

SPAHM: the Spectrum of Approximated Hamiltonian Matrices representations

Clémence Cominboeuf (✉ clemence.cominboeuf@epfl.ch)

École Polytechnique Fédérale de Lausanne (EPFL) <https://orcid.org/0000-0001-7993-2879>

Alberto Fabrizio

École Polytechnique Fédérale de Lausanne <https://orcid.org/0000-0002-4440-3149>

Ksenia Briling

École Polytechnique Fédérale de Lausanne (EPFL) <https://orcid.org/0000-0003-2250-9898>

Article

Keywords: SPAHM, quantum machine learning, chemical information

Posted Date: November 3rd, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-1017031/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

SPA^HM: the Spectrum of Approximated Hamiltonian Matrices representations

Alberto Fabrizio,^{1, 2, a)} Ksenia R. Briling,^{1, a)} and Clemence Corminboeuf^{1, 2, b)}

¹⁾*Laboratory for Computational Molecular Design, Institute of Chemical Sciences and Engineering, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland*

²⁾*National Centre for Computational Design and Discovery of Novel Materials (MARVEL), École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland*

(Dated: 25 October 2021)

Physics-based molecular representations are the cornerstone of all modern machine learning techniques applied to solve chemical problems. While several approaches exist to design ever more accurate fingerprints, the majority resolves in including more physics to construct larger and more complex representations. Here, we present an alternative approach to harness the complexity of chemical information into a lightweight numerical form, naturally invariant under real-space transformations, and seamlessly including the information about the charge state of a molecule. The Spectrum of Approximated Hamiltonian Matrices (SPA^HM) leverages the information contained in widely-used and readily-evaluated “guess” Hamiltonians to form a robust fingerprint for quantum machine learning. Relying on the origin of the SPA^HM fingerprints and a hierarchy of approximate Hamiltonians, we analyze the relative merits of adding physics into molecular representations and find that alternative strategies, focusing more on the machine learning task, represent a clear route towards direct improvements.

^{a)}These authors contributed equally to this work.

^{b)}Electronic mail: clemence.corminboeuf@epfl.ch

I. INTRODUCTION

Modern machine learning (ML) techniques are at the forefront of an unprecedented methodological shift affecting virtually all fields of chemistry.¹⁻⁵ Regardless of the chosen application or algorithm, the predicting power of artificial intelligence in chemistry is ultimately related to the choice of a molecular representation, *i.e.* of a numerical descriptor encoding all the relevant information about the chemical system.⁶⁻⁸

The crucial role of representations is mirrored by the intensive work that has been dedicated to find ever more reliable and widely applicable fingerprints.^{7,8} Although there are effectively infinite ways to input the information about a molecule into a machine learning algorithm, conceptually molecular representations could be subdivided into two main categories: cheminformatics and physics-based [or quantum machine learning (QML)] descriptors. The first class relies on string-based fingerprints, such as SMILES^{9,10} and SELFIES,¹¹ or on readily available and descriptive properties, such as the number of aromatic carbon atoms, the shape index of a molecule, and its size,¹²⁻¹⁶ which are usually chosen using an *a priori* knowledge about their correlation with the specific target.¹⁷ In contrast, the second category of representations is inspired by the fundamental laws of physics that govern molecular systems, in particular, the laws of quantum mechanics and the basic laws of symmetry. A third category of chemical representations has been introduced very recently, relying on artificial neural networks (and no human input) to infer suitable descriptors for the learning exercise.¹⁸

As physics-based representations are rooted in fundamental laws, they are directly applicable to any learning task, ranging from the regression of molecular properties to revealing the relationship between molecules in large chemical databases. Although each of the existing descriptors has its own mathematical form and is conceptually different from the others, they all share a common initial ingredient: the structure of the molecule in Cartesian coordinates and the type of atoms that it contains. This starting point is not arbitrary and it is intimately related to the connection between (static) molecular properties and the electronic Hamiltonian $\hat{\mathbf{H}}$.

For a fixed nuclear configuration, the information about all the electronic properties of a molecule is contained in the many-body electronic wavefunction $\Psi(\mathbf{x}_1, \dots, \mathbf{x}_n)$, as defined by the Schrödinger equation. Since the electronic Hamiltonian defines $\Psi(\mathbf{x}_1, \dots, \mathbf{x}_n)$, the

molecular information necessary to fix $\hat{\mathbf{H}}$ is in principle sufficient for a non-linear model to access any electronic property. The expression for all the universal (*i.e.* non-molecule specific) terms of the Hamiltonian (*e.g.* kinetic energy) only require the knowledge of the total number of electrons (N). In contrast, the external potential (the electron-nuclear attraction potential) also depends on the position of the nuclei $\{\mathbf{R}_I\}$ and their charges $\{Z_I\}$.¹⁹ Under the assumption of charge neutrality (*i.e.* $N = \sum_I Z_I$), \mathbf{R}_I and Z_I uniquely fix the form of the Hamiltonian and thus represent the only required information to characterize the electronic wavefunction and molecular quantum chemical properties.

