

This document is confidential and is proprietary to the American Chemical Society and its authors. Do not copy or disclose without written permission. If you have received this item in error, notify the sender and delete all copies.

Genetic Optimization of Training Sets for Improved Machine Learning Models of Molecular Properties

Journal:	<i>The Journal of Physical Chemistry Letters</i>
Manuscript ID	jz-2017-00038v.R1
Manuscript Type:	Letter
Date Submitted by the Author:	02-Mar-2017
Complete List of Authors:	Browning, Nicholas; Ecole Polytechnique Federale de Lausanne, Laboratory of Computational Chemistry and Biochemistry Ramakrishnan, Raghunathan; University of Basel, Institute of Physical Chemistry von Lilienfeld, O. Anatole; University of Basel, Institute of Physical Chemistry Rothlisberger, Ursula; Ecole Polytechnique Federale Lausanne, ;

SCHOLARONE™
Manuscripts

Genetic Optimization of Training Sets for Improved Machine Learning Models of Molecular Properties

Nicholas J. Browning,^{†,‡} Raghunathan Ramakrishnan,^{¶,‡,§} O. Anatole von Lilienfeld,^{*,¶,‡} and Ursula Röthlisberger^{*,†,‡}

Laboratory of Computational Chemistry and Biochemistry, Institute of Chemical Sciences and Engineering. Swiss Federal Institute of Technology (EPFL) CH-1015 Lausanne, Switzerland, National Center for Computational Design and Discovery of Novel Materials (MARVEL), and Institute of Physical Chemistry, Department of Chemistry, University of Basel, Klingelbergstrasse 80, CH-4056 Basel, Switzerland

E-mail: anatole.vonlilienfeld@unibas.ch; ursula.roethlisberger@epfl.ch

*To whom correspondence should be addressed

[†] Laboratory of Computational Chemistry and Biochemistry, Institute of Chemical Sciences and Engineering. Swiss Federal Institute of Technology (EPFL) CH-1015 Lausanne, Switzerland

[‡] National Center for Computational Design and Discovery of Novel Materials (MARVEL)

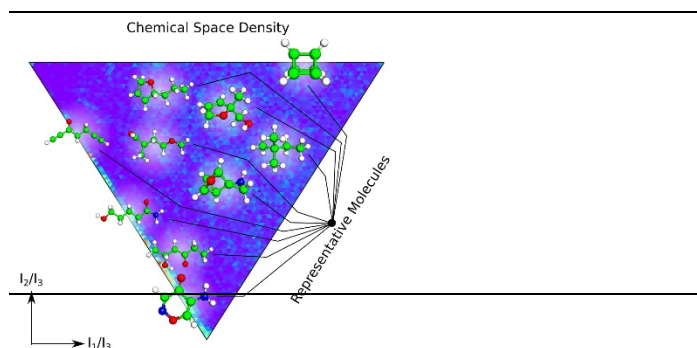
[¶] Institute of Physical Chemistry, Department of Chemistry, University of Basel, Klingelbergstrasse 80, CH-4056 Basel, Switzerland

[§] Current address: Tata Institute of Fundamental Research, Centre for Interdisciplinary Sciences, 21 Brundavan Colony, Narsingi, Hyderabad 500075, India

Abstract

The training of molecular models of quantum mechanical properties based on statistical machine learning requires large datasets which exemplify the map from chemical structure to molecular property. Intelligent *a priori* selection of training examples is often difficult or impossible to achieve as prior knowledge may be unavailable. Ordinarily representative selection of training molecules from such datasets is achieved through random sampling. We use genetic algorithms for the optimization of training set composition consisting of tens of thousands of small organic molecules. The resulting machine learning models are considerably more accurate: in the limit of small training sets mean absolute errors for out-of-sample predictions are reduced by up to $\sim 75\%$. We discuss and present optimized training sets consisting of 10 molecular classes for all molecular properties studied. We show that these classes can be used to design improved training sets for the generation of machine learning models of the same properties in similar but unrelated molecular sets.

Graphical TOC Entry



Experimentally accurate solutions to the time-independent non-relativistic electronic Schrödinger equation (SE) $H(\{Z_I, \mathbf{R}_I\}, N_e)\Psi = E \Psi$ for N_e electrons and a collection of atoms involve numerically challenging calculations¹. This limits routine electronic structure elucidation and accurate high-throughput screening. Previous works^{2–4} have found that the task of repetitiously solving the SE can be mapped onto a computationally efficient, data-driven supervised machine-learning (ML) problem instead. In these models, expectation values of quantum-mechanical (QM) operators are inferred in the subset of chemical space spanned by a set of reference molecular graphs, enabling a speedup of several orders of magnitude⁵ for predicting relevant molecular properties such as enthalpies, polarizabilities and electronic excitations^{4,6,7}. Here, QM reference calculations provide training examples $\{(x_\alpha, y_\alpha)\}_{\alpha=1}^N$, where x_α are molecular structures and y_α the expectation values for a chemical property, for interpolation in a $4N$ dimensional space, aka. chemical space. After training, accurate property predictions for new as of yet unseen molecules can be obtained at the base cost of the underlying ML model, provided that the new query molecule lies close to the space spanned by the reference data. A key issue in the validation of ML models is then the selection of appropriate data to use for training. Here, we will tackle this problem in the context of ML models of quantum chemistry–based estimates of molecular properties. Training examples are typically chosen from a uniform random distribution, thus there is no guarantee that the selected data will produce an optimal model. In this work we study the effect of optimizing the composition of the set of training examples used for learning molecular properties, by maximizing the predictive power of the underlying ML model for a given training set size. Previous work has found that evolutionary algorithms may be used to reduce training set sizes for artificial neural networks^{8,9} with improved model error. Here, however, we investigate the limits and composition of tailored training sets for learning a variety of molecular properties. There are numerous alternative strategies to design ML models with reduced property prediction errors. In previous work k-fold cross-validation has been used¹⁰ to reduce the predictive error of the ML model by optimizing the hyperparam-

eters σ and λ . The molecular representation could also be exchanged without any loss of generality. For this study we chose to rely on the sorted Coulomb matrix representation – a well established and generic representation which meets the crucial uniqueness criterion¹¹. We note that more advantageous descriptors, such as the Bag-of-Bonds (BoB)¹², or Bonds-and-Angles Machine Learning (BAML) models¹³ could have been used just as well. Finally, instead of focussing on properties directly, Δ -learning³ can be used to focus on the difference between a baseline and higher level of theory. In this study we include this model for learning the differences between PM7 and B3LYP atomization enthalpies.

