




SpookyNet: Learning force fields with electronic degrees of freedom and nonlocal effects

Oliver T. Unke ^{1,2}✉, Stefan Chmiela ¹, Michael Gastegger ^{1,2}, Kristof T. Schütt¹, Huziel E. Saucedo ^{1,3} & Klaus-Robert Müller ^{1,4,5,6,7}✉

Machine-learned force fields combine the accuracy of ab initio methods with the efficiency of conventional force fields. However, current machine-learned force fields typically ignore electronic degrees of freedom, such as the total charge or spin state, and assume chemical locality, which is problematic when molecules have inconsistent electronic states, or when nonlocal effects play a significant role. This work introduces SpookyNet, a deep neural network for constructing machine-learned force fields with explicit treatment of electronic degrees of freedom and nonlocality, modeled via self-attention in a transformer architecture. Chemically meaningful inductive biases and analytical corrections built into the network architecture allow it to properly model physical limits. SpookyNet improves upon the current state-of-the-art (or achieves similar performance) on popular quantum chemistry data sets. Notably, it is able to generalize across chemical and conformational space and can leverage the learned chemical insights, e.g. by predicting unknown spin states, thus helping to close a further important remaining gap for today's machine learning models in quantum chemistry.

¹Machine Learning Group, Technische Universität Berlin, 10587 Berlin, Germany. ²DFG Cluster of Excellence “Unifying Systems in Catalysis” (UniSysCat), Technische Universität Berlin, 10623 Berlin, Germany. ³BASLEARN, BASF-TU joint Lab, Technische Universität Berlin, 10587 Berlin, Germany. ⁴Department of Artificial Intelligence, Korea University, Anam-dong, Seongbuk-gu, Seoul 02841, Korea. ⁵Max Planck Institute for Informatics, Stuhlsatzenhausweg, 66123 Saarbrücken, Germany. ⁶BIFOLD—Berlin Institute for the Foundations of Learning and Data, Berlin, Germany. ⁷Google Research, Brain team, Berlin, Germany. ✉email: oliver.unke@googlemail.com; klaus-robot.mueller@tu-berlin.de

Molecular dynamics (MD) simulations of chemical systems allow to gain insights on many intricate phenomena, such as reactions or the folding of proteins¹. To perform MD simulations, knowledge of the forces acting on individual atoms at every time step of the simulation is required². The most accurate way of deriving these forces is by (approximately) solving the Schrödinger equation (SE), which describes the physical laws governing chemical systems³. Unfortunately, the computational cost of accurate ab initio approaches⁴ makes them impractical when many atoms are studied, or the simulation involves thousands (or even millions) of time steps. For this reason, it is common practice to use force fields (FFs)—analytical expressions for the potential energy of a chemical system, from which forces are obtained by derivation—instead of solving the SE⁵. The remaining difficulty is to find an appropriate functional form that gives forces at the required accuracy.

Recently, machine learning (ML) methods have gained increasing popularity for addressing this task^{6–15}. They allow to automatically learn the relation between chemical structure and forces from ab initio reference data. The accuracy of such ML-FFs (also known as machine learning potentials) is limited by the quality of the data used to train them and their computational efficiency is comparable to conventional FFs^{11,16}.

One of the first methods for constructing ML-FFs for high-dimensional systems was introduced by Behler and Parrinello for studying the properties of bulk silicon¹⁷. The idea is to encode the local (within a certain cutoff radius) chemical environment of each atom in a descriptor, e.g., using symmetry functions¹⁸, which is used as input to an artificial neural network¹⁹ predicting atomic energies. The total potential energy of the system is modeled as the sum of the individual contributions, and forces are obtained by derivation with respect to atom positions. Alternatively, it is also possible to directly predict the total energy (or forces) without relying on a partitioning into atomic contributions^{20–22}. However, an atomic energy decomposition makes predictions size extensive and the learned model applicable to systems of different size. Many other ML-FFs follow this design principle, but rely on different descriptors^{23–25} or use other ML methods^{6,8,26}, such as kernel machines^{7,27–32}, for the prediction. An alternative to manually designed descriptors is to use the raw atomic numbers and Cartesian coordinates as input instead. Then, suitable atomic representations can be learned from (and adapted to) the reference data automatically. This is usually achieved by “passing messages” between atoms to iteratively build increasingly sophisticated descriptors in a deep neural network architecture. After the introduction of the deep tensor neural network (DTNN)³³, such message-passing neural networks (MPNNs)³⁴ became highly popular and the original architecture has since been refined by many related approaches^{35–37}.

However, atomic numbers and Cartesian coordinates (or descriptors derived from them) do not provide an unambiguous description of chemical systems³⁸. They only account for the nuclear degrees of freedom, but contain no information about electronic structure, such as the total charge or spin state. This is of no concern when all systems of interest have a consistent electronic state (e.g., they are all neutral singlet structures), but leads to an ill-defined learning problem otherwise (Fig. 1a). Further, most ML-FFs assume that atomic properties are dominated by their local chemical environment¹¹. While this approximation is valid in many cases, it still neglects that quantum systems are inherently nonlocal in nature, a quality which Einstein famously referred to as “*spooky actions at a distance*”³⁹. For example, electrons can be delocalized over a chemical system and charge or spin density may instantaneously redistribute to specific atoms based on distant structural changes (Fig. 1b)^{40–44}.

ML-FFs have only recently begun to address these issues. For example, the charge equilibration neural network technique (CENT)⁴⁵ was developed to construct interatomic potentials for ionic systems. In CENT, a neural network predicts atomic electronegativities (instead of energy contributions), from which partial charges are derived via a charge equilibration scheme^{46–48} that minimizes the electrostatic energy of the system and models nonlocal charge transfer. Then, the total energy is obtained by an analytical expression involving the partial charges. Since they are constrained to conserve the total charge, different charge states of the same chemical system can be treated by a single model. The recently proposed fourth-generation Behler-Parinello neural network (4G-BPNN)⁴⁹ expands on this idea using two separate neural networks: The first one predicts atomic electronegativities, from which partial charges are derived using the same method as in CENT. The second neural network predicts atomic energy contributions, receiving the partial charges as additional inputs, which contain implicit information about the total charge. The charge equilibration scheme used in CENT and 4G-BPNNs involves the solution of a system of linear equations, which formally scales cubically with the number of atoms, although iterative solvers can be used to reduce the complexity⁴⁹. Unfortunately, only different total charges, but not spin states, can be distinguished with this approach. In contrast, neural spin equilibration (NSE)⁵⁰, a recently proposed modification to the AIMNet model⁵¹, distinguishes between α and β -spin charges, allowing it to also treat different spin states. In the NSE method, a neural network predicts initial (spin) charges from descriptors that depend only on atomic numbers and coordinates. The discrepancy between predictions and true (spin) charges is then used to update the descriptors and the procedure is repeated until convergence. An alternative approach to include information about the charge and spin state is followed by the BpopNN model⁵². In this method, electronic information is encoded indirectly by including spin populations when constructing atomic descriptors. However, this requires running (constrained) density functional theory calculations to derive the populations before the model can be evaluated. A similar approach is followed by OrbNet⁵³: Instead of spin populations, the atomic descriptors are formed from the expectation values of several quantum mechanical operators in a symmetry-adapted atomic orbital basis.