Since no two different molecules have the same Hamiltonian, any representation that relies upon \mathbf{R}_I and Z_I is guaranteed to satisfy the injectivity requirement of machine learning, *i.e.* there must be a one-to-one map between the representation of a molecule and its properties. Nonetheless, injectivity is not the only condition necessary for efficient and transferable learning. A representation must encode the same symmetries as the target property upon any transformation of real-space coordinates: rotation, reflection, translation, and permutation of atoms of the same species.^{6,20,21} Moreover, a representation should be continuous and differentiable, as an underlying assumption of statistical learning is the smoothness of the target properties.^{22,23} Existing physics-based representations differ in the way they use or transform the \mathbf{R}_I and Z_I to comply with these two additional requirements.

Despite the differences in mathematical formulation and physical basis, the QML representations can be classified into three main families. The first contains all those representations that simply use an internal system of coordinates (directly or upon non-linear transformations) to account for the space-transformation invariances of scalar molecular properties. This class was historically the first appearing in the field and includes well-known representations such as Coulomb Matrix (CM),^{24,25} Bag of Bonds (BoB),²⁶ many-body tensor representation (MBTR),²⁷ permutation invariant polynomials (PIPs),^{28–32} and graph-based representations.³³ The second family of physics-based representations is characterized by descriptors mimicking a density of particles. Most notable representations in this category are Behler–Parrinello symmetry functions,^{34–36} the smooth overlap of atomic positions (SOAP),^{6,21} the N -body iterative contraction of equivariant features (NICE)³⁷ and the atomic cluster expansion (ACE).^{38–40} The third and last category includes all those representations that mimic a potential, among which probably the most well-known are (atomic) Spectrum of London and Axilrod–Teller–Muto potential [(a)SLATM],⁴¹ the long-distance

equivariant (LODE) representation,⁴² FCHL18,⁴³ and FCHL19.⁴⁴

Modern QML representations try to incorporate ever more physics by, for instance, increasing the many-body order expansions,^{37–40} including non-local information,⁴² or even by adding equivariance directly in the fingerprints.^{21,37} While this approach has improved the overall accuracy of the machine learning models, it has also led to a generalized increase in the size of the descriptors.⁸ Larger representations are not only detrimental from a computational perspective, but they also pose a potential threat to the robustness of the regression models.^{45,46} As these problems undermine the efficiency of the learning, most of the recent research has focused on the development of reliable feature selection and sparsification techniques, capable of filtering out of the representation all the unnecessary information.^{47–49}

A conceptually different approach is to design simpler, yet physically robust, QML representations. In this work, we propose the Spectrum of Approximated Hamiltonian Matrices (SPA^HM). SPA^HM looks back at the origin of physics-based fingerprints and uses the diagonalization of the electronic Hamiltonian as the central ingredient to generate an input for machine learning algorithms. Naturally, the effective Hamiltonian matrix (Fock matrix) is not a suitable choice to build a representation, as it is the final product of a converged self-consistent field (SCF) computation. Nonetheless, any quantum chemical computation needs a reasonable low-cost Ansatz to initiate the SCF procedure, usually in the form of an approximated Hamiltonian matrix. These matrices are all cheaper to compute than a single SCF iteration and form a controlled hierarchy of increasing complexity and accuracy that is readily computed for any given molecular system. Using well-established SCF guesses, we construct a collection of molecular representations and test their accuracy and their learning abilities for the regression of several quantum chemical properties. The SPA^HM representations are lightweight, physically and chemically motivated, naturally invariant under the basic symmetries of physics, and, in contrast to existing QML representations, include seamlessly the information about the number of electrons and, in principle, the spin state of a molecule.