This work is based on a quantum chemistry database published in 2014¹⁴, which contains relaxed geometries and 13 molecular properties computed at the DFT/B3LYP/6-31(2df,p) level of theory for 133'885 small organic molecules of up to 9 heavy atoms, extracted from the GDB-17 list of 166B SMILES strings¹⁵, and will herein be referred to as the GDB9 database. In this work we use both this and a smaller subset, denoted as the GDB8 database, containing molecules of up to 8 heavy atoms, resulting in 21800 molecules. For the GDB8 subset, the following properties are considered for training set composition optimization: enthalpy H and free energy G of atomization; heat capacity C_v ; isotropic molecular polarizability α ; electronic spatial extent (R^2); harmonic zero-point vibrational energy ZPVE; energy of the highest occupied E_{HOMO} and lowest unoccupied E_{LUMO} molecular orbitals; HOMO-LUMO gap ΔE and dipole moment μ . In addition, the modeling of electronic spectra of GDB8 has been studied in⁷. Further details and compositional analysis of the GDB-X databases can be found in¹⁶.

Similarly to the Hamiltonian used for electronic structure calculation, molecular information in the ML model has Cartesian coordinate and nuclear charge information inherently encoded through sorted Coulomb matrices^{10,17}, which are used to represent molecular structures for training and property prediction. The L^1 norm $d_{\alpha,\beta} = \|\mathbf{M}^\alpha - \mathbf{M}^\beta\|_{L^1} = \sum_{i,j} |\mathbf{M}_{ij}^\alpha - \mathbf{M}_{ij}^\beta|$ then serves as a metric measure of similarity between any two molecules α and β in the set of all sorted Coulomb matrices $\{\mathbf{M}\}$. Note that sorted Coulomb matrices

encode (except among enantiomers) the external potential of any given molecule uniquely, such that it is invariant with respect to molecular translation, rotation, and atom-indexing. The ML model attempts to construct a non-linear mapping between molecular characteristics and molecular properties. Here, we model the property P^{ML} of a new out-of-sample molecule \mathbf{M} as a linear combination of weighted Laplacian kernel functions $k(\mathbf{M}^a, \mathbf{M})$, located on each training instance a :

$$P^{\text{ML}}(\mathbf{M}) = \sum_{\alpha=1}^N c_{\alpha} k(\mathbf{M}^{\alpha}, \mathbf{M}) \quad (1)$$

where a runs over all molecules characterized by sorted Coulomb matrices \mathbf{M}^a in the training set of size N , and $k(\mathbf{M}^{\alpha}, \mathbf{M}^{\beta}) = \exp(-d_{\alpha,\beta}/\sigma)$ is the Laplacian kernel. For the Coulomb matrix descriptor, the combination of the Laplacian kernel and L^1 norm has been shown to yield an optimal combination of low computational cost and good predictive accuracy for models of atomization enthalpies¹⁰. The regression coefficient vector \mathbf{c} is obtained by training on $\{\mathbf{M}^a, \mathbf{P}^{\text{ref}}\}$. Note that each molecule a contributes to the property estimate not only according to its distance, but also according to its regression weight c_a . The global hyperparameter σ corresponds to the kernel-width. The regression coefficients are the solutions to the kernel ridge regression (KRR)¹⁸ minimization problem for a given kernel width σ and regularization parameter λ ,

$$\mathcal{L} = ||\mathbf{P}^{\text{ref}} - \mathbf{P}^{\text{ML}}||^2 + \lambda \mathbf{c}^T \mathbf{K} \mathbf{c} \quad (2)$$

where \mathcal{L} is the Lagrangian to be minimized with respect to the coefficient vector \mathbf{c} and \mathbf{P}^{ref} is the vector of all reference training data. The solution for the coefficient vector then reads $\mathbf{c} = (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{P}^{\text{ref}}$, with \mathbf{K} being the kernel matrix, and \mathbf{I} the identity matrix. Frequently, 5-fold cross validation is used to estimate optimal values of hyperparameters σ and λ , however, this is prohibitively expensive to perform during an optimization procedure. Instead, we employ the single-kernel method¹⁹, using constant values of $\sigma = 1000$ and $\lambda = 0$. The choice

of λ is warranted by the lack of outliers within the GDB8 database and thus any GA-optimized set will not be affected by the presence of these samples. We have additionally confirmed that $\sigma = 1000$ is approximately the minimum for all properties by scanning a logarithmic and fine grid. We also choose to fix σ values for each property as here we only wish to investigate the effect of training set selection, and thus it is preferable to keep the kernels fixed.

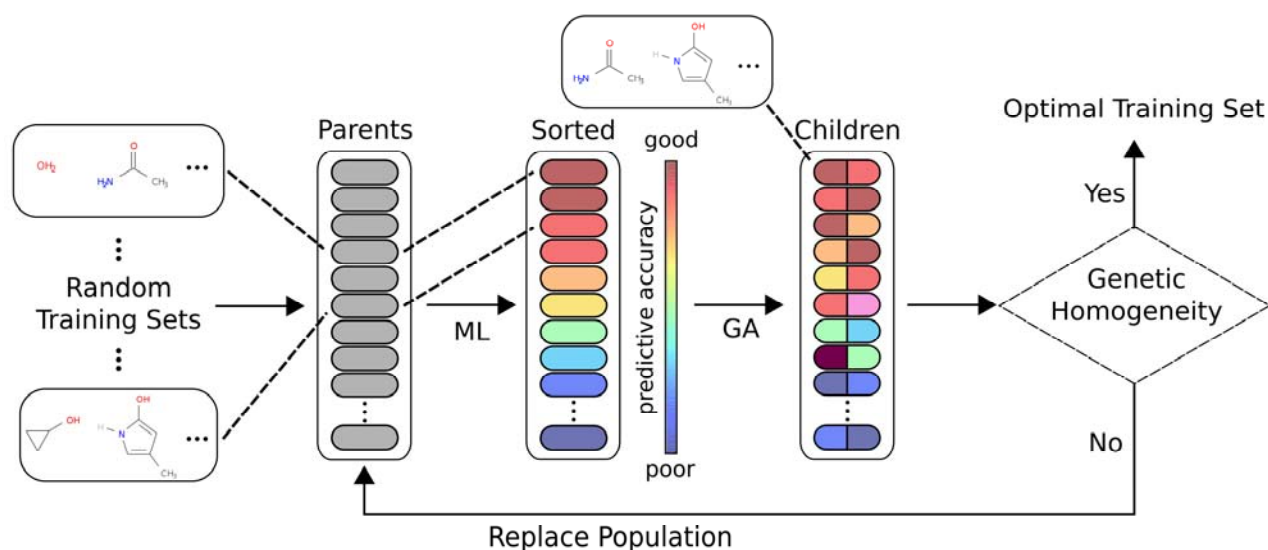


Figure 1: GA procedure to optimize molecular sets for training the ML model. A population of trial training sets containing an identical number of molecules are randomly sampled from the molecular database. This initial population of trial training sets serves as the first iteration of the algorithm, and is denoted as the parent population. An ML model is trained for each training set, and the mean out-of-sample prediction error is assigned to the training set as a fitness metric. The population then undergoes fitness-based selection and variation operators to create a child population. A portion of the worst children are replaced by the best parents. Finally the children are labelled as parents, and the algorithm repeats until there exists limited information difference between subsequent iterations.

