 85%




FIG. 2: (Color online) Top: Cross-validated ML errors as a function of number of molecules in training set,  $N$ . Bottom: For  $N = 1000$ , correlation of DFT-PBE0 [17, 18] results ( $E^{ref}$ ) with ML (cross validated) based estimates ( $E^{est}$ ) of atomization energies. Correlations for bond counting [30] and semi-empirical quantum chemistry (PM6[31]) are also shown. Corresponding RMSE (root mean square error)/MAE (mean absolute error) for bond counting, PM6, and ML are 75.0/71.0, 75.1/73.1, 30.1/14.9 kcal/mol, respectively.

with an accuracy of roughly 15 kcal/mol.

For probing the applicability, we investigated whether the 1k model can also be useful beyond the equilibrium geometries. Specifically, we calculated the functional dependence of atomization energies on scaling Cartesian geometries by a factor,  $f$ . From the 6k molecules (not used for training) we picked four which exhibit chemical diversity. Specifically, these molecules contain single bonds and branching only ( $C_7H_{16}$ ), a double bond ( $C_6H_{12}$ ), triple bonds including nitrogen ( $C_6NH_5$ ), and a sulfur containing cycle with a hydroxy group ( $C_4SH_3OH$ ). The resulting ML atomization energy curves (Fig. 3) correctly distinguish between the molecules, closely reproduce the DFT energy at  $f = 1$ , and appear continuous and differentiable throughout relevant bonding distances. For comparison, corresponding Morse potential curves are

also displayed. Their well-depth and exponential factor were explicitly fitted to the molecular DFT minimum, as well as repulsive wall and dissociative limit [34]. Albeit slightly overestimating equilibrium distance and well depth for  $C_4SH_3OH$  and  $C_6NH_5$ , the ML model is in overall good agreement with the Morse potential curves. One can speculate if the better performance of the ML model for the larger molecules is due to the fact that in the total set larger molecules are more frequent than smaller molecules. Again, we stress the contrast that while the Morse potential curves were explicitly fitted for these four molecules, the ML model was obtained for a training set based on one thousand *other* molecules.




FIG. 3: (Color online) Energy of atomization curves of four molecules containing single bonds and branching only ( $C_7H_{16}$ ), a double bond ( $C_6H_{12}$ ), triple bonds including nitrogen ( $C_6NH_5$ ), and a sulfur containing cycle with a hydroxy group ( $C_4SH_3OH$ ). (bottom to top in insets; black: Carbon; blue: Nitrogen; yellow: Sulfur; red: Oxygen; white: Hydrogen) (DFT-PBE0, Morse potential, and ML model 1k).

We have developed a ML approach for modeling atomization energies in the chemical compound space of small organic molecules. For larger training sets, accuracies have been achieved that are competitive with mean-field electronic structure theory, at a fraction of the computational cost. We find good performance when making predictions for new molecules (transferability) and when predicting atomization energies beyond the equilibrium geometry. Our representation of molecules as Coulomb matrices is inspired by the nuclear repulsion term in the molecular Hamiltonian, and free atom energies. Future extensions of our approach might permit rational compound design applications [35–37] as well as geometrical relaxations, chemical reactions, or molecular dynamics in various ensembles [38]. Finally, our results suggest that the Coulomb matrix, or improvements thereof, could be of interest as a descriptor beyond the presented application.

We are thankful for helpful discussions with K. Burke, M. Cuendet, K. Hansen, J-L. Reymond, B. C. Rinderspacher, M. Rozgic, M. Scheffler, A. P. Thompson, M. E. Tuckerman, S. Varma. All authors acknowledge support from the long program “Navigating Chemical Compound Space for Materials and Bio Design”, IPAM, UCLA. This research used resources of the Argonne Leadership Computing Facility at Argonne National Laboratory, which is supported by the Office of Science of the U.S. DOE under contract DE-AC02-06CH11357. M.R. and K.-R. M. acknowledge partial support by DFG (MU 987/4-2) and EU (PASCAL2).