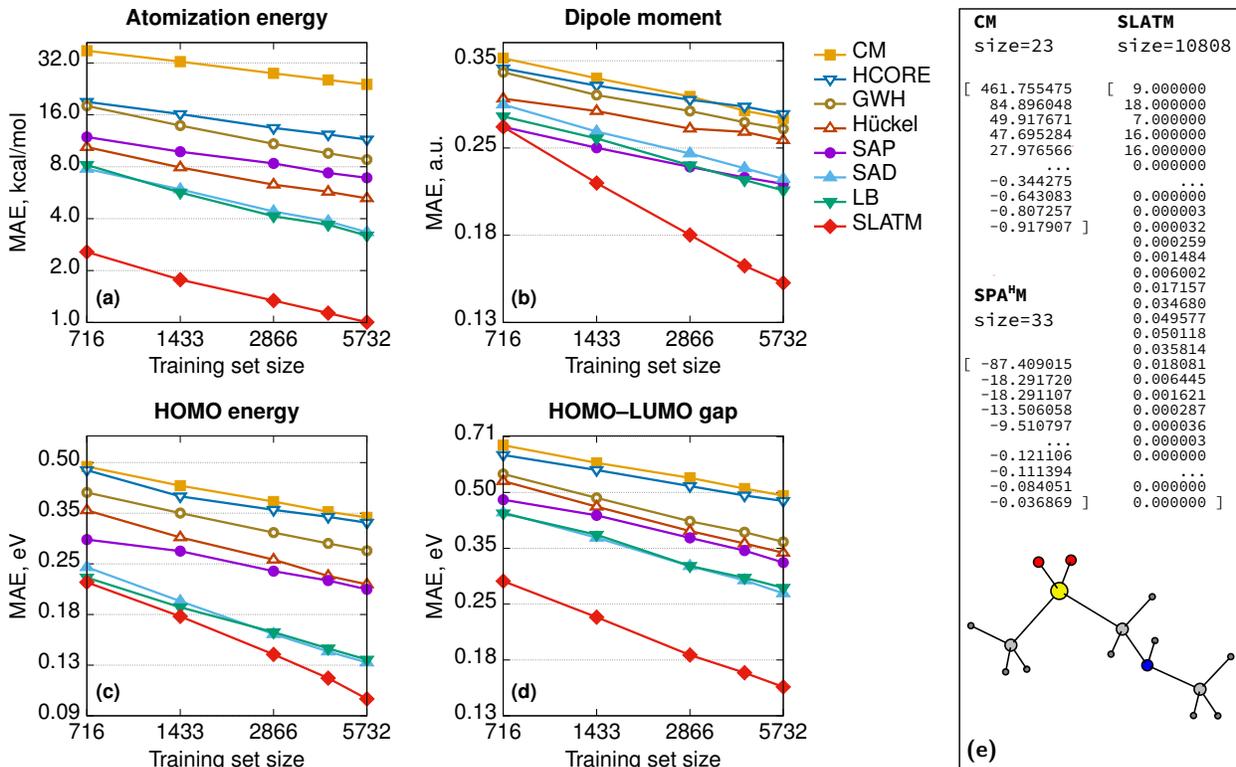


FIG. 1. *Left*. Learning curves in logarithmic scale for (a) atomization energies, (b) dipole moments, (c) HOMO energies and (d) HOMO-LUMO gaps. The color code reflects the different representations. *Right*. Illustrative example of the sizes of the CM, SPA^M, and SLATM representations. All the Hamiltonians were evaluated in the MINAO⁵⁰ minimal basis.

II. RESULTS AND DISCUSSION

A. Learning Curves

To assess the ability of the SPA^M representations to learn and their overall accuracy, we trained a kernel ridge regression (KRR) model on the QM7 database^{24,51} to target four quantum chemical properties. Each of these quantities has been chosen as it is representative of a particular category. Atomization energies are both routinely used to assess the quality of ML models and represent a broader class of extensive (*i.e.* size-dependent) thermodynamic properties.^{24,52} Dipole moments are traditionally used in quantum chemistry as proxies of the quality of the wavefunction. The HOMO energies are intensive (*i.e.* size-independent) quantities. Finally, the HOMO-LUMO gap allows probing the quality of both frontier orbitals and simultaneously tests the additivity of errors in the KRR models.

The QM7 dataset was randomly divided into a training set of 5732 molecules and a test set containing the remaining 1433 compounds, corresponding to an 80%–20% split. For each molecule, we constructed the SPA^HM representations by diagonalizing the different “guess” Hamiltonians in a minimal basis set (MINAO,⁵⁰ see Section IV) and using the sorted occupied eigenvalues as the KRR fingerprint. The occupied orbital energies carry information about both the atom types (core-orbital eigenvalues), the general electronic structure of the molecule (core and valence), and the total number of electrons. In addition, the eigenvalues of a Hamiltonian are naturally invariant under all real-space transformations (permutation, translation, rotation) and the size of the occupied set is independent of the choice of the atomic orbital basis.

The learning curves for all the SPA^HM representations are reported in Figure 1. In addition, we report the curves of the original version (eigenvalue) Coulomb matrix (CM)²⁴ and SLATM,⁴¹ as the first has a similar size and diagonalization philosophy as SPA^HM and the second is an example of a widely-used global representation.

In the Figure 1, the SPA^HM representations are indicated by the type of approximate Hamiltonian used for their construction. Although it is rather complex to establish a definitive hierarchy of self-consistent-field guesses, it is always possible to provide a more qualitative trend based on the amount of physics that each guess includes. The diagonalization of the core Hamiltonian (\mathbf{H}_{core}) and the generalized Wolfsberg–Helmholz (GWH)⁵³ are the simplest approximations, as they do not try to model any two-electron term. Building on the GWH guess, the extended Hückel method uses approximate ionization potentials as diagonal terms and it is generally more robust.⁵⁴ The superposition of atomic densities (SAD)^{55,56} is another popular choice that however only produces a density matrix (DM). Nonetheless, it is rather straightforward to construct a complete Hamiltonian matrix (including the one- and two-electron terms) by contracting the SAD density matrix with a potential of choice (Hartree–Fock or any exchange–correlation density functional). We report the SAD learning curve in Figure 1 using the PBE0 potential, as all the properties were computed with this functional. Finally, the superposition of atomic potentials (SAP)^{57,58} and the Laikov–Briling (LB)⁵⁹ guesses use effective one-electron potentials to construct sophisticated, yet computationally lightweight, guess Hamiltonians.