The present work introduces SpookyNet, a deep MPNN which takes atomic numbers, Cartesian coordinates, the number of electrons, and the spin state as direct inputs. It does not rely on equilibration schemes, which often involve the costly solution of a linear system, or values derived from ab initio calculations, to encode the electronic state. Our end-to-end learning approach is shared by many recent ML methods that aim to solve the Schrödinger equation^{54–56} and mirrors the inputs that are also used in ab initio calculations. To model local interactions between atoms, early MPNNs relied on purely distance-based messages^{33,35,37}, whereas later architectures such as DimeNet⁵⁷ proposed to include angular information in the feature updates. However, explicitly computing angles between all neighbors of an atom scales quadratically with the number of neighbors. To achieve linear scaling, SpookyNet encodes angular information implicitly via the use of basis functions based on Bernstein polynomials⁵⁸ and spherical harmonics. Spherical harmonics are also used in neural network architectures for predicting rotationally equivariant quantities, such as tensor field networks⁵⁹, Cormorant⁶⁰, PaiNN⁶¹, or NequIP⁶². However, since only scalar quantities (energies) need to be predicted for constructing ML-FFs, SpookyNet projects rotationally equivariant features to invariant representations for computational efficiency. Many methods for constructing descriptors of atomic environments use similar approaches to derive rotationally invariant quantities

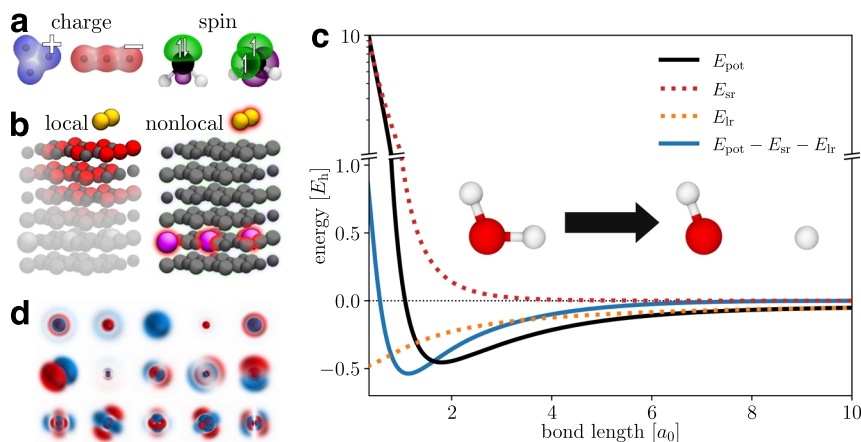


Fig. 1 Main features of SpookyNet and problems addressed in this work. **a** Optimized geometries of $\text{Ag}_3^+/\text{Ag}_3^-$ (left) and singlet/triplet CH_2 (right). Without information about the electronic state (charge/spin), machine learning models are unable to distinguish between the different structures. **b** Au_2 dimer on a $\text{MgO}(001)$ surface doped with Al atoms (Au: yellow, Mg: gray, O: red, Al: pink). The presence of Al atoms in the crystal influences the electronic structure and affects Au_2 binding to the surface, an effect which cannot be adequately described by only local interactions. **c** Potential energy E_{pot} (solid black) for O-H bond dissociation in water. The asymptotic behavior of E_{pot} for very small and very large bond lengths can be well-approximated by analytical short-ranged E_{sr} (dotted red) and long-ranged E_{lr} (dotted orange) energy contributions, which follow known physical laws. When they are subtracted from E_{pot} , the remaining energy (solid blue) covers a smaller range of values and decays to zero quicker, which simplifies the learning problem. **d** Visualization of a random selection of learned interaction functions for SpookyNet trained on the QM7-X⁷¹ dataset. They are designed to closely resemble atomic orbitals, facilitating SpookyNet's ability to extract chemical insight from data.

from spherical harmonics^{25,63,64}. In addition, SpookyNet allows to model quantum nonlocality and electron delocalization explicitly by introducing a nonlocal interaction between atoms, which is independent of their distance. Following previous works, its energy predictions are augmented with physically motivated corrections to improve the description of long-ranged electrostatic and dispersion interactions^{37,65,66}. Additionally, SpookyNet explicitly models short-ranged repulsion between nuclei. Such corrections simplify the learning problem and guarantee correct asymptotic behaviour (Fig. 1c). As such, SpookyNet is a hybrid between a pure ML approach and a classical FF. However, it is much closer to a pure ML approach than methods like IPML⁶⁷, which rely exclusively on parametrized functions (known from classical FFs) for modeling the potential energy, but predict environment-dependent parameters with ML methods. Further, inductive biases in SpookyNet's architecture encourage learning of atomic representations which capture similarities between different elements and interaction functions designed to resemble atomic orbitals, allowing it to efficiently extract meaningful chemical insights from data (Fig. 1d).

Results

SpookyNet architecture. SpookyNet takes sets of atomic numbers $\{Z_1, \dots, Z_N | Z_i \in \mathbb{N}\}$ and Cartesian coordinates $\{\vec{r}_1, \dots, \vec{r}_N | \vec{r}_i \in \mathbb{R}^3\}$, which describe the element types and positions of N atoms, as input. Information about the electronic wave function, which is necessary for an unambiguous description of a chemical system, is provided via two additional inputs: The total charge $Q \in \mathbb{Z}$ encodes the number of electrons (given by $Q + \sum_i Z_i$), whereas the total angular momentum is encoded as the number of unpaired electrons $S \in \mathbb{N}_0$. For example, a singlet state is indicated by $S = 0$, a doublet state by $S = 1$, and so on. The nuclear charges Z , total charge Q and spin state S are transformed to F -dimensional embeddings and combined to form initial atomic feature representations

$$\mathbf{x}^{(0)} = \mathbf{e}_Z + \mathbf{e}_Q + \mathbf{e}_S. \quad (1)$$

Here, the nuclear embeddings \mathbf{e}_Z contain information about the ground state electron configuration of each element and the

electronic embeddings \mathbf{e}_Q and \mathbf{e}_S contain delocalized information about the total charge and spin state, respectively. A chain of T interaction modules iteratively refines these representations through local and nonlocal interactions

$$\mathbf{x}_i^{(t)} = \mathbf{x}_i^{(t-1)} + \text{local}^{(t)}(\{\mathbf{x}_j^{(t-1)}, \vec{r}_{ij}\}_{j \in \mathcal{N}(i)}) + \text{nonlocal}^{(t)}(\{\mathbf{x}^{(t-1)}\}), \quad (2)$$

where $\mathcal{N}(i)$ contains all atom indices within a cutoff distance r_{cut} of atom i and $\vec{r}_{ij} = \vec{r}_j - \vec{r}_i$ is the relative position of atom j with respect to atom i . The local interaction functions are designed to resemble s, p, and d atomic orbitals (see Fig. 1d) and the model learns to encode different distance and angular information about the local environment of each atom with the different interaction functions (see Fig. 2a). The nonlocal interactions on the other hand model the delocalized electrons. The representations $\mathbf{x}^{(t)}$ at each stage are further refined through learned functions \mathcal{F}_t according to $\mathbf{y}_i^{(t)} = \mathcal{F}_t(\mathbf{x}^{(t)})$ and summed to the atomic descrip-

$$\mathbf{f} = \sum_{i=1}^T \mathbf{y}_i^{(t)}, \quad (3)$$

from which atomic energy contributions E_i are predicted with linear regression. The total potential energy is given by

$$E_{\text{pot}} = \sum_{i=1}^N E_i + E_{\text{rep}} + E_{\text{ele}} + E_{\text{vdw}}, \quad (4)$$

where E_{rep} , E_{ele} , and E_{vdw} are empirical corrections, which augment the energy prediction with physical knowledge about nuclear repulsion, electrostatics, and dispersion interactions. Energy-conserving forces $\vec{F}_i = -\partial E_{\text{pot}} / \partial \vec{r}_i$ can be obtained via automatic differentiation. A schematic depiction of the SpookyNet architecture is given in Fig. 3.

Electronic states. Most existing ML-FFs can only model structures with a consistent electronic state, e.g., neutral singlets. An exception are systems for which the electronic state can be inferred via structural cues, e.g., in the case of protonation/