Simple systematic enumeration to find training sets which minimize the out-of-sample predictive error of the ML model is a computationally demanding optimization problem. Finding 1000 optimal training molecules to train a machine which can best reproduce the properties of all molecules in the GDB8 validation database requires the training of $\binom{10800}{1000} = 1.21 \times 10^{1445}$ machines for a complete search. Clearly, complete optimization of training set composition in such a nearly-infinitely large space is impossible, thus an intelligent search

method is required to find near-optimal solutions. Here we employ a genetic algorithm (GA)^{20–22}, a biologically inspired meta-heuristic optimization technique, which has proved a successful optimization scheme in highly-dimensional and/or large spaces²³. The GA optimization process is pictorially summarized in figure 1. Bounded chemical compound space representations for the ML model, i.e training sets, are represented as a collection of unique molecular pointers,

$$\mathbf{X}_i = \{x_1, x_2, \dots, x_a, \dots, x_N\}.$$

\mathbf{X}_i is termed a position vector, where each element x_a uniquely maps to each Coulomb matrix \mathbf{M}_a , and N is the training set size. Each position vector \mathbf{X}_i is used to train the ML model and therefore facilitates the calculation of the mean absolute error (MAE) of predicting out-of-sample properties. Starting with a population of training sets with unique molecules sampled from a uniform distribution, individuals are stochastically selected using MAE as a fitness criterion. These initial training sets are then successively evolved by applying selection, variation and replacement. Selection determines which training sets should remain in the population to produce the training sets of the following generation. Variation results in new representations of chemical compound space and includes two operations. The first, termed crossover, mixes two training sets uniquely such that two new "child" training sets are produced which contain training molecules from both "parents". The second operator, termed mutation, randomly changes training molecules to introduce new information into the population. Finally, a portion of the worst children are replaced by the best parents to create a new parent population. After a number of iterations, genetic homogeneity is reached within the population and out-of-sample MAEs improve no further. These training sets are termed GA-optimized training sets and are considered to be near-optimal.

All reported relative mean absolute errors (RMAEs) refer to out-of-sample prediction of 10900 molecules from a machine trained on N in-sample molecules. Target accuracies

(for which $\text{RMAE} = 1$) for thermochemical properties and orbital energies are 1 kcal mol^{-1} . For ZPVE a target accuracy of 10 cm^{-1} was selected, comparable to the average accuracy of coupled cluster methods with converged basis sets²⁴ for predicting harmonic vibrational frequencies of small molecules. For isotropic polarizability and norm of dipole moment, target accuracies of $0.1 a_0^3$ and 0.1 D were used. Again, these values are within the predictive uncertainty using CC level of theory²⁵.




Figure 2: Relative mean absolute error (RMAE, target accuracy = 1 kcal mol^{-1}) as a function of training set size for enthalpy of atomization H using direct-learning P^{ML} (green), GA-optimized direct-learning P^{GA} (yellow), B3LYP-PM7 delta-learning Δ^{ML} (blue) and GA-optimized delta-learning Δ^{GA} (red). Dashed horizontal lines show training set sizes required to reach given accuracies. Left inset: scatter plot of B3LYP-PM7 reference energies of atomization Δ^{ref} and predicted values using both a randomly generated $N = 1000$ training set Δ^{ML} (blue), and the GA-optimized counterpart Δ^{GA} (red). Right inset: typical error distribution over B3LYP-PM7 reference atomization enthalpies for aforementioned random (blue) and GA-optimized (red) $N = 1000$ training sets.

Figure 2 displays learning curves of randomly generated and GA-optimized training sets for out-of-sample predictions of B3LYP level enthalpies of atomization, H , using both direct learning as well as the $\Delta H_{\text{PM7}}^{\text{B3LYP}}$ -ML model³. Due to KRR based ML errors decaying as inverse powers of training set size, we present RMAEs as a function of training set size on a log-log scale. Upon GA optimization we note a substantial lowering of the learning curves for both properties and all training set sizes. The combination of GA optimization with Δ -ML

model is the most promising: for machines trained on 3k molecules the GA optimization reduces the RMAE from 3.1 to 1.9. Corresponding scatterplots and error distribution plots are on display as insets in figure 2. They also indicate that the GA-optimized Δ -ML models approach the ideal model much faster than randomly sampled training sets. It is particularly encouraging to note that the relative gain in predictive power, obtained by using GA, appears to converge towards a constant rather than to vanish.

RMAEs for ML models with and without using GA are summarized in table 1 for various training set sizes and all the aforementioned properties. GA optimization of training set composition systematically improves RMAEs for all properties and training set sizes. To compare differences in training set size for a given target accuracy for both GA-optimized and randomly sampled training sets, please see supporting information (SI), table 1. Some properties experience greater improvement than others. The reduction in RMAE is most prevalent for smaller training set sizes, particularly of size $N = 10$. More specifically, properties related to chemical bonding, such as enthalpies and free energies of atomization as well as ZPVE improve by $\sim 75\%$. Other extensive properties, such as heat capacity or polarizability, improve by $\sim 50\%$. In contrast, intensive electronic properties experience much less improvement. Overall, the smallest RMAE reduction is found for the norm of the dipole moment. While percentage wise, seemingly small, the error reduction for ΔL_{PM7}^{B3LYP} is still very relevant due to its outstanding accuracy in absolute terms. This finding could possibly be also due to the comparatively more demanding target accuracies, the standard deviation of energetic properties in the data set is orders of magnitude larger than the standard deviation of electronic properties, such as the dipole moment. See figure 4 for comparison.