Besides the internal hierarchy, the accuracy of all the SPA^HMs is always comprised between SLATM and the eigenvalues of Coulomb Matrix. While SLATM consistently out-

performs the SPA^HM representations, the difference with the most robust guesses (LB and SAD) is usually much smaller than the accuracy of the functional itself (~ 5 kcal/mol).⁶⁰ Importantly, SLATM is also three orders of magnitude larger than SPA^HM on QM7. The significant difference in the extent of the representation is crucial from an efficiency perspective, as the number of features dictates the computational effort of constructing the kernel matrix for an equal size of the training set. While lightweight, efficient, and naturally accounting for the charge state of molecules, we only tested a few well-known Hamiltonians for building SPA^HM. As the performance of SPA^HM is largely independent from the choice of the basis set or potential (see Supplementary Information), it is necessary to consider alternative strategies to improve its accuracy. The heavy dependence of the learning on the quality of the parent Hamiltonian suggests that the construction of “better guesses” is the correct direction. Nonetheless, as discussed in Section II B, *better guesses* does not necessarily mean “improved quantum chemical approximate Hamiltonians” (*i.e.* closer to the converged Fock matrix), but rather the construction of simpler, systematic Hamiltonians specifically optimized for the learning task.

B. Physics and noise

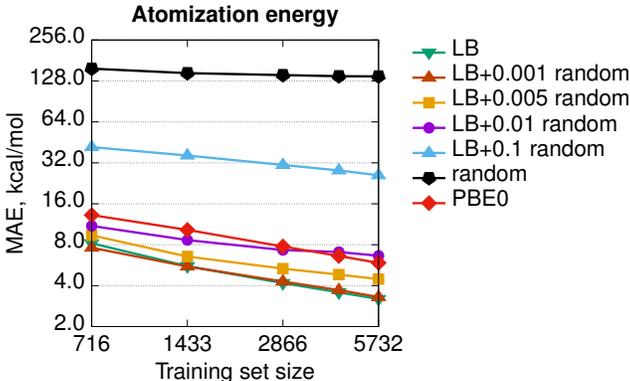


FIG. 2. Learning curves (in logarithmic scale) of atomization energies for SPA^HM based on the converged PBE0 Fock matrix and on the LB Hamiltonian, with an increasing (pseudo-)random perturbation added to the representation vector. The perturbation max. magnitude is reflected in the legend. For comparison, a learning curve for a fully-random representation is shown. All the Hamiltonians (including converged PBE0) were evaluated in the MINAO⁵⁰ minimal basis.

In general, the accuracy of the different SPA^HM representations (Figure 1) follows the same trend as the complexity of the underlying SCF guess. This result seems to suggest that the more physics is included in the approximate Hamiltonian, the easier is the learning exercise for the corresponding SPA^HM representation. To test the robustness of this conclusion, we constructed a *test* representation using the converged PBE0 Hamiltonian matrix (Figure 2, label PBE0). As already mentioned, any representation based on the converged Fock matrix is both too expensive and worthless for practical machine learning, since it is always possible to (upon diagonalization) use the converged wavefunction to compute any desired quantum chemical property. Nonetheless, this test is essential, as it pushes the physics of SPA^HM to the limit. Figure 2 shows that PBE0 is not the best representation when regressing the atomization energies and even some SPA^HMs outperform its accuracy. As the SCF changes the eigenvalues of each molecule independently from the others, the relationship between the feature vectors also varies unconcertedly. This sparsification of the data in the representation space effectively decreases the correlation between the features and the target properties and worsens the learning.

The performance of the converged Fock matrix *versus* more approximated (and computationally cheaper) Hamiltonians shows that “more physics” is not necessarily the key to better learning. Yet, the relative ordering of the guess Hamiltonians suggests that higher-quality potentials correlate with the best representations. The question associated with the relevance of the physics could be generalized and one could ask if there is the need for any physics at all or random featurization could lead to the same (or better) results.⁶¹ In addition, SPA^HM representations are so small (33 features for QM7) with respect to the size of the dataset (7165 molecules), that the learning could be the result of random correlations between the features and the target properties. Overall, it is still unclear if any random (*i.e.* not physically motivated) perturbation of SPA^HM could lead to better learning. To analyze the behavior of our most robust representation upon random perturbation, we modify SPA^HM(LB) by adding an increasing (pseudo-)random perturbation sampled from a uniform distribution and testing its accuracy on the QM7 database.

Figure 2 shows that, for the smallest perturbation tested (magnitude max. 0.001), the original and the modified learning curves are almost indistinguishable, except for a non-significant difference due to the shuffling of the training set. As the allowed random noise increases, we observe the systematic and progressive worsening of the learning exercise towards

the limit of physically meaningless random numbers. As the magnitude of the perturbation increases, the physics in the representation fades, the error increases, and the learning curves become flatter. This demonstrates that the performance of SPA^HM is not just a consequence of a random correlation between the feature vectors and the properties.