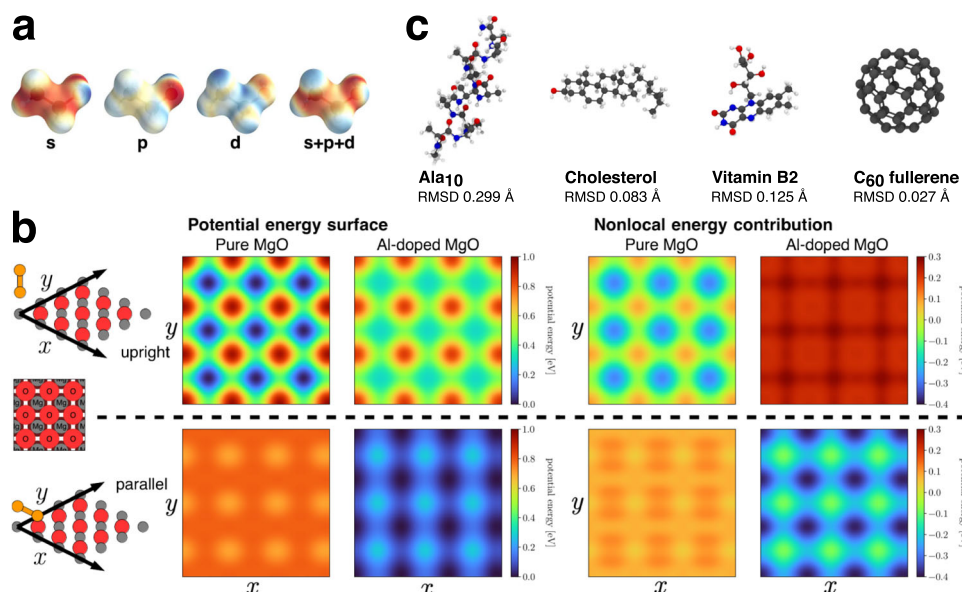


Fig. 2 Examples of chemical insights extracted by SpookyNet. **a** Visualization of the learned local chemical potential for ethanol (see methods). The individual contributions of s-, p-, and d-orbital-like interactions are shown (red: low energy, blue: high energy). **b** Potential energy surface scans obtained by moving an Au₂ dimer over an (Al-doped) MgO surface in different (upright/parallel) configurations (the positions of Mg and O atoms are shown for orientation). SpookyNet learns to distinguish between local and nonlocal contributions to the potential energy, allowing it to model changes of the potential energy surface when the crystal is doped with Al atoms far from the surface. **c** A model trained on small organic molecules learns general chemical principles that can be transferred to much larger structures outside the chemical space covered by the training data. Here, optimized geometries obtained from SpookyNet trained on the QM7-X database (opaque in color) are shown and compared with reference geometries obtained from ab initio calculations (transparent in gray). As indicated by the low root mean square deviations (RMSD), geometries predicted by SpookyNet are very similar to the reference.

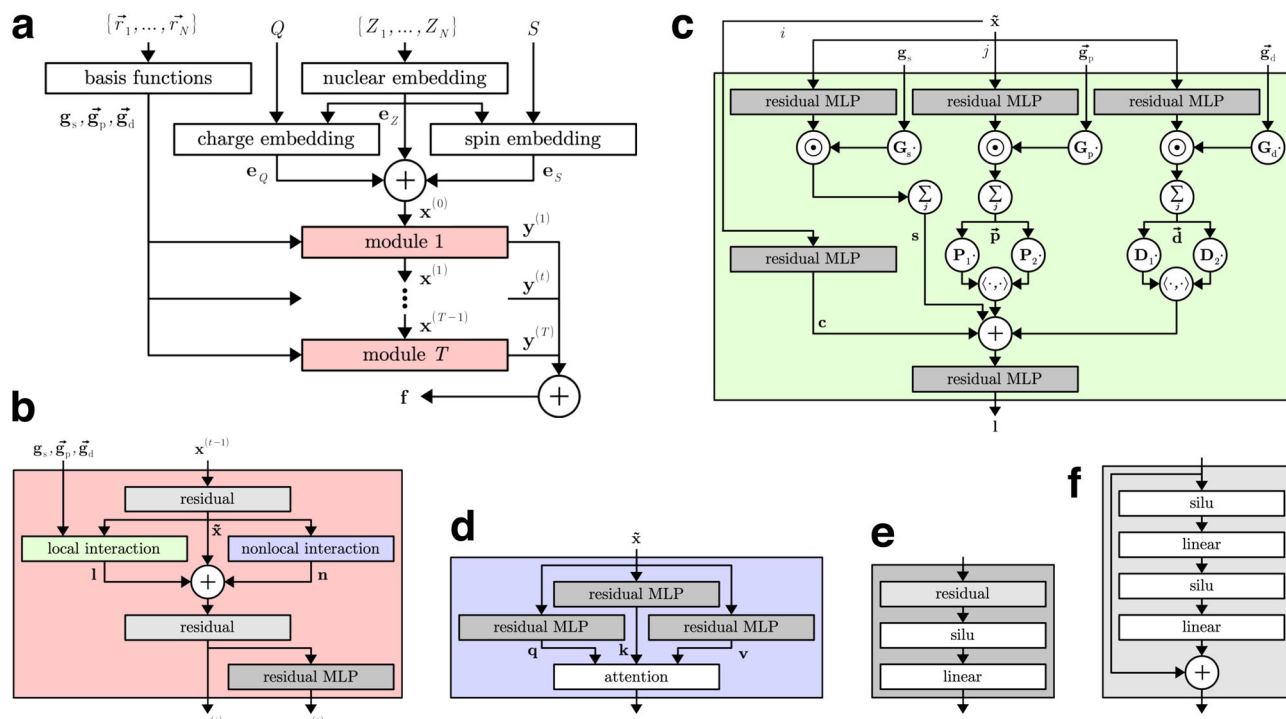


Fig. 3 Schematic depiction of the SpookyNet architecture with color-coded view of individual components. **a** Architecture overview, for details on the nuclear and electronic (charge/spin) embeddings and basis functions, refer to Eqs. (9), (10), and (13), respectively. **b** Interaction module, see Eq. (11). **c** Local interaction block, see Eq. (12). **d** Nonlocal interaction block, see Eq. (18). **e** Residual multilayer perceptron (MLP), see Eq. (8). **f** Pre-activation residual block, see Eq. (7).

deprotonation³⁷. In most cases, however, this is not possible, and ML-FFs that do not model electronic degrees of freedom are unable to capture the relevant physics. Here, this problem is solved by explicitly accounting for different electronic states (see

Eq. (1)). To illustrate their effects on potential energy surfaces, two exemplary systems are considered: Ag₃⁺/Ag₃⁻ and singlet/triplet CH₂, which can only be distinguished by their charge, respectively, their spin state. SpookyNet is able to faithfully