Overall, these results clearly amount to numerical evidence that the choice of how to select training set members can have a dramatic effect on the predictive power of the resulting ML model. It follows that substantially fewer training examples would be needed for the generation of ML models which reach the same accuracy as ML models trained on a much larger training set sampled at random. In order for this insight to be useful, however, one

Table 1: Randomized and GA-optimized out-of-sample relative mean absolute errors (RMAEs) for all properties. All target chemical accuracies are 1 kcal mol⁻¹, except for ZPVE, dipole moment and isotropic polarizability, which have target accuracies of 10cm⁻¹, 0.1D and 0.1a³ respectively. GA-optimized RMAEs are denoted by P^{GA} while randomly generated training set RMAEs are denoted as P^{ML} . Final row corresponds to out-of-sample RMAEs for enthalpy of atomization H using Δ -learning (aka ΔH_{PM7}^{B3LYP}), and bracketed the GA optimized counterpart, referred to as Δ^{ML} and Δ^{GA} respectively.

$P^{ML} (P^{GA})$	N													
	10		50		100		500		1k		2k		3k	
H	113.0	(31.6)	48.0	(18.3)	33.3	(14.3)	14.8	(7.5)	10.2	(5.8)	6.8	(4.5)	5.1	(3.9)
G	101.8	(28.8)	44.0	(17.7)	31.4	(14.1)	14.3	(7.5)	9.9	(5.6)	6.7	(4.3)	5.0	(3.9)
C_v	27.3	(14.5)	18.2	(9.4)	14.6	(7.8)	7.4	(4.0)	5.2	(2.9)	3.4	(2.3)	2.5	(2.0)
ZPVE	353.2	(83.2)	150.4	(38.4)	97.9	(28.0)	31.5	(14.0)	20.9	(10.5)	13.9	(7.0)	10.5	(3.5)
$\langle R^2 \rangle$	168.5	(92.2)	117.0	(44.2)	85.6	(33.2)	35.7	(19.1)	25.7	(15.5)	18.3	(12.7)	14.5	(11.6)
μ	11.3	(8.5)	10.3	(7.7)	9.9	(7.4)	8.4	(6.3)	7.5	(5.7)	6.2	(5.1)	5.2	(4.7)
α	40.8	(16.3)	23.1	(12.0)	18.5	(10.8)	11.8	(7.8)	9.6	(6.5)	7.2	(5.4)	5.8	(4.9)
E_{HOMO}	13.0	(9.0)	11.2	(8.1)	10.4	(7.3)	7.7	(5.2)	6.3	(4.5)	4.9	(3.8)	4.0	(3.5)
E_{LUMO}	22.3	(15.8)	18.8	(12.7)	17.0	(11.1)	11.9	(8.0)	9.7	(6.7)	7.4	(5.6)	5.9	(5.0)
gap	24.0	(17.8)	20.8	(15.0)	19.5	(13.5)	14.3	(9.8)	11.8	(8.1)	9.0	(6.8)	7.3	(6.2)
$\Delta^{ML}(\Delta^{GA})H$	6.6	(5.0)	6.0	(4.4)	5.7	(4.1)	4.6	(3.2)	4.1	(2.6)	3.4	(2.1)	3.1	(1.9)

would also require access to the solution of the selection optimization process in an *a priori* fashion, and not *a posteriori* as it is the case in this study.

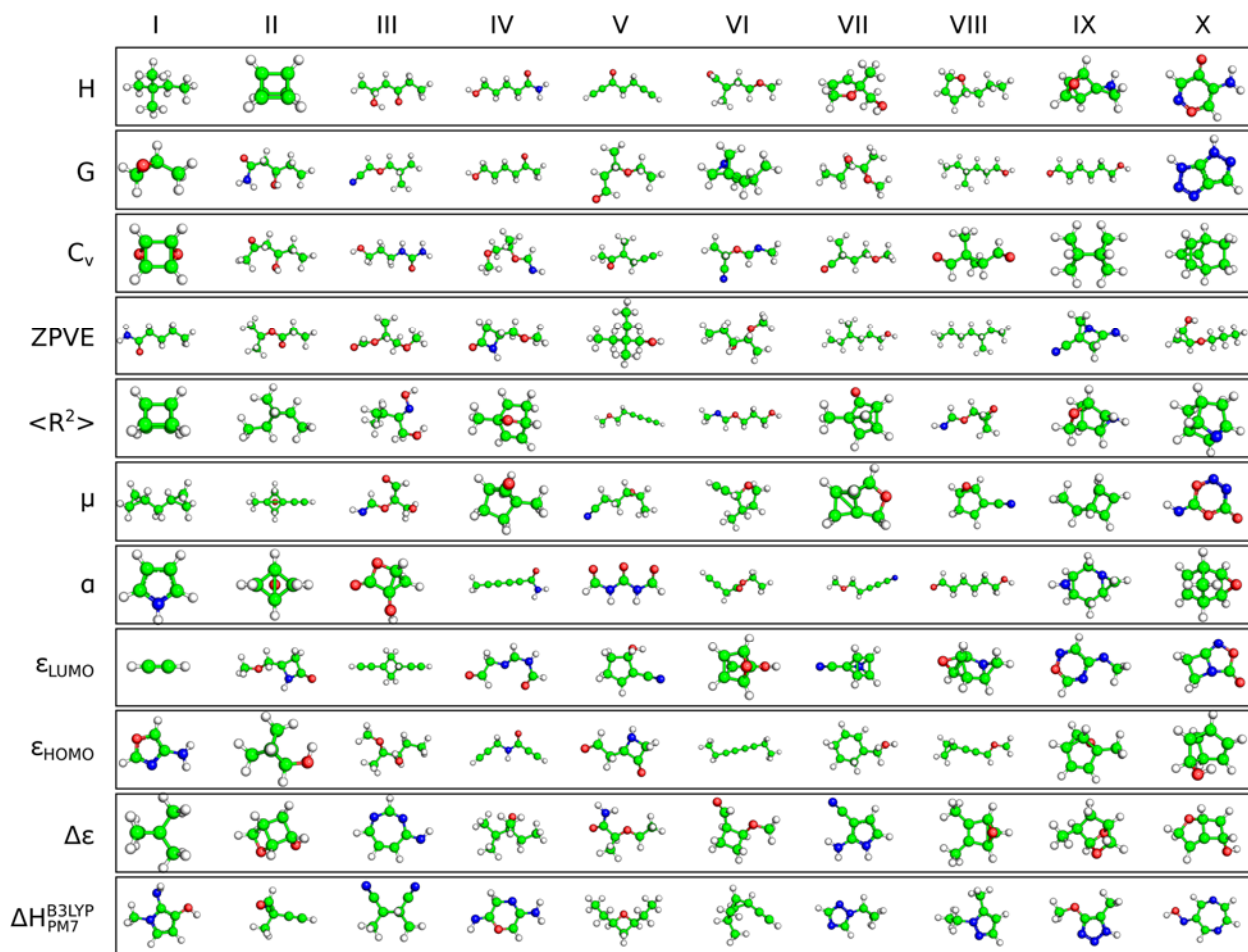


Figure 3: Molecular classes I-X for all properties listed in table 1. These classes correspond to the $N = 10$ column, where each molecule and its respective training set is collectively presumed to optimally represent the entire space of validation molecules. Each training set is sorted by its index in the GDB8 database, thus chemical weight approximately increases from left to right.