C. Core and Valence

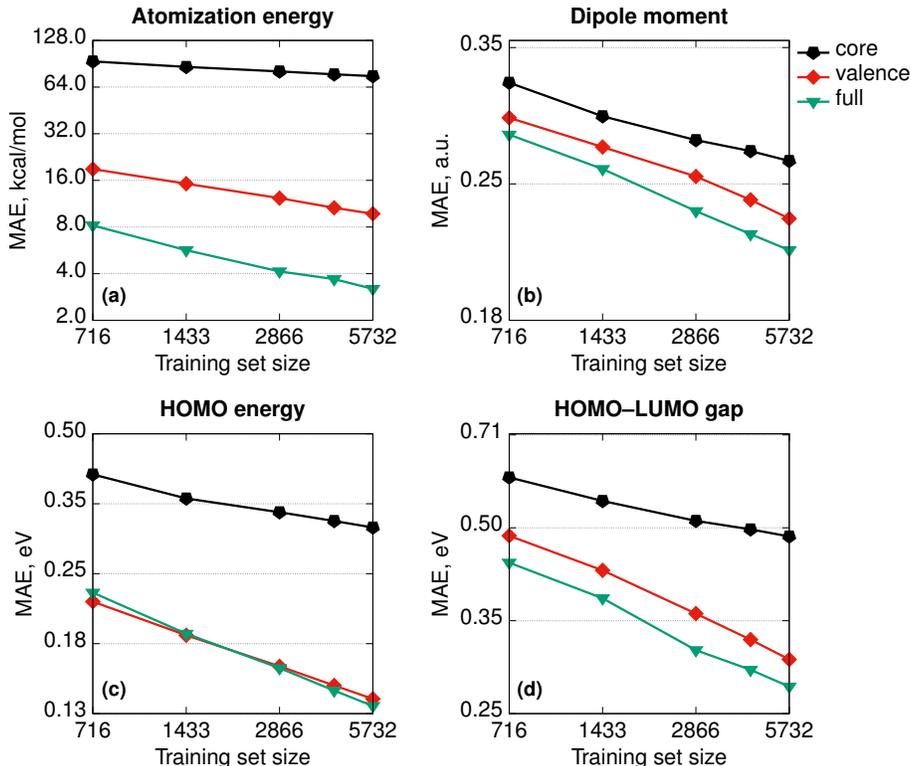


FIG. 3. Learning curves (in logarithmic scale) for the SPA^HM based on the LB Hamiltonian with core, valence, and the full occupied eigenvalues sets used as representations. All the Hamiltonians were evaluated in the MINAO⁵⁰ minimal basis.

Besides any consideration about the relative importance of physics, the SPA^HM representations are still the eigenvalues of (approximate) electronic Hamiltonians. As such, it is relevant to try to rationalize how the different parts of SPA^HM contribute to the learning. As for any set of eigenvalues, it is always possible to divide any SPA^HM representation into its core and valence parts. The core orbital energies do not vary significantly for the same atom in different molecules and therefore naturally track the number and the types of nuclei

across the database. In contrast, the valence set is a fingerprint of the chemistry and the bonding patterns proper to each molecule.

Figure 3 shows the learning curves for the SPA^HMs based on the LB Hamiltonian with core, valence, and the full occupied eigenvalues sets taken as representation vectors. While the valence set results in a consistently better learning than the core, both are necessary to achieve the overall accuracy of SPA^HM(LB) with the exception of the HOMO energies. For the HOMO eigenvalues, the valence set can be considered an alternative type of Δ -learning,⁶² where the approximate baseline (approximate HOMO energies) are the input of the kernel itself, rather than corrected *a posteriori*. Importantly, core orbital energies are not sufficient information to accurately regress the atomization energies, but the valence set error is also twice as large as the total SPA^HM. Therefore, while the information about the chemical bonding is more relevant for the general performance, the information about the number and the type of nuclei in the molecules is essential to improve the learning.

III. CONCLUSIONS

In this work, we proposed a lightweight and efficient quantum machine learning representation, capable of naturally accounting for the charge state of molecules by leveraging the information contained in standard quantum chemical “guess” Hamiltonians. Using the QM7 database, we tested the performance of a hierarchy of representations for a set of four representative quantum chemical properties. The performance of the SPA^HM representations follows the same qualitative trend as the one describing the amount of physics encoded in the parent approximate Hamiltonian. Nonetheless, we also find that the trend stops when pushing the physics to the limit and using the fully converged Fock matrices to construct the representation. Since increasingly adding physics is not the roadmap for the potential improvement of the SPA^HM representations, alternative strategies have to be analyzed. A particular appealing plan is the development of ML-oriented “guesses” that sacrifice resemblance to the converged Hamiltonians, to promote formal simplicity and being systematic across molecular databases. Sharing the same conceptual origin of this work, an alternative strategy consists in representing molecules for ML using the approximate Hamiltonian matrix itself, its eigenvectors, or the resulting density matrices. Together with the SPA^HMs, these descriptors form a more comprehensive class of Hamiltonian-centered fingerprints that

is conceptually alternative to any existing quantum machine learning representation.

IV. METHODS

Out of the total set of 7165 molecules²⁴, 5732 molecules were randomly selected as the training set and the remaining 1433 as the test set. The hyperparameters for each representation were optimized with a grid search using a 5-fold cross validation procedure and the learning curves were computed using random sub-sampling (5 times per point). The optimization and regression code was written in Python using the `numpy`⁶³ and `scikit-learn`⁶⁴ libraries. The `QML`⁶⁵ package was used to construct the CM and SLATM representations.