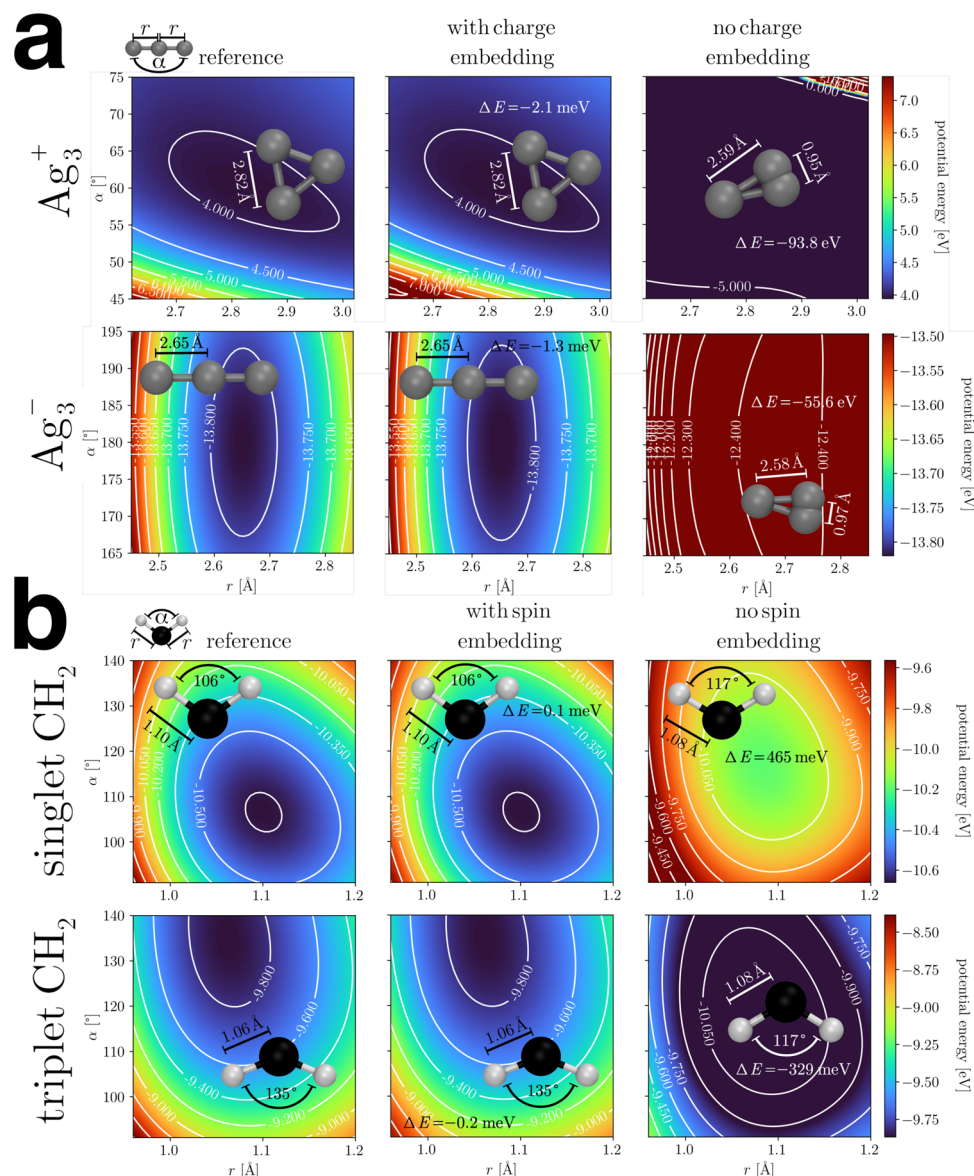


Fig. 4 Potential energy surfaces of **a** Ag_3^+ and Ag_3^- and **b** singlet and triplet CH_2 . The middle and right columns show the prediction of SpookyNet with and without charge/spin embedding, respectively, whereas the reference ground truth is shown in the left column. Minimum energy structures and prediction errors (ΔE) for the minimum energy are also shown.

reproduce the reference potential energy surface for all systems. When the charge/spin embeddings $\mathbf{e}_Q/\mathbf{e}_S$ (Eq. (1)) are removed, however, the model becomes unable to represent the true potential energy surface and its predictions are qualitatively different from the reference (see Fig. 4). As a consequence, wrong global minima are predicted when performing geometry optimizations with a model trained without the charge/spin embeddings, whereas they are virtually indistinguishable from the reference when the embeddings are used. Interestingly, even without a charge embedding, SpookyNet can predict different potential energy surfaces for $\text{Ag}_3^+/\text{Ag}_3^-$. This is because explicit point charge electrostatics are used in the energy prediction (see Eq. (4)) and the atomic partial charges are constrained such that the total molecular charge is conserved. However, such implicit information is insufficient to distinguish both charge states adequately. In the case of singlet/triplet CH_2 , there is no such implicit information and both systems appear identical to a model without electronic embeddings, i.e., it predicts the same energy surface for both systems, which neither reproduces the true singlet nor triplet reference.

Models with electronic embeddings even generalize to unknown electronic states. As an example, the QMspin database⁶⁸ is considered. It consists of $\sim 13\text{k}$ carbene structures with at most nine non-hydrogen atoms (C, N, O, F), which were optimized either in a singlet or triplet state. For each of these, both singlet and triplet energies computed at the MRCISD+Q-F12/cc-pVDZ-F12 level of theory are reported, giving a total of $\sim 26\text{k}$ energy-structure pairs in the database (see ref. ⁶⁹ for more details). For the lack of other models evaluated on this dataset, SpookyNet is compared to itself without electronic embeddings. This baseline model only reaches a mean absolute prediction error (MAE) of 444.6 meV for unknown spin states. As expected, the performance is drastically improved when the electronic embeddings are included, allowing SpookyNet to reach an MAE of 68.0 meV. Both models were trained on 20k points, used 1k samples as validation set, and were tested on the remaining data. An analysis of the local chemical potential (see methods) reveals that a model with electronic embeddings learns a feature-rich internal representation of molecules, which significantly differs

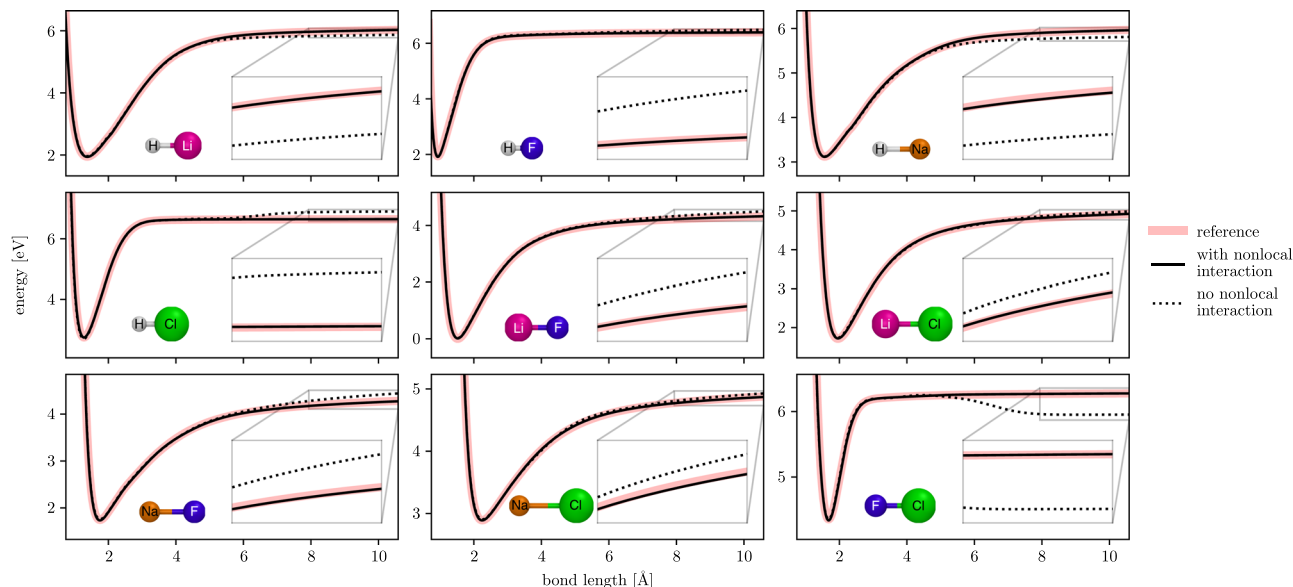


Fig. 5 Dissociation curves of different diatomic molecules predicted by SpookyNet with/without nonlocal interactions. Individual panels show the dissociation curves for different species. From top-left to bottom/right: HLi, HF, HNa, HCl, LiF, LiCl, NaF, NaCl, and FCl.

between singlet and triplet states (see Supplementary Fig. 3). In contrast, the local chemical potential learned by a model without electronic embeddings is almost uniform and (necessarily) identical between both states, suggesting that the electronic embeddings are crucial to extract the relevant chemistry from the data.

Nonlocal effects. For many chemical systems, local interactions are sufficient for an accurate description. However, there are also cases where a purely local picture breaks down. To demonstrate that nonlocal effects can play an important role even in simple systems, the dissociation curves of nine (neutral singlet) diatomic molecules made up of H, Li, F, Na, and Cl atoms are considered (Fig. 5). Once the bonding partners are separated by more than the chosen cutoff radius r_{cut} , models that rely only on local information will always predict the same energy contributions for atoms of the same element (by construction). However, since electrons are free to distribute unevenly across atoms, even when they are separated (e.g., due to differences in their electronegativity), energy contributions should always depend on the presence of other atoms in the structure. Consequently, it is difficult for models without nonlocal interactions to predict the correct asymptotic behavior for all systems simultaneously. As such, when the nonlocal interactions are removed from interaction modules (Eq. (2)), SpookyNet predicts an unphysical “step” for large interatomic separations, even when a large cutoff is used for the local interactions. In contrast, the reference is reproduced faithfully when nonlocal interactions are enabled. Note that such artifacts—occurring if nonlocal interactions are not modeled—are problematic e.g., when simulating reactions. Simply increasing the cutoff is no adequate solution to this problem, since it just shifts the artifact to larger separations. In these specific examples, even the inclusion of long-range corrections is insufficient to avoid artifacts in the asymptotic tails (analytical corrections for electrostatics and dispersion are enabled for both models), although they can help in some cases^{37,70}.