The additional GA optimization layer can be seen as providing the underlying machine learning model with the capacity to intelligently select its own data for optimal out-of-sample prediction. This in turn naturally means that the machine, through the kernel, L^1 distance metric and CM representation, is finding optimal maps from molecular structure and composition to property with respect to the diverse chemical space of the validation database. As such, it is interesting to inspect the outcome of this optimization. While this

is rapidly overwhelming for larger training sets, for $N = 10$ training sets the most essential chemistry encoded can easily be grasped. Figure 3 shows the optimized sets for all properties.

It is important to keep in mind that the molecules shown likely do not represent the global optimum, rather they represent one of many near-optimal combinations. We have obtained them by first resampling unique ensembles of $N = 10$ training sets many times, with a biased probability proportional to the L^1 distance from each molecule in the best GA-optimized training set for each property. All training sets sampled in this way produced significantly better RMAEs comparatively to random sampling. Thus, each GA-optimal molecule should be interpreted as being representative for very similar compounds which could be selected in optimized training sets, and are reported in figure 3. We therefore find it appropriate to relabel each index in the GA-optimal training sets as molecular classes I-X. Note though that, due to lack of an iterative optimisation procedure, ML models based on such biased training sets composed of these similar molecules yield slightly worse performance than ML models resulting from the underlying GA-optimized training sets.

From inspection it is obvious that the chemistry, represented by these classes, is very rich, including linear, planar, cage like, branched, and even strained structures. Variety in chemical composition is maintained through the occasional replacement of carbon units by functional groups containing oxygen or nitrogen. Furthermore, all hybridization states, sp^3 , sp^2 , and sp are represented. It can also be seen that many properties share very similar molecular classes, with some molecules being shared across optimal training sets for different properties. For example classes VII and VIII for ZPVE and G , or classes VIII and IV for α and G , respectively. We note that for energetic properties in particular, very stable and unstrained saturated molecules are selected as well as rather exotic unsaturated and strained systems, suggesting a bias towards the extremes of the distribution. Additionally, these exotic molecules tend to appear more so for more complex properties, such as HOMO-LUMO gap or ΔH_{PM7}^{B3LYP} . However, we do not find it obvious to further rationalize the specific selection of chemistries.

While the visual interpretation of the chemical classes discussed previously is not obvious, we can use the statistics from the GA optimization runs to systematically identify the effect of the bias. More specifically, to understand selection pressures we analyzed the regions of chemical space which produce optimal training sets for a given training set size. To this end, we have run 1000 identical GA optimizations for training sets of size $N = 10$ to obtain sufficient statistics to infer selection probabilities. In figure 4 we plot these results in terms of the principal moment of inertia density of molecular structures (left) for enthalpies of atomization, H , as well as in terms of the cumulative density plots (right) for properties H , μ , ΔE and ZPVE. Additionally on each cumulative density plot, we overlay the inferred selection probabilities of all molecules within the training database. It can be seen from the principal moment of inertia density plot, that while the training database disproportionately contains linear and oblate/prolate structures, GA optimization does not necessarily drive training set composition towards regions of high density within the GDB8 database. Indeed, many of the molecules which are frequently selected during training set composition optimization are from medium to low density regions and appear to rather homogeneously cover the entire plane.

In contrast, from the cumulative property density plots it can be seen that GA optimization preferentially selects regions of low property density (tails), while regions of high density (linear) are less likely to be sampled. We have investigated if introducing a deliberate bias into the training set sample, through Monte Carlo sampling of all the property distributions shown in figure 4 to specifically over-represent these regions, affords ML models with increased predictive power. Unfortunately, the introduction of such bias does not necessarily yield any improvement over random sampling (and indeed can even be worse) — thus this technique was not extended to other property distributions. While there are systematic trends upon GA optimization, there do not seem to be simple biasing rules purely based on properties, and the effective change in sampling chemical space has to be taken into account. We illustrate the effect of the GA-optimization on training set distribution

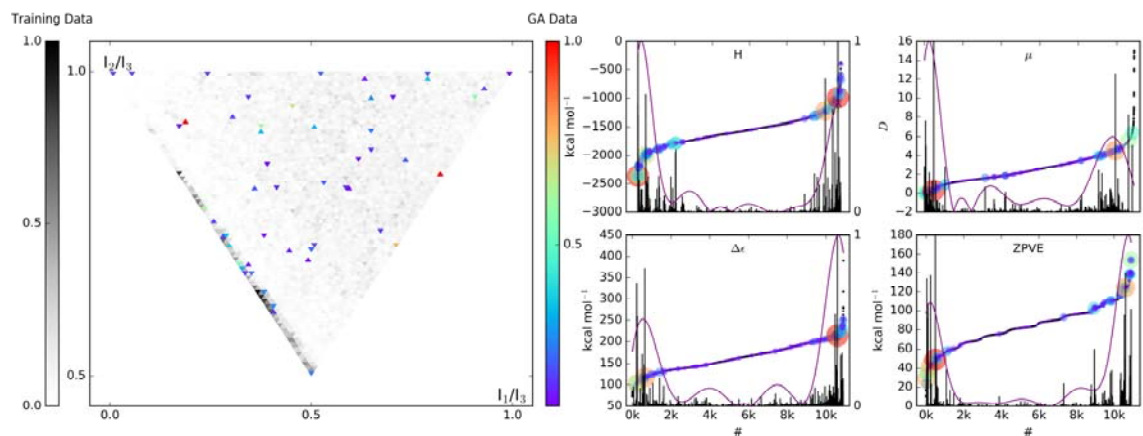


Figure 4: Left: Principal moment of inertia density plots for the training database (grayscale) and 1000 GA-optimized $N = 10$ training sets (colored) for enthalpy of atomization, H . Color corresponds to selection probability upon GA optimization. Right: training data cumulative density plots for the properties: H , μ , ΔE , and ZPVE. Normalised selection probabilities inferred from the ensemble of GA optimisations for $N = 10$ shown in black, with a corresponding 8th order polynomial fit shown in purple. Circle sizes and colors also correspond to selection probabilities.

in figure 5: we plot the difference in L^1 distance distribution between an average of 1000 randomly generated training sets of various sizes and their GA-optimized counterparts for the same properties. It can be seen that there is a small yet systematic GA-induced outward shift, i.e. to predominantly sample training molecules which are further apart from each other across chemical space. As with the reduction in RMAE upon GA-optimization these changes are not identical across properties, and appear to slowly converge to a constant ΔP per property. In view of this, we note that we have additionally tried to optimize training sets through maximal L^1 spread and thus generate training sets which are more chemically diverse, however, this yielded no improvement over random sampling.