While the atomization energies were taken as computed in the original QM7 reference²⁴ (PBE0 in a converged numerical basis), the other three properties (norm of dipole moment, HOMO energies and HOMO–LUMO gap) were computed at PBE0⁶⁶/cc-pVQZ^{67,68} level. The initial guesses were obtained in a minimal basis (MINAO⁵⁰). All quantum chemical computations were made with a locally modified version of PySCF.^{69,70} The codes used in this paper are included in a more comprehensive package called Q-stack (<https://github.com/lcmd-epfl/Q-stack>). Q-stack is a library and a collection of stand-alone codes, mainly programmed in Python, that provides custom quantum chemistry operations to promote quantum machine learning.

DATA AVAILABILITY

The data and the model that support the findings of this study are freely available in Materials Cloud.

CODE AVAILABILITY

The code is available in Q-stack (<https://github.com/lcmd-epfl/Q-stack>) and as a separate Github repository at <https://github.com/lcmd-epfl/SPAHM>.

REFERENCES

- ¹Huang, B. & von Lilienfeld, O. A. Ab initio machine learning in chemical compound space. *Chem. Rev.* **121** (16), 10001–10036 (2021).
- ²Dral, P. O. & Barbatti, M. Molecular excited states through a machine learning lens. *Nat. Rev. Chem.* **5** (6), 388–405 (2021).
- ³Jorner, K., Tomberg, A., Bauer, C., Sköld, C. & Norrby, P.-O. Organic reactivity from mechanism to machine learning. *Nat. Rev. Chem.* **5** (4), 240–255 (2021).
- ⁴Unke, O. T. *et al.* Machine learning force fields. *Chem. Rev.* **121** (16), 10142–10186 (2021).
- ⁵Meuwly, M. Machine learning for chemical reactions. *Chem. Rev.* **121** (16), 10218–10239 (2021).
- ⁶Bartók, A. P., Kondor, R. & Csányi, G. On representing chemical environments. *Phys. Rev. B* **87** (18), 184115 (2013).
- ⁷Huang, B. & von Lilienfeld, O. A. Communication: Understanding molecular representations in machine learning: The role of uniqueness and target similarity. *J. Chem. Phys.* **145** (16), 161102 (2016).
- ⁸Musil, F. *et al.* Physics-inspired structural representations for molecules and materials. *Chem. Rev.* **121** (16), 9759–9815 (2021).
- ⁹Weininger, D. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *J. Chem. Inf. Model.* **28** (1), 31–36 (1988).
- ¹⁰Weininger, D., Weininger, A. & Weininger, J. L. SMILES. 2. Algorithm for generation of unique SMILES notation. *J. Chem. Inf. Comput. Sci.* **29** (2), 97–101 (1989).
- ¹¹Krenn, M., Häse, F., Nigam, A., Friederich, P. & Aspuru-Guzik, A. Self-referencing embedded strings (SELFIES): A 100% robust molecular string representation. *Mach. Learn. Sci. Technol.* **1** (4), 045024 (2020).
- ¹²Karelson, M., Lobanov, V. S. & Katritzky, A. R. Quantum-chemical descriptors in QSAR/QSPR studies. *Chem. Rev.* **96** (3), 1027–1044 (1996).
- ¹³Karelson, M. *Molecular descriptors in QSAR/QSPR* (Wiley-Interscience, New York, 2000).
- ¹⁴Todeschini, R. & Consonni, V. *Handbook of molecular descriptors* Vol. 11 of *Methods and Principles in Medicinal Chemistry* (Wiley, 2000).
- ¹⁵Todeschini, R. & Consonni, V. *Molecular Descriptors for Chemoinformatics* Vol. 41 of *Methods and Principles in Medicinal Chemistry* (Wiley, 2009).