More complex nonlocal effects may arise for larger structures. For example, Ko et al. recently introduced four datasets for systems exhibiting nonlocal charge transfer effects⁴⁹. One of these systems consists of a diatomic Au cluster deposited on the surface of a periodic MgO(001) crystal ($\text{Au}_2\text{-MgO}$). In its minimum

energy configuration, the Au_2 cluster “stands upright” on the surface on top of an O atom. When some of the Mg atoms (far from the surface) are replaced by Al (see Fig. 1b), however, the Au_2 cluster prefers to “lie parallel” to the surface above two Mg atoms (the distance between the Au and Al atoms is above 10 Å). In other words, the presence of Al dopant atoms nonlocally modifies the electronic structure at the surface in such a way that a different Au_2 configuration becomes more stable. This effect can be quantified by scanning the potential energy surface of $\text{Au}_2\text{-MgO}$ by moving the Au_2 cluster above the surface in different configurations (see Fig. 2b). Upon introduction of dopant Al atoms, nonlocal energy contributions destabilize the upright configuration of Au_2 , particularly strongly above the positions of oxygen atoms. In contrast, the parallel configuration is lowered in energy, most strongly above positions of Mg atoms.

When applied to the $\text{Au}_2\text{-MgO}$ system, SpookyNet significantly improves upon the values reported for models without any treatment of nonlocal effects and also achieves lower prediction errors than 4G-BPNNs⁴⁹, which model nonlocal charge transfer via charge equilibration (see Table 1). For completeness, values for the other three systems introduced in ref. ⁴⁹ are also reported in Table 1, even though they could be modeled without including nonlocal interactions (as long as charge embeddings are used). For details on the number of training/validation data used for each dataset, refer to ref. ⁴⁹ (all models use the same settings).

Generalization across chemical and conformational space. For more typical ML-FF construction tasks where nonlocal effects are negligible and all molecules have consistent electronic states, SpookyNet improves upon the generalization capabilities of comparable ML-FFs. Here, the QM7-X database⁷¹ is considered as a challenging benchmark. This database was generated starting from ~7k molecular graphs with up to seven non-hydrogen atoms (C, N, O, S, Cl) drawn from the GDB13 chemical universe⁷². Structural and constitutional (stereo)isomers were sampled and optimized for each graph, leading to ~42k equilibrium structures. For each of these, an additional 100 non-equilibrium structures were generated by displacing atoms along linear combinations of normal modes corresponding to a temperature of 1500 K, which leads to ~4.2M structures in total. For each of these, QM7-X contains data for 42 physicochemical properties (see ref. ⁷¹ for a

Table 1 Root mean square errors (RMSEs) of energies (meV/atom), forces (meV Å⁻¹) and charges (me) for the datasets introduced in ref. 49.

		2G-BPNN	3G-BPNN	4G-BPNN	SpookyNet
C ₁₀ H ₂ /C ₁₀ H ₃ ⁺	Energy	1.619	2.045	1.194	0.364
	Forces	129.5	231.0	78.00	5.802
	Charges	—	20.08	6.577	0.117
Na _{8,9} Cl ₈ ⁺	Energy	1.692	2.042	0.481	0.135
	Forces	57.39	76.67	32.78	1.052
	Charges	—	20.80	15.83	0.111
Ag ₃ ^{+/-}	Energy	352.0	320.2	1.323	0.220
	Forces	1803	1913	31.69	26.64
	Charges	—	26.48	9.976	0.459
Au ₂ -MgO	Energy	2.287	—	0.219	0.107
	Forces	153.1	—	66.0	5.337
	Charges	—	—	5.698	1.013

The values for 2G-, 3G-, and 4G-BPNNs are taken from ref. 49. Best results in bold.

Table 2 Mean absolute errors for energy (meV) and force (meV Å⁻¹) predictions for the QM7-X⁷¹ dataset.

task		SchNet ³⁵	PaiNN ⁶¹	SpookyNet
Known molecules/	Energy	50.847	15.691	10.620 (0.403)
Unknown conformations	Forces	53.695	20.301	14.851 (0.430)
Unknown molecules/	Energy	51.275	17.594	13.151 (0.423)
Unknown conformations	Forces	62.770	24.161	17.326 (0.701)

Results for SpookyNet are averaged over four runs, the standard deviation between runs is given in brackets. Best results in bold.

complete list). For constructing ML-FFs, however, the properties E_{tot} and F_{tot} , which correspond to energies and forces computed at the PBE0+MBD^{73,74} level of theory, are the most relevant.

Because of the variety of molecules and the strongly distorted conformers contained in the QM7-X dataset, models need to be able to generalize across both chemical and conformational space to perform well. Here, two different settings are considered: In the more challenging task (unknown molecules/unknown conformations), a total of 10,100 entries corresponding to all structures sampled for 25 out of the original ~7k molecular graphs are reserved as test set and models are trained on the remainder of the data. In this setting, all structures in the test set correspond to unknown molecules, i.e., the model has to learn general principles of chemistry to perform well. As comparison, an easier task (known molecules/unknown conformations) is constructed by randomly choosing 10,100 entries as test set, so it is very likely that the training set contains at least some structures for all molecules contained in QM7-X (only unknown conformations need to be predicted). SpookyNet achieves lower prediction errors than both SchNet³⁵ and PaiNN⁶¹ for both tasks and is only marginally worse when predicting completely unknown molecules, suggesting that it successfully generalizes across chemical space (see Table 2). Interestingly, a SpookyNet model trained on QM7-X also generalizes to significantly larger chemical structures: Even though it was trained on structures with at most seven non-hydrogen atoms, it can be used e.g., for geometry optimizations of molecules like vitamin B2, cholesterol, or deca-alanine (see Fig. 2c). The optimized structures predicted by SpookyNet have low root mean square deviations (RMSD) from the ab initio reference geometries and are of higher quality than structures obtained from other models trained on QM7-X (see Supplementary Fig. 6). Remarkably, it even predicts the correct structures for fullerenes, although the QM7-X dataset contains no training data for any pure carbon structure. As an additional test, the trained model was also applied to structures from the conformer benchmark introduced in ref. 75, which contains structures with

up to 48 non-hydrogen atoms. Here, the relative energies between different conformers are predicted with sub-kcal accuracy, although absolute energies are systematically overestimated for large structures (see the conformer benchmark in the Supplementary Discussion for details).

Since the QM7-X dataset has only recently been published, the performance of SpookyNet is also benchmarked on the well-established MD17 dataset²¹. MD17 consists of structures, energies, and forces collected from ab initio MD simulations of small organic molecules at the PBE+TS^{76,77} level of theory. Prediction errors for several models published in the literature are summarized in Table 3 and compared to SpookyNet, which reaches lower prediction errors or closely matches the performance of other models for all tested molecules.

Discussion

The present work innovates by introducing SpookyNet, an MPNN for constructing ML-FFs, which models electronic degrees of freedom and nonlocal interactions using attention in a transformer architecture^{78,79}. SpookyNet includes physically motivated inductive biases that facilitate the extraction of chemical insight from data. For example, element embeddings in SpookyNet include the ground state electronic configuration, which encourages alchemically meaningful representations. An analytical short-range correction based on the Ziegler-Biersack-Littmark stopping potential⁸⁰ improves the description of nuclear repulsion, whereas long-range contributions to the potential energy are modeled with point charge electrostatics and an empirical dispersion correction, following previous works^{37,65,81–84}. These empirical augmentations allow SpookyNet to extrapolate beyond the data it was trained on based on physical knowledge from data.

While the general method to combine analytical long-range contributions with the predictions of a neural network is inspired by PhysNet³⁷, there are several differences between PhysNet and

Table 3 Mean absolute errors for energy (kcal mol⁻¹) and force (kcal mol⁻¹ Å⁻¹) predictions for the MD17 benchmark.

		sGDML ²²	SchNet ³⁵	PhysNet ³⁷	FCHL19 ²⁶	NequIP ⁶²	PaiNN ⁶¹	SpookyNet
Aspirin	Energy	0.19	0.37	0.230	0.182	—	0.159	0.151 (0.008)
	Forces	0.68	1.35	0.605	0.478	0.348	0.371	0.258 (0.034)
Ethanol	Energy	0.07	0.08	0.059	0.054	—	0.063	0.052 (0.001)
	Forces	0.33	0.39	0.160	0.136	0.208	0.230	0.094 (0.011)
Malondialdehyde	Energy	0.10	0.13	0.094	0.081	—	0.091	0.079 (0.002)
	Forces	0.41	0.66	0.319	0.245	0.337	0.319	0.167 (0.015)
Naphthalene	Energy	0.12	0.16	0.142	0.117	—	0.117	0.116 (0.001)
	Forces	0.11	0.58	0.310	0.151	0.096	0.083	0.089 (0.018)
Salicylic acid	Energy	0.12	0.20	0.126	0.114	—	0.114	0.114 (0.004)
	Forces	0.28	0.85	0.337	0.221	0.238	0.209	0.180 (0.040)
Toluene	Energy	0.10	0.12	0.100	0.098	—	0.097	0.094 (0.001)
	Forces	0.14	0.57	0.191	0.203	0.101	0.102	0.087 (0.014)
Uracil	Energy	0.11	0.14	0.108	0.104	—	0.104	0.105 (0.001)
	Forces	0.24	0.56	0.218	0.105	0.172	0.140	0.119 (0.021)