In order to investigate the transferability of the above insights gained upon the molecular classes we have studied the effect of designing an artificially constructed training set, derived from the classes discovered through the use of GA on the GDB8 database. More specifically, we define "projection" rules to sample the larger and independent GDB9-8 database for which no GA has ever been performed. For each of the GDB8 molecules used to train the




Figure 5: GA-optimization tends to increase molecular distances: Normalised L^1 distance distribution differences (GA-optimized training set distribution minus randomly selected counterpart) for properties H (top left, green), μ (top right, purple), ΔE (bottom left, cyan) and ZPVE (bottom right, orange) as a function of training set size N .

ML model of H , on display in figure 3, we individually replace each hydrogen with the heavy substituents -OH, NH_2 and CH_3 . Figure 6 illustrates the procedure for the projection of GDB8 molecule H -V in figure 3 to GDB9-8. This procedure leads to hundreds of pseudo-GDB9-8 molecules, all having the same number of heavy elements as the other molecules present in GDB9-8, containing a total of $\sim 100\text{k}$ molecules. For each of these new, pseudo-GDB9-8 structures, we search the GDB9-8 training database (discussed in Computational Details) for the closest set of molecules which represent each of the ten classes. With the resulting molecules we trained new ML models for H which are applicable to GDB9-8.

While there are many pseudo-GDB9-8 structures found during the search which do exist (with some deviation in L^1 due to small geometrical differences), many of them do not have real GDB9-8 counterparts and in these cases a closest analogue is selected. When creating training sets for use in the GDB9-8 database, we do not weight these unfavourably and instead simply generate them with molecules uniformly sampled from each projected molecular class. On average, the resulting ML model affords a reduction of RMAE for out-of-sample predictions by more than 20% (decrease from 116.2 to 90.4) with respect to randomly generated training sets. Furthermore, upon use of the projected molecular classes, the error




Figure 6: Exemplification of applying the projection rule GDB8-GDB9-8 to the molecular class *H-V* (lower left, see figure 3) for -OH substitution. The box contains the pseudo-GDB9-8 molecules obtained by substitution of each alkyl hydrogen (shown in orange). For each of these pseudo molecules the entire GDB9-8 database is queried and the closest molecule is selected to represent the projection of this molecular class into the GDB9-8 database (lower right).

of the most extreme outlier in GDB9-8 is reduced from 275.4 to 116.1, while the error of the best training set reduces from 75.3 to 61.4. One should keep in mind that this still represents a substantial reduction: Due to the exponential scaling of chemical space with molecular size, GDB9-8 is roughly an order of magnitude larger than GDB8. As such, one should not expect that a simple projection from the 10 best molecules in GDB8 towards GDB9-8 results in the optimal set of 10 molecules out of GDB9-8, but rather in substantial sub-optimality. Finally, extending the projection rule beyond the GDB9-8 database would result in increasingly poorer accuracy due to a vanishing contribution of the GDB8 derived rules to ideal sampling criteria for training sets of larger molecules. While beyond the scope of this work, if the representation was constructed such that local environments were encoded, rather than entire molecules, we believe that analogous rules would scale well.

To investigate the generality of the relative improvement of tailored over uniformly randomly sampled training sets, we have calculated enthalpies of atomization for a distinct set of experimentally validated molecules, differing substantially from GDB. The set corresponds




Figure 7: Relative mean absolute errors (RMAEs, target accuracy = 1 kcal mol^{-1}) of randomly generated (green, shaded: 1-standard deviation) and GA-optimized (blue) GDB8-trained models, validated against the PubChem-8 database, as a function of training set size. Insets: Histograms of L^1 distances (top) and enthalpies of atomization (bottom) for GDB8 and PubChem-8 data base.

to approximately 10k relaxed molecules consisting of 8 heavy atoms (C, N, O, F) or less, and is known as the PubChem database²⁶ (see computational methods below for details), which we denote here within as the PubChem-8 database. To test transferability, we have applied the ML models trained on GA-optimized and randomly generated GDB8 training sets to predict Pubchem-8 enthalpies of atomization. While the prediction error is substantial, as one would expect for ML models trained on very distant training instances, the learning curves in Fig. 7 indicate systematic improvement with GDB8 training set and, more importantly, that the relative improvement is conserved, i.e. ML models based on randomly sampled training sets consistently result in worse performance than ML models based on GA-optimized training sets. Consequently, the effects of GA-optimization do not stem from overfitting the GDB-8 database. At this point one should note that this new database exhibits very different biases than those present in the GDBX series. Figure 7 insets show that the kernel densities estimates^{27,28} of both the distance matrix (left inset) between all molecules in the PubChem-8 database and the corresponding enthalpies of atomization (right inset) exhibit markedly different probability densities than those present within the GDB-8 database. In-

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

deed, the minimum L^1 distance to any molecule in the GDB8 database is 40 a.u., meaning that this PubChem database is disjoint from the GDB8 database. Additionally, the level of theory differs from the GDB8 data set in that basis sets used for the PubChem-8 DFT calculations are much smaller than in Ref.¹⁴, and hence the calculated enthalpies can differ by a non-trivial amount. As such, one should not expect ML models trained on GDB8 sets to be competitive with machines trained directly on PubChem-8 molecules.

Here, genetic algorithms (GA) have been employed to optimize the composition of molecular training sets that can be used for the training of machine learning models of molecular properties. Application of GA to a training set of a given size improves the performance of the ML model substantially, when compared to a model trained on randomly selected molecules coming from the GDB subspace of organic molecules that follows a distribution that is comprehensive with respect to chemical intuition. Conversely, for achieving the same accuracy, dramatically less training examples are necessary—provided that the user has the possibility to bias the selection of training examples prior to training. Ensembles of GA optimization procedures have shown that, while there is evidence of systematic bias towards low density regions of property distributions, attempting to construct a beneficially biased training instance distribution for improved ML models is not obvious. However, the design of improved training sets containing structures identified through the use of simple molecular projection rules applied to GA-optimized training sets seems to be possible. We have exemplified this for $\sim 100\text{k}$ molecules in GDB9-8 using constructed training set molecules obtained by GA for $\sim 20\text{k}$ molecules in GDB8.