- ¹⁶David, L., Thakkar, A., Mercado, R. & Engkvist, O. Molecular representations in AI-driven drug discovery: a review and practical guide. *J. Cheminformatics* **12** (1), 1–22 (2020).
- ¹⁷Ghiringhelli, L. M. *et al.* Learning physical descriptors for materials science by compressed sensing. *New J. Phys.* **19** (2), 023017 (2017).
- ¹⁸Schütt, K. T., Unke, O. T. & Gastegger, M. Equivariant message passing for the prediction of tensorial properties and molecular spectra (2021). arXiv:2102.03150.
- ¹⁹Szabo, A. & Ostlund, N. S. *Modern quantum chemistry: Introduction to advanced electronic structure theory* (McGraw-Hill, New York, 1989).
- ²⁰Glielmo, A., Sollich, P. & De Vita, A. Accurate interatomic force fields via machine learning with covariant kernels. *Phys. Rev. B* **95** (21), 214302 (2017).
- ²¹Grisafi, A., Wilkins, D. M., Csányi, G. & Ceriotti, M. Symmetry-adapted machine learning for tensorial properties of atomistic systems. *Phys. Rev. Lett.* **120** (3), 036002 (2018).
- ²²Smola, A. J., Schölkopf, B. & Müller, K.-R. The connection between regularization operators and support vector kernels. *Neural Netw.* **11** (4), 637–649 (1998).
- ²³Rudi, A. & Rosasco, L. *Generalization properties of learning with random features. Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, 3218–3228 (Curran Associates Inc., Red Hook, NY, USA, 2017).
- ²⁴Rupp, M. *et al.* Fast and accurate modeling of molecular atomization energies with machine learning. *Phys. Rev. Lett.* **108** (5), 058301 (2012).
- ²⁵Rupp, M., Ramakrishnan, R. & von Lilienfeld, O. A. Machine learning for quantum mechanical properties of atoms in molecules. *J. Phys. Chem. Lett.* **6** (16), 3309–3313 (2015).
- ²⁶Hansen, K. *et al.* Machine learning predictions of molecular properties: Accurate many-body potentials and nonlocality in chemical space. *J. Phys. Chem. Lett.* **6** (12), 2326–2331 (2015).
- ²⁷Huo, H. & Rupp, M. Unified representation of molecules and crystals for machine learning (2017). arXiv:1704.06439.
- ²⁸Braams, B. J. & Bowman, J. M. Permutationally invariant potential energy surfaces in high dimensionality. *Int. Rev. Phys. Chem.* **28** (4), 577–606 (2009).
- ²⁹Brown, A., McCoy, A. B., Braams, B. J., Jin, Z. & Bowman, J. M. Quantum and classical studies of vibrational motion of CH_5^+ on a global potential energy surface obtained from a

- novel *ab initio* direct dynamics approach. *J. Chem. Phys.* **121** (9), 4105–4116 (2004).
- ³⁰Bowman, J. M. *et al.* Ab-initio-based potential energy surfaces for complex molecules and molecular complexes. *J. Phys. Chem. Lett.* **1** (12), 1866–1874 (2010).
- ³¹Xie, Z. & Bowman, J. M. Permutationally invariant polynomial basis for molecular energy surface fitting via monomial symmetrization. *J. Chem. Theory Comput.* **6** (1), 26–34 (2010).
- ³²Jiang, B. & Guo, H. Permutation invariant polynomial neural network approach to fitting potential energy surfaces. *J. Chem. Phys.* **139** (5), 054112 (2013).
- ³³Pietrucci, F. & Andreoni, W. Graph theory meets *ab initio* molecular dynamics: Atomic structures and transformations at the nanoscale. *Phys. Rev. Lett.* **107** (8), 085504 (2011).
- ³⁴Behler, J. & Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **98** (14), 146401 (2007).
- ³⁵Behler, J. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J. Chem. Phys.* **134** (7), 074106 (2011).
- ³⁶Zhang, L., Han, J., Wang, H., Car, R. & E, W. Deep potential molecular dynamics: A scalable model with the accuracy of quantum mechanics. *Phys. Rev. Lett.* **120** (14), 143001 (2018).
- ³⁷Nigam, J., Pozdnyakov, S. & Ceriotti, M. Recursive evaluation and iterative contraction of n -body equivariant features. *J. Chem. Phys.* **153** (12), 121101 (2020).
- ³⁸Drautz, R. Erratum: Atomic cluster expansion for accurate and transferable interatomic potentials [Phys. Rev. B 99, 014104 (2019)]. *Phys. Rev. B* **100** (24), 249901 (2019).
- ³⁹Drautz, R. Atomic cluster expansion for accurate and transferable interatomic potentials. *Phys. Rev. B* **99** (1), 014104 (2019).
- ⁴⁰Dusson, G. *et al.* Atomic cluster expansion: Completeness, efficiency and stability (2019). arXiv:1911.03550.
- ⁴¹Huang, B. & von Lilienfeld, O. A. Quantum machine learning using atom-in-molecule-based fragments selected on the fly. *Nat. Chem.* **12** (10), 945–951 (2020).
- ⁴²Grisafi, A. & Ceriotti, M. Incorporating long-range physics in atomic-scale machine learning. *J. Chem. Phys.* **151** (20), 204105 (2019).
- ⁴³Faber, F. A., Christensen, A. S., Huang, B. & von Lilienfeld, O. A. Alchemical and structural distribution based representation for universal quantum machine learning. *J. Chem. Phys.* **148** (24), 241717 (2018).