Results for SpookyNet are averaged over ten random splits, the standard deviation between runs is given in brackets. All models are trained on 1000 data points (separate models are used for each molecule), best results in bold.

the present approach. Notably, PhysNet does not include analytical short-range corrections and uses interaction functions that rely on purely radial information (instead of incorporating higher-order angular information). In addition, PhysNet cannot model different electronic states or nonlocal effects. In contrast, SpookyNet can predict different potential energy surfaces for the same molecule in different electronic states and is able to model nonlocal changes to the properties of materials such as MgO upon introduction of dopant atoms. Further, it successfully generalizes to structures well outside the chemical and conformational space covered by its training data and improves upon existing models in different quantum chemical benchmarks. The inductive biases incorporated into the architecture of SpookyNet encourage learning a chemically intuitive representation of molecular systems (see Fig. 2a). For example, the interaction functions learned by SpookyNet are designed to resemble atomic orbitals (see Fig. 1d). Obtaining such an understanding of how ML models⁸⁵, here SpookyNet, solve a prediction problem is crucially important in the sciences as a low test set error³² alone cannot rule out that a model may overfit or for example capitalize on various artifacts in data⁸⁶ or show “Clever Hans” effects⁸⁷

So far, most ML-FFs rely on nuclear charges and atomic coordinates as their only inputs and are thus unable to distinguish chemical systems with different electronic states. Further, they often rely on purely local information and break down when nonlocal effects cannot be neglected. The additions to MPNN architectures introduced in this work solve both of these issues, extending the applicability of ML-FFs to a much wider range of chemical systems than was previously possible and allow to model properties of quantum systems that have been neglected in many existing ML-FFs.

Remaining challenges in the construction of ML-FFs pertain to their successful application to large and heterogeneous condensed phase systems, such as proteins in aqueous solution. This is a demanding task, among others, due to the difficulty of performing ab initio calculations for such large systems, which is necessary to generate appropriate reference data. Although models trained on small molecules may generalize well to larger structures, it is not understood how to guarantee that all relevant regions of the potential energy surface, visited e.g. during a dynamics simulation, are well described. We conjecture that the inclusion of physically motivated inductive biases, which is a crucial ingredient in the SpookyNet architecture, may serve as a general design principle to improve the next generation of ML-FFs and tackle such problems.

Methods

Details on the neural network architecture. In the following, basic neural network building blocks and components of the SpookyNet architecture are described in detail (see Fig. 3 for a schematic depiction). A standard building block of most neural networks are linear layers, which take input features $\mathbf{x} \in \mathbb{R}^{n_{\text{in}}}$ and transform them according to

$$\text{linear}(\mathbf{x}) = \mathbf{W}\mathbf{x} + \mathbf{b}, \quad (5)$$

where $\mathbf{W} \in \mathbb{R}^{n_{\text{out}} \times n_{\text{in}}}$ and $\mathbf{b} \in \mathbb{R}^{n_{\text{out}}}$ are learnable weights and biases, and n_{in} and n_{out} are the dimensions of the input and output feature space, respectively (in this work, $n_{\text{in}} = n_{\text{out}}$ unless otherwise specified). Since Eq. (5) can only describe linear transformations, an activation function is required to learn nonlinear mappings between feature spaces. Here, a generalized SiLU (Sigmoid Linear Unit) activation function^{88,89} (also known as “swish”⁹⁰) given by

$$\text{silu}(x) = \frac{\alpha x}{1 + e^{-\beta x}} \quad (6)$$

is used. Depending on the values of α and β , Eq. (6) smoothly interpolates between a linear function and the popular ReLU (Rectified Linear Unit) activation⁹¹ (see Supplementary Fig. 4). Instead of choosing arbitrary fixed values, α and β are learnable parameters in this work. Whenever the notation $\text{silu}(\mathbf{x})$ is used, Eq. (6) is applied to the vector \mathbf{x} entry-wise and separate α and β parameters are used for each entry. Note that a smooth activation function is necessary for predicting potential energies, because the presence of kinks would introduce discontinuities in the atomic forces.

In theory, increasing the number of layers should never decrease the performance of a neural network, since in principle, superfluous layers could always learn the identity mapping. In practice, however, deeper neural networks become increasingly difficult to train due to the vanishing gradients problem⁹², which often degrades performance when too many layers are used. To combat this issue, it is common practice to introduce “shortcuts” into the architecture that skip one or several layers⁹³, creating a residual block⁹⁴. By inverting the order of linear layers and activation functions, it is even possible to train neural networks with several hundreds of layers⁹⁵. These “pre-activation” residual blocks transform input features \mathbf{x} according to

$$\text{residual}(\mathbf{x}) = \mathbf{x} + \text{linear}_2(\text{silu}_2(\text{linear}_1(\text{silu}_1(\mathbf{x})))) . \quad (7)$$

Throughout the SpookyNet architecture, small feedforward neural networks consisting of a residual block, followed by an activation and a linear output layer, are used as learnable feature transformations. For conciseness, such residual multilayer perceptrons (MLPs) are written as

$$\text{resmlp}(\mathbf{x}) = \text{linear}(\text{silu}(\text{residual}(\mathbf{x}))) . \quad (8)$$

The inputs to SpookyNet are transformed to initial atomic features (Eq. (1)) via embeddings. A nuclear embedding is used to map atomic numbers $Z \in \mathbb{N}$ to vectors $\mathbf{e}_Z \in \mathbb{R}^F$ given by

$$\mathbf{e}_Z = \mathbf{M}\mathbf{d}_Z + \tilde{\mathbf{e}}_Z . \quad (9)$$

Here, $\mathbf{M} \in \mathbb{R}^{F \times 20}$ is a parameter matrix that projects constant element descriptors $\mathbf{d}_Z \in \mathbb{R}^{20}$ to an F -dimensional feature space and $\tilde{\mathbf{e}}_Z \in \mathbb{R}^F$ are element-specific bias parameters. The descriptors \mathbf{d}_Z encode information about the ground state electronic configuration of each element (see Supplementary Table 3 for details). Note that the term $\tilde{\mathbf{e}}_Z$ by itself allows to learn arbitrary embeddings for different elements, but including $\mathbf{M}\mathbf{d}_Z$ provides an inductive bias to learn representations

that capture similarities between different elements, i.e., contain alchemical knowledge.

Electronic embeddings are used to map the total charge $Q \in \mathbb{Z}$ and number of unpaired electrons $S \in \mathbb{N}_0$ to vectors $\mathbf{e}_Q, \mathbf{e}_S \in \mathbb{R}^F$, which delocalize this information over all atoms via a mechanism similar to attention⁷⁸. The mapping is given by

$$\mathbf{q} = \text{linear}(\mathbf{e}_Z), \mathbf{k} = \begin{cases} \tilde{\mathbf{k}}_+ & \Psi \geq 0 \\ \tilde{\mathbf{k}}_- & \Psi < 0 \end{cases}, \mathbf{v} = \begin{cases} \tilde{\mathbf{v}}_+ & \Psi \geq 0 \\ \tilde{\mathbf{v}}_- & \Psi < 0 \end{cases},$$

$$a_i = \frac{\Psi \ln(1 + \exp(\mathbf{q}_i^T \mathbf{k} / \sqrt{F}))}{\sum_{j=1}^N \ln(1 + \exp(\mathbf{q}_j^T \mathbf{k} / \sqrt{F}))}, \mathbf{e}_\Psi = \text{resmlp}(a\mathbf{v}), \quad (10)$$

where $\tilde{\mathbf{k}}, \tilde{\mathbf{v}} \in \mathbb{R}^F$ are parameters and $\Psi = Q$ for charge embeddings, or $\Psi = S$ for spin embeddings (independent parameters are used for each type of electronic embedding). Separate parameters indicated by subscripts \pm are used for positive and negative total charge inputs Q (since S is always positive or zero, only the $+$ parameters are used for spin embeddings). Here, all bias terms in the resmlp transformation (Eq. (8)) are removed, such that when $a\mathbf{v} = \mathbf{0}$, the electronic embedding $\mathbf{e}_\Psi = \mathbf{0}$ as well. Note that $\sum_i a_i = \Psi$, i.e. the electronic information is distributed across atoms with weights proportional to the scaled dot product $\mathbf{q}_i^T \mathbf{k} / \sqrt{F}$.