\* Electronic address: anatole@alcf.anl.gov

- [1] P. Ragué von Schleyer, N. Allinger, T. Clark, J. Gasteiger, P. Kollman, H.F. Schaefer III., and P. Schreiner, eds., *Encyclopedia of Computational Chemistry* (John Wiley and Sons, 1998).
- [2] P. Hohenberg and W. Kohn, *Phys. Rev.* **136**, B864 (1964).
- [3] P. Kirkpatrick and C. Ellis, *Nature* **432**, 823 (2004).
- [4] O. A. von Lilienfeld and M. E. Tuckerman, *J. Chem. Phys.* **125**, 154104 (2006).
- [5] R. R. Coifman, I. G. Kevrekidis, S. Lafon, M. Maggioni, and B. Nadler, *Multiscale Mod. Sim.* **7**, 842 (2008).
- [6] A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, *Phys. Rev. Lett.* **104**, 136403 (2010).
- [7] C. M. Handley and P. L. A. Popelier, *J. Chem. Theory Comput.* **5**, 1474 (2009).
- [8] G. Hautier, C. C. Fischer, A. Jain, T. Mueller, and G. Ceder, *Chem. Mater.* **22**, 3762 (2010).
- [9] A. Brown, B. J. Braams, K. Christoffel, Z. Jin, and J. M. Bowman, *J. Chem. Phys.* **119**, 8790 (2003).
- [10] S. Lorenz, A. Gross, and M. Scheffler, *Chem. Phys. Lett.* **395**, 210 (2004).
- [11] J. Behler and M. Parrinello, *Phys. Rev. Lett.* **98**, 146401 (2007).
- [12] J. Behler, R. Martonak, D. Donadio, and M. Parrinello, *Phys. Rev. Lett.* **100**, 185501 (2008).
- [13] W. Kohn and L. J. Sham, *Phys. Rev.* **140**, A1133 (1965).
- [14] A. D. Becke, *J. Chem. Phys.* **98**, 5648 (1993).
- [15] L. C. Blum and J.-L. Reymond, *J. Am. Chem. Soc.* **131**, 8732 (2009).
- [16] R. Guha, M. T. Howard, G. R. Hutchison, P. Murray-Rust, H. Rzepa, C. Steinbeck, J. K. Wegner, and E. Wilhagen, *J. Chem. Inf. Model.* **46**, 991 (2006).
- [17] J. P. Perdew, M. Ernzerhof, and K. Burke, *J. Chem. Phys.* **105**, 9982 (1996).
- [18] M. Ernzerhof and G. E. Scuseria, *J. Chem. Phys.* **110**, 5029 (1999).
- [19] V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler, *Comput. Phys. Comm.* **180**, 2175 (2009).
- [20] B. J. Lynch and D. G. Truhlar, *J. Phys. Chem. A* **107**, 3898 (2003).
- [21] G. Schneider, *Nature Reviews* **9**, 273 (2010).
- [22] O. Ivanciuc, *J. Chem. Inf. Comp. Sci.* **40**, 1412 (2000).
- [23] J.-L. Faulon, D. P. Visco, Jr., and R. S. Pophale, J.

- Chem. Inf. Comp. Sci. **43**, 707 (2003).
- [24] J. Behler, J. Chem. Phys. **134**, 074106 (2011).
- [25] B. Schölkopf and A. J. Smola, *Learning with Kernels* (MIT Press, Cambridge, 2002).
- [26] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning. Data Mining, Inference, and Prediction* (Springer, New York, 2009), 2nd ed.
- [27] K.-R. Müller, S. Mika, G. Rätsch, K. Tsuda, and B. Schölkopf, IEEE Transactions on Neural Networks **12**, 181 (2001).
- [28] Stratification was done by sorting energies of the training data, grouping corresponding sorted compounds into blocks of five compounds each, and, for each such block, randomly assigning one compound to one cross-validation fold. This procedure ensures that each fold covers the whole energy range.
- [29] K. R. Müller, M. Finke, N. Murata, K. Schulten, and S. Amari, Neural Comp. **8**, 1085 (1996).
- [30] As summarized on [http://www.wiredchemist.com/chemistry/data/bond\\_energies\\_lengths.html](http://www.wiredchemist.com/chemistry/data/bond_energies_lengths.html): Huheey, pps. A-21 to A-34; T.L. Cottrell, "The Strengths of Chemical Bonds," 2nd ed., Butterworths, London, 1958; B. deB. Darwent, "National Standard Reference Data Series," National Bureau of Standards, No. 31, Washington, DC, 1970; S.W. Benson, J. Chem. Educ., 42, 502 (1965).
- [31] J. J. P. Stewart, J. Mol. Modeling **13**, 1173 (2007).
- [32] MOPAC2009, James J. P. Stewart, Stewart Computational Chemistry, Colorado Springs, CO, USA, [HTTP://OpenMOPAC.net](http://OpenMOPAC.net) (2008).
- [33] T.-C. Lim, Mol. Phys. **108**, 1589 (2010).
- [34] For the repulsive wall, atomization energies for Coulomb matrices were set to zero at typical roots for covalent bonds,  $f = 2/3$ . [33] For the minimum atomization energies, the finite difference derivative,  $dE/df = 1$ , was set to zero at  $f = 0$ , and using  $df = 0.005$ . For the dissociative tail we assume zero atomization energies at  $f = 3$ .
- [35] O. A. von Lilienfeld, R. Lins, and U. Rothlisberger, Phys. Rev. Lett. **95**, 153002 (2005).
- [36] O. A. von Lilienfeld, J. Chem. Phys. **131**, 164102 (2009).
- [37] D. Sheppard, G. Henkelman, and O. A. von Lilienfeld, J. Chem. Phys. **133**, 084104 (2010).
- [38] M. E. Tuckerman, *Statistical mechanics: Theory and molecular simulation* (Oxford University Press, 2010).