- ⁴⁴Christensen, A. S., Bratholm, L. A., Faber, F. A. & von Lilienfeld, O. A. FCHL revisited: Faster and more accurate quantum machine learning. *J. Chem. Phys.* **152** (4), 044107 (2020).
- ⁴⁵Tibshirani, R. Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B* **58** (1), 267–288 (1996).
- ⁴⁶Hastie, T., Friedman, J. & Tibshirani, R. *The Elements of Statistical Learning* (Springer, New York, 2001).
- ⁴⁷Guyon, I. & Elisseeff, A. An introduction to variable and feature selection. *J. Mach. Learn. Res.* **3**, 1157–1182 (2003).
- ⁴⁸Eklund, M., Norinder, U., Boyer, S. & Carlsson, L. Choosing feature selection and learning algorithms in QSAR. *J. Chem. Inf. Model.* **54** (3), 837–843 (2014).
- ⁴⁹Cersonsky, R. K., Helfrecht, B. A., Engel, E. A., Kliavinek, S. & Ceriotti, M. Improving sample and feature selection with principal covariates regression. *Mach. Learn. Sci. Technol.* **2** (3), 035038 (2021).
- ⁵⁰Knizia, G. Intrinsic atomic orbitals: An unbiased bridge between quantum theory and chemical concepts. *J. Chem. Theory Comput.* **9** (11), 4834–4843 (2013).
- ⁵¹Blum, L. C. & Reymond, J.-L. 970 million druglike small molecules for virtual screening in the chemical universe database GDB-13. *J. Am. Chem. Soc.* **131** (25), 8732–8733 (2009).
- ⁵²Ruddigkeit, L., Van Deursen, R., Blum, L. C. & Reymond, J. L. Enumeration of 166 billion organic small molecules in the chemical universe database GDB-17. *J. Chem. Inf. Model.* **52** (11), 2864–2875 (2012).
- ⁵³Wolfsberg, M. & Helmholz, L. The spectra and electronic structure of the tetrahedral ions MnO_4^- , CrO_4^{2-} , and ClO_4^- . *J. Chem. Phys.* **20** (5), 837–843 (1952).
- ⁵⁴Hoffmann, R. An extended Hückel theory. I. Hydrocarbons. *J. Chem. Phys.* **39** (6), 1397–1412 (1963).
- ⁵⁵Amat, L. & Carbó-Dorca, R. Use of promolecular ASA density functions as a general algorithm to obtain starting MO in SCF calculations. *Int. J. Quantum Chem.* **87** (2), 59–67 (2001).
- ⁵⁶Van Lenthe, J. H., Zwaans, R., Van Dam, H. J. J. & Guest, M. F. Starting SCF calculations by superposition of atomic densities. *J. Comput. Chem.* **27** (8), 926–932 (2006).
- ⁵⁷Lehtola, S. Assessment of initial guesses for self-consistent field calculations. Superposition of atomic potentials: Simple yet efficient. *J. Chem. Theory Comput.* **15** (3), 1593–1604

- (2019).
- ⁵⁸Lehtola, S., Visscher, L. & Engel, E. Efficient implementation of the superposition of atomic potentials initial guess for electronic structure calculations in Gaussian basis sets. *J. Chem. Phys.* **152** (14), 144105 (2020).
- ⁵⁹Laikov, D. N. & Briling, K. R. Atomic effective potentials for starting molecular electronic structure calculations. *Theor. Chem. Acc.* **139** (1), 17 (2020).
- ⁶⁰Lynch, B. J. & Truhlar, D. G. Robust and affordable multicoefficient methods for thermochemistry and thermochemical kinetics: The MCCM/3 suite and SAC/3. *J. Phys. Chem. A* **107** (19), 3898–3906 (2003).
- ⁶¹Chuang, K. V. & Keiser, M. J. Comment on “Predicting reaction performance in C–N cross-coupling using machine learning”. *Science* **362** (6416), eaat8603 (2018).
- ⁶²Ramakrishnan, 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.* **11** (5), 2087–2096 (2015).
- ⁶³Harris, C. R. *et al.* Array programming with NumPy. *Nature* **585** (7825), 357–362 (2020).
- ⁶⁴Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
- ⁶⁵Christensen, A. S. *et al.* QML: A Python toolkit for quantum machine learning (2017).
- ⁶⁶Adamo, C. & Barone, V. Toward reliable density functional methods without adjustable parameters: The PBE0 model. *J. Chem. Phys.* **110** (13), 6158–6170 (1999).
- ⁶⁷Dunning, T. H. Gaussian basis sets for use in correlated molecular calculations. I. The atoms boron through neon and hydrogen. *J. Chem. Phys.* **90** (2), 1007–1023 (1989).
- ⁶⁸Woon, D. E. & Dunning, T. H. Gaussian basis sets for use in correlated molecular calculations. III. The atoms aluminum through argon. *J. Chem. Phys.* **98** (2), 1358–1371 (1993).
- ⁶⁹Sun, Q. Libcint: An efficient general integral library for Gaussian basis functions. *J. Comput. Chem.* **36** (22), 1664–1671 (2015).
- ⁷⁰Sun, Q. *et al.* PySCF: the Python-based simulations of chemistry framework. *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **8** (1), e1340 (2017).

ACKNOWLEDGEMENTS

The authors acknowledge the National Centre of Competence in Research (NCCR) “Materials’ Revolution: Computational Design and Discovery of Novel Materials (MARVEL)” of the Swiss National Science Foundation (SNSF), the European Research Council (ERC, grant agreement no 817977), and the EPFL for financial support.

AUTHOR CONTRIBUTIONS

A.F. and K.R.B. performed the computations and wrote the computer code. A.F. and C.C. designed the representations and the project. All the authors contributed to the writing of the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [SupplementaryInformation.pdf](#)