The initial atomic representations $\mathbf{x}^{(0)}$ (Eq. (1)) are refined iteratively by a chain of T interaction modules according to

$$\begin{aligned} \tilde{\mathbf{x}} &= \text{residual}_1(\mathbf{x}^{(t-1)}), \\ \mathbf{x}^{(t)} &= \text{residual}_2(\tilde{\mathbf{x}} + \mathbf{I} + \mathbf{n}), \\ \mathbf{y}^{(t)} &= \text{resmlp}(\mathbf{x}^{(t)}). \end{aligned} \quad (11)$$

Here, $\tilde{\mathbf{x}} \in \mathbb{R}^F$ are temporary atomic features and $\mathbf{I}, \mathbf{n} \in \mathbb{R}^F$ represent interactions with other atoms. They are computed by local (Eq. (12)) and nonlocal (Eq. (18)) interaction blocks, respectively, which are described below. Each module t produces two outputs $\mathbf{x}^{(t)}, \mathbf{y}^{(t)} \in \mathbb{R}^F$, where $\mathbf{x}^{(t)}$ is the input to the next module in the chain and all $\mathbf{y}^{(t)}$ outputs are accumulated to the final atomic descriptors \mathbf{f} (Eq. (3)).

The features \mathbf{I} in Eq. (11) represent a local interaction of atoms within a cutoff radius $r_{\text{cut}} \in \mathbb{R}$ and introduce information about the atom positions $\vec{r} \in \mathbb{R}^3$. They are computed from the temporary features $\tilde{\mathbf{x}}$ (see Eq. (11)) according to

$$\begin{aligned} \mathbf{c} &= \text{resmlp}_c(\tilde{\mathbf{x}}), \\ \mathbf{s}_i &= \sum_{j \in \mathcal{N}(i)} \text{resmlp}_s(\tilde{\mathbf{x}}_j) \odot (\mathbf{G}_s \mathbf{g}_s(\vec{r}_{ij})), \\ \tilde{\mathbf{p}}_i &= \sum_{j \in \mathcal{N}(i)} \text{resmlp}_p(\tilde{\mathbf{x}}_j) \odot (\mathbf{G}_p \tilde{\mathbf{g}}_p(\vec{r}_{ij})), \\ \tilde{\mathbf{d}}_i &= \sum_{j \in \mathcal{N}(i)} \text{resmlp}_d(\tilde{\mathbf{x}}_j) \odot (\mathbf{G}_d \tilde{\mathbf{g}}_d(\vec{r}_{ij})), \\ \mathbf{I} &= \text{resmlp}_1(\mathbf{c} + \mathbf{s} + (\mathbf{P}_1 \tilde{\mathbf{p}}, \mathbf{P}_2 \tilde{\mathbf{p}}) + (\mathbf{D}_1 \tilde{\mathbf{d}}, \mathbf{D}_2 \tilde{\mathbf{d}})), \end{aligned} \quad (12)$$

where, $\mathcal{N}(i)$ is the set of all indices $j \neq i$ for which $\|\vec{r}_{ij}\| < r_{\text{cut}}$ (with $\vec{r}_{ij} = \vec{r}_j - \vec{r}_i$). The parameter matrices $\mathbf{G}_s, \mathbf{G}_p, \mathbf{G}_d \in \mathbb{R}^{F \times K}$ are used to construct feature-wise interaction functions as linear combinations of basis functions $\mathbf{g}_s \in \mathbb{R}^K$, $\tilde{\mathbf{g}}_p \in \mathbb{R}^{K \times 3}$, and $\tilde{\mathbf{g}}_d \in \mathbb{R}^{K \times 5}$ (see Eq. (13)), which have the same rotational symmetries as s-, p-, and d-orbitals. The features $\mathbf{s} \in \mathbb{R}^F$, $\tilde{\mathbf{p}} \in \mathbb{R}^{F \times 3}$, and $\tilde{\mathbf{d}} \in \mathbb{R}^{F \times 5}$ encode the arrangement of neighboring atoms within the cutoff radius and $\mathbf{c} \in \mathbb{R}^F$ describes the central atom in each neighborhood. Here, \mathbf{s} stores purely radial information, whereas $\tilde{\mathbf{p}}$ and $\tilde{\mathbf{d}}$ allow to resolve angular information in a computationally efficient manner (see Supplementary Discussion for details). The parameter matrices $\mathbf{P}_1, \mathbf{P}_2, \mathbf{D}_1, \mathbf{D}_2 \in \mathbb{R}^{F \times F}$ are used to compute two independent linear projections for each of the rotationally equivariant features $\tilde{\mathbf{p}}$ and $\tilde{\mathbf{d}}$, from which rotationally invariant features are obtained via a scalar product $\langle \cdot, \cdot \rangle$. The basis functions (see Fig. 6) are given by

$$\begin{aligned} \mathbf{g}_s(\vec{r}) &= \begin{bmatrix} 0\mathbf{g}_0^0 \\ \vdots \\ K-1\mathbf{g}_0^0 \end{bmatrix}, \\ \tilde{\mathbf{g}}_p(\vec{r}) &= \begin{bmatrix} 0\mathbf{g}_1^{-1} & 0\mathbf{g}_1^0 & 0\mathbf{g}_1^1 \\ \vdots & \vdots & \vdots \\ K-1\mathbf{g}_1^{-1} & K-1\mathbf{g}_1^0 & K-1\mathbf{g}_1^1 \end{bmatrix}, \\ \tilde{\mathbf{g}}_d(\vec{r}) &= \begin{bmatrix} 0\mathbf{g}_2^{-2} & 0\mathbf{g}_2^{-1} & 0\mathbf{g}_2^0 & 0\mathbf{g}_2^1 & 0\mathbf{g}_2^2 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ K-1\mathbf{g}_2^{-2} & K-1\mathbf{g}_2^{-1} & K-1\mathbf{g}_2^0 & K-1\mathbf{g}_2^1 & K-1\mathbf{g}_2^2 \end{bmatrix}, \\ k\mathbf{g}_l^m &= \rho_k(\|\vec{r}\|) \cdot Y_l^m(\vec{r}), \end{aligned} \quad (13)$$

where the radial component ρ_k is

$$\rho_k(r) = b_{k,K-1}(\exp(-\gamma r)) \cdot f_{\text{cut}}(r) \quad (14)$$

and

$$b_{k,K-1}(x) = \binom{K-1}{k} x^k (1-x)^{K-1-k} \quad (15)$$

are Bernstein polynomials ($k = 0, \dots, K-1$). The hyper-parameter K determines the total number of radial components (and the degree of the Bernstein polynomials). For $K \rightarrow \infty$, linear combinations of $b_{k,K-1}(x)$ can approximate any continuous function on the interval $[0, 1]$ uniformly³⁸. An exponential function $\exp(-\gamma r)$ maps distances r from $[0, \infty)$ to the interval $(0, 1]$, where $\gamma \in \mathbb{R}_{>0}$ is a radial decay parameter shared across all basis functions (for computational efficiency). A desirable side effect of this mapping is that the rate at which learned interaction functions can vary decreases with increasing r , which introduces a chemically meaningful inductive bias (electronic wave functions also decay exponentially with increasing distance from a nucleus)^{37,55}. The cutoff function

$$f_{\text{cut}}(r) = \begin{cases} \exp\left(-\frac{r^2}{(r_{\text{cut}}-r)(r_{\text{cut}}+r)}\right) & r < r_{\text{cut}} \\ 0 & r \geq r_{\text{cut}} \end{cases} \quad (16)$$

ensures that basis functions smoothly decay to zero for $r \geq r_{\text{cut}}$, so that no discontinuities are introduced when atoms enter or leave the cutoff radius. The angular component $Y_l^m(\vec{r})$ in Eq. (13) is given by