It is worth noting that all optimized training sets reported are unlikely to correspond to global optima, but instead are rather near-optimal. The collection of training molecules producing such optimized RMAEs is dependent on the choice of GA parameters, but in particular upon population size, i.e the extent of training set sampling. Conversely the speed of optimization is directly linked to population and training set sizes, thus there is a trade-off between RMAE reduction and computation time. Nevertheless this error can

be further reduced using larger population sizes, or more training data. Tightening GA convergence criteria may also yield some improvement.

The distribution of optimal training instances depends strongly on the chosen fitting function, in our case the specific combination of Slater-type kernel function, L^1 norm, and Coulomb-matrix. Modification of the model (through representation/regressor/hyperparameter-selection) will likely influence the details of the selection pressure of particular molecules for a given training set size and property. Any other ML model is likely to lead to another optimal distribution. However, the overall insight that substantial model improvements can be obtained through GA optimization should be general since the selection bias present in training sets is independent of model details, thus our method is extensible to other property-learning techniques²⁹. Additionally, the choice of data set is arbitrary, and we could have equally investigated more diverse chemical spaces³⁰. The exact nature of the relationship between ML model specifics and selection bias remains to be elucidated in future work.

Finally, we note that while we have investigated how to remove the bias in a training set of a certain size for a given ML model, the inherent bias in the entire data set has not been explored—it is an open question, however, if that is even possible, due to the practically infinite number of possible molecular compositions and geometries. Future work will show if our conclusions also apply to ML models trained on other subsets of chemical space, e.g. corresponding to ZINC-molecules or the “representative set of molecules”³⁰.

Computational Methods

For the GA optimization of training sets in the GDB8 database, the entire database of 21800 molecules is partitioned into independent training and validation databases, each consisting of 10900 molecules. GA optimization procedures for all properties and training set sizes sample solely from the training database, while the validation set is used for calculating out-of-sample MAEs, used to guide the GA optimization procedure. We also define a third database in which all GDB8 molecules are removed from the GDB9 database, and name it

the GDB9-8 database. The GDB9 and GDB9-8 databases have been similarly partitioned into training and validation databases of 65k and 55k molecules each, respectively, which are used for projection rule development discussed later.

Genetic optimization is performed using a population of 2000 trial training sets for each property and training set size, for a maximum number of 500 iterations. Tournament selection without replacement³¹ is used with a tournament size of 2. Crossover is performed by partially-matched uniform crossover^{32–34} with a per-gene crossover probability of 0.5. Mutation is performed by partially-matched polynomial mutation³⁵ with a distribution index of 7 and a per-gene mutation probability of $\frac{0.5}{N}$ where N is the training set size. Finally, population replacement is achieved through fitness-proportionate elitism with an elitist ratio of 0.7. These parameters result from rigorous benchmarking, targeting maximal RMAE reduction, and typically result in convergence after ~300 iterations.

Finally, we have extracted approximately 95k molecules consisting of 8 heavy atoms (C, N, O, F) or less from the experimental PubChem database. We have performed geometry optimizations at the PM7 level with MOPAC³⁶ on all structures. From these, we have randomly extracted 10k molecules for further geometry optimization with DFT/B3LYP in an SVP basis using TURBOMOLE³⁷. Of these, approximately 56 molecules did not converge, while another 45 had converged to saddle points. We excluded these from further analysis. The resulting database, termed PubChem-8, contains 9899 converged geometries for which we have computed enthalpies of atomization.

Supporting Information

Table showing required training set sizes to reach GDB8 relative mean absolute errors (RMAEs) for both GA-optimized and randomly generated training sets, with both direct and Δ -learning models.

Acknowledgments

This research was supported by the NCCR MARVEL, funded by the Swiss National Science Foundation. OAvLacknowledges funding from the Swiss National Science Foundation (No. PP00P2_138932). Some calculations were performed at sciCORE (<http://scicore.unibas.ch/>) Scientific a Core Facility at the University of Basel and at the Swiss National Supercomputing Centre (CSCS, <http://www.cscs.ch/>). The authors would also like to thank P. F. Riley for discussions.

References

- (1) Cramer, C. J. *Essentials of Computational Chemistry: Theories and Models* (2nd Edition); John Wiley & Sons: West Sussex, England, 2004.
- (2) Rupp, M.; Tkatchenko, A.; Müller, K.-R.; von Lilienfeld, O. A. Fast and Accurate Modelling of Molecular Atomization Energies with Machine Learning. *Phys. Rev. Lett.* **2012**, *108*, 058301.
- (3) Ramakrishnan, R.; Dral, P. O.; Rupp, M.; von Lilienfeld, O. A. Big Data meets Quantum Chemistry Approximations: The Δ -Machine learning Approach. *J. Chem. Theory Comput.* **2015**, *11*, 2087–2096.
- (4) Montavon, G.; Rupp, M.; Gobre, V.; Vazquez-Mayagoitia, A.; Hansen, K.; Tkatchenko, A.; Müller, K.-R.; von Lilienfeld, O. A. Machine learning of molecular electronic properties in chemical compound space. *New J. Phys.* **2013**, *15*, 095003.
- (5) von Lilienfeld, O. A. First Principles View on Chemical Compound Space: Gaining Rigorous Atomistic Control of Molecular Properties. *Int. J. Quantum Chem.* **2013**, *113*, 1676–1689.
- (6) Hansen, K.; Biegler, F.; Ramakrishnan, R.; Pronobis, W.; von Lilienfeld, O. A.; Müller, K.-R.; Tkatchenko, A. Machine Learning Predictions of Molecular Properties:

- Accurate Many-Body Potentials and Nonlocality in Chemical Space. *J. Phys. Chem. Lett.* **2015**, 6, 2326–2331.
- (7) Ramakrishnan, R.; Hartmann, M.; Tapavicza, E.; von Lilienfeld, O. A. Electronic Spectra from TDDFT and Machine Learning in Chemical Space. *J. Chem. Phys.* **2015**, 143.
- (8) Reeves, C. R.; Bush, D. R. In *Instance Selection and Construction for Data Mining*; Liu, H., Motoda, H., Eds.; Springer US: Boston, MA, 2001; pp 339–356.
- (9) Ventura, D.; Andersen, T.; Martinez, T. R. *Artificial Neural Nets and Genetic Algorithms: Proceedings of the International Conference in Alès, France, 1995*; Springer Vienna: Vienna, 1995; pp 468–471.
- (10) Hansen, K.; Montavon, G.; Biegler, F.; Fazli, S.; Rupp, M.; Scheffler, M.; von Lilienfeld, O. A.; Tkatchenko, A.; Müller, K.-R. Assessment and Validation of Machine Learning Methods for Predicting Molecular Atomization Energies. *J. Chem. Theory Comput.* **2013**, 9, 3404–3419.
- (11) von Lilienfeld, O. A.; Ramakrishnan, R.; Rupp, M.; Knoll, A. Fourier series of Atomic Radial Distribution Functions: A Molecular Fingerprint for Machine Learning Models of Quantum Chemical Properties. *Int. J. Quantum Chem.* **2015**, 115, 1084–1093.
- (12) Hansen, K.; Biegler, F.; Ramakrishnan, R.; Pronobis, W.; von Lilienfeld, O. A.; Müller, K.-R.; Tkatchenko, A. Machine Learning Predictions of Molecular Properties: Accurate Many-Body Potentials and Nonlocality in Chemical Space. *J. Phys. Chem. Lett.* **2015**, 6, 2326–2331.
- (13) Huang, B.; von Lilienfeld, O. A. Communication: Understanding Molecular Representations in Machine Learning: The Role of Uniqueness and Target Similarity. *J. Chem. Phys.* **2016**, 145.

- (14) Ramakrishnan, R.; Dral, P. O.; Rupp, M.; von Lilienfeld, O. A. Quantum Chemistry Structures and Properties of 134 kilo Molecules. *Scientific Data* **2014**, 1.
- (15) Ruddigkeit, L.; Deurson, R. v.; Blum, L. C.; Reymond, J.-L. Enumeration of 166 Billion Organic Small Molecules in the Chemical Universe Database GDB-17. *J. Chem. Inf. Model.* **2012**, 52.
- (16) Reymond, J.-L.; Visini, R.; Awale, M. *Fragment-based Drug Discovery Lessons and Outlook*; Wiley, 2016; pp 57–74.
- (17) Dral, P. O.; von Lilienfeld, O. A.; Thiel, W. Machine Learning of Parameters for Accurate Semiempirical Quantum Chemical Calculations. *J. Chem. Theory Comput.* **2015**, 11, 2120–2125.
- (18) Hastie, T.; Tibshirani, R.; Friedman, J. *The Elements of Statistical Learning* ; Springer Series in Statistics; Springer New York Inc.: New York, NY, USA, 2001.
- (19) Ramakrishnan, R.; von Lilienfeld, O. A. Many Molecular Properties from One Kernel in Chemical Space. *Chimia* **2015**, 69, 182–186.
- (20) Turing, A. M. In *Computers & Thought* ; Feigenbaum, E. A., Feldman, J., Eds.; MIT Press: Cambridge, MA, USA, 1995; Chapter Computing Machinery and Intelligence, pp 11–35.
- (21) Barricelli, N. Esempi Numerici di Processi di Evoluzione. *Methodos* **1954**, 45–68.
- (22) Fraser, A. Simulation of Genetic Systems by Automatic Digital Computers I. Introduction. *Aus. J. Bio. Sci.* **1957**, 484–491.
- (23) Godefroid, P.; Khurshid, S. In *Tools and Algorithms for the Construction and Analysis of Systems*; Katoen, J.-P., Stevens, P., Eds.; Lecture Notes in Computer Science; Springer Berlin Heidelberg, 2002; Vol. 2280; pp 266–280.

- (24) Tew, D. P.; Klopper, W.; Heckert, M.; Gauss, J. Basis Set Limit CCSD(T) Harmonic Vibrational Frequencies. *J. Phys. Chem. A* **2007**, *111*, 11242–11248.
- (25) Hickey, A. L.; Rowley, C. N. Benchmarking Quantum Chemical Methods for the Calculation of Molecular Dipole Moments and Polarizabilities. *J. Phys. Chem. A* **2014**, *118*, 3678–3687.
- (26) Kim, S.; Thiessen, P. A.; Bolton, E. E.; Chen, J.; Fu, G.; Gindulyte, A.; Han, L.; He, J.; He, S.; Shoemaker, B. A. et al. PubChem Substance and Compound databases. *Nucleic Acids Research* **2016**, *44*, D1202.
- (27) Rosenblatt, M. Remarks on Some Nonparametric Estimates of a Density Function. *Ann. Math. Statist.* **1956**, *27*, 832–837.
- (28) Parzen, E. On Estimation of a Probability Density Function and Mode. *Ann. Math. Statist.* **1962**, *33*, 1065–1076.
- (29) Botu, V.; Ramprasad, R. Learning Scheme to Predict Atomic Forces and Accelerate Materials Simulations. *Phys. Rev. B* **2015**, *92*, 094306.
- (30) Virshup, A. M.; Contreras-García, J.; Wipf, P.; Yang, W.; Beratan, D. N. Stochastic Voyages into Uncharted Chemical Space Produce a Representative Library of All Possible Drug-Like Compounds. *J. Am. Chem. Soc.* **2013**, *135*, 7296–7303.
- (31) Sastry, K.; Goldberg, D. E. *Modeling Tournament Selection With Replacement Using Apparent Added Noise*; 2001.
- (32) Spears, W.; Anand, V. In *Methodologies for Intelligent Systems*; Ras, Z., Zemankova, M., Eds.; Lecture Notes in Computer Science; Springer Berlin Heidelberg, 1991; Vol. 542; pp 409–418.
- (33) Spears, W. M.; Jong, K. A. D. *An Analysis of Multi-Point Crossover ; Foundations of Genetic Algorithms*; Elsevier, 1991; Vol. 1; pp 301 – 315.

- (34) Eshelman, L. J. *The {CHC} Adaptive Search Algorithm: How to Have Safe Search When Engaging in Nontraditional Genetic Recombination; Foundations of Genetic Algorithms*; Elsevier, 1991; Vol. 1; pp 265 – 283.
- (35) Deb, K.; Goyal, M. A Combined Genetic Adaptive Search (GeneAS) for Engineering Design. *Computer Science and Informatics* **1996**, 26, 30–45.
- (36) MOPAC2012, J. P. Stewart, J., Stewart Computational Chemistry, Colorado Springs, 2012.
- (37) TURBOMOLE V7.1 2016, a development of University of Karlsruhe and Forschungszentrum Karlsruhe GmbH, 1989-2007, TURBOMOLE GmbH, since 2007; available from <http://www.turbomole.com>.