$$\begin{aligned} Y_l^m(\vec{r}) &= \begin{cases} \sqrt{2} \cdot \Pi_l^m(z) \cdot A_{|m|}(x, y) & m < 0 \\ \Pi_l^0(z) & m = 0 \\ \sqrt{2} \cdot \Pi_l^m(z) \cdot B_m(x, y) & m > 0 \end{cases}, \\ A_m(x, y) &= \sum_{p=0}^m \binom{m}{p} x^p y^{m-p} \sin\left(\frac{\pi}{2}(m-p)\right), \\ B_m(x, y) &= \sum_{p=0}^m \binom{m}{p} x^p y^{m-p} \cos\left(\frac{\pi}{2}(m-p)\right), \\ \Pi_l^m(z) &= \sqrt{\frac{(l-m)!}{(l+m)!}} \sum_{p=0}^{\lfloor (l-m)/2 \rfloor} c_{plm} r^{2p-1} z^{l-2p-m}, \\ c_{plm} &= \frac{(-1)^p}{2^l} \binom{m}{p} \binom{2l-2p}{l} \frac{(l-2p)!}{(l-2p-m)!}, \end{aligned} \quad (17)$$

where $\vec{r} = [x \ y \ z]^T$ and $r = \|\vec{r}\|$. Note that the Y_l^m in Eq. (17) omit the normalization constant $\sqrt{(4\pi)/(2l+1)}$, but are otherwise identical to the standard (real) spherical harmonics.

Although locality is a valid assumption for many chemical systems¹¹, electrons may also be delocalized across multiple distant atoms. Starting from the temporary features $\tilde{\mathbf{x}}$ (see Eq. (11)), such nonlocal interactions are modeled via self-attention⁷⁸ as

$$\begin{aligned} \mathbf{q}_i &= \text{resmlp}_q(\tilde{\mathbf{x}}_i), & \mathbf{Q} &= [\mathbf{q}_1 \cdots \mathbf{q}_N]^T, \\ \mathbf{k}_i &= \text{resmlp}_k(\tilde{\mathbf{x}}_i), & \mathbf{K} &= [\mathbf{k}_1 \cdots \mathbf{k}_N]^T, \\ \mathbf{v}_i &= \text{resmlp}_v(\tilde{\mathbf{x}}_i), & \mathbf{V} &= [\mathbf{v}_1 \cdots \mathbf{v}_N]^T, \\ \mathbf{N} &= \text{attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}), & \mathbf{N} &= [\mathbf{n}_1 \cdots \mathbf{n}_N]^T, \end{aligned} \quad (18)$$

where the features \mathbf{n} in Eq. (11) are the (transposed) rows of the matrix $\mathbf{N} \in \mathbb{R}^{N \times F}$. The idea of attention is inspired by retrieval systems⁹⁶, where a query is mapped against keys to retrieve the best-matched corresponding values from a database. Standard attention is computed as

$$\begin{aligned} \mathbf{A} &= \exp\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{F}}\right), & \mathbf{D} &= \text{diag}(\mathbf{A}\mathbf{1}_N), \\ \text{attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) &= \mathbf{D}^{-1}\mathbf{A}\mathbf{V}, \end{aligned} \quad (19)$$

where $\mathbf{Q}, \mathbf{K}, \mathbf{V} \in \mathbb{R}^{N \times F}$ are queries, keys, and values, $\mathbf{1}_N$ is the all-ones vector of length N , and $\text{diag}(\cdot)$ is a diagonal matrix with the input vector as the diagonal. Unfortunately, computing attention with Eq. (19) has a time and space complexity of $\mathcal{O}(N^2F)$ and $\mathcal{O}(N^2 + NF)$ ⁷⁹, respectively, because the attention matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ has to be stored explicitly. Since quadratic scaling with the number of atoms N is problematic for large chemical systems, the FAVOR+ (Fast Attention Via positive Orthogonal Random features) approximation⁷⁹ is used instead:

$$\begin{aligned} \hat{\mathbf{Q}} &= [\phi(\mathbf{q}_1) \cdots \phi(\mathbf{q}_N)]^T, & \hat{\mathbf{K}} &= [\phi(\mathbf{k}_1) \cdots \phi(\mathbf{k}_N)]^T, \\ \hat{\mathbf{D}} &= \text{diag}(\hat{\mathbf{Q}}(\hat{\mathbf{K}}^T \mathbf{1}_N)), & & \\ \text{attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) &= \hat{\mathbf{D}}^{-1}(\hat{\mathbf{Q}}\hat{\mathbf{K}}^T \mathbf{V}). \end{aligned} \quad (20)$$

Here $\phi: \mathbb{R}^F \rightarrow \mathbb{R}_{>0}^f$ is a mapping designed to approximate the softmax kernel via f random features, see ref. ⁷⁹ for details (here, $f = F$ for simplicity). The time and space complexities for computing attention with Eq. (20) are $\mathcal{O}(NFf)$ and $\mathcal{O}(NF + Nf + Ff)$ ⁷⁹, i.e., both scale linearly with the number of atoms N . To make




Fig. 6 Visualization of basis functions. All basis functions ${}_k g_l^m$ with $K = 4$ (see Eq. (13)) with different radial and angular components ρ_k (Eq. (14)) and Y_l^m (Eq. (17)) are shown.

the evaluation of SpookyNet deterministic, the random features of the mapping ϕ are drawn only once at initialization and kept fixed afterwards (instead of redrawing them for each evaluation).

Once all interaction modules are evaluated, atomic energy contributions E_i are predicted from the atomic descriptors \mathbf{f}_i via linear regression

$$E_i = \mathbf{w}_E^T \mathbf{f}_i + \tilde{E}_{Z_i}, \quad (21)$$

and combined to obtain the total potential energy (see Eq. (4)). Here, $\mathbf{w}_E \in \mathbb{R}^F$ are the regression weights and $\tilde{E}_Z \in \mathbb{R}$ are element-dependent energy biases.

The nuclear repulsion term E_{rep} in Eq. (4) is based on the Ziegler-Biersack-Littmark stopping potential⁸⁰ and given by

$$E_{\text{rep}} = k_e \sum_{i > j \in \mathcal{N}(i)} \frac{Z_i Z_j}{r_{ij}} f_{\text{cut}}(r_{ij}) \cdot \left(\sum_{k=1}^4 c_k e^{-a_k r_{ij} (Z_i^p + Z_j^p)/d} \right). \quad (22)$$

Here, k_e is the Coulomb constant and a_k, c_k, p , and d are parameters (see Eqs. (12) and (16) for the definitions of $\mathcal{N}(i)$ and f_{cut}). Long-range electrostatic interactions are modeled as

$$E_{\text{ele}} = k_e \sum_i \sum_{j > i} q_i q_j \left(\frac{f_{\text{switch}}(r_{ij})}{\sqrt{r_{ij}^2 + 1}} + \frac{1 - f_{\text{switch}}(r_{ij})}{r} \right), \quad (23)$$

where q_i are atomic partial charges predicted from the atomic features \mathbf{f}_i according to

$$q_i = \mathbf{w}_q^T \mathbf{f}_i + \tilde{q}_{Z_i} + \frac{1}{N} \left[Q - \sum_{j=1}^N (\mathbf{w}_q^T \mathbf{f}_j + \tilde{q}_{Z_j}) \right]. \quad (24)$$

Here, $\mathbf{w}_q \in \mathbb{R}^F$ and $\tilde{q}_Z \in \mathbb{R}$ are regression weights and element-dependent biases, respectively. The second half of the equation ensures that $\sum_i q_i = Q$, i.e., the total charge is conserved. Standard Ewald summation⁹⁷ can be used to evaluate E_{ele} when periodic boundary conditions are used. Note that Eq. (23) smoothly interpolates between the correct r^{-1} behavior of Coulomb's law at large distances ($r > r_{\text{off}}$) and a damped ($r_{ij}^2 + 1$)^{-1/2} dependence at short distances ($r < r_{\text{on}}$) via a smooth switching function f_{switch} given by

$$\sigma(x) = \begin{cases} \exp(-\frac{1}{x}) & x > 0 \\ 0 & x \leq 0 \end{cases}, \quad (25)$$

$$f_{\text{switch}}(r) = \frac{\sigma\left(1 - \frac{r - r_{\text{on}}}{r_{\text{off}} - r_{\text{on}}}\right)}{\sigma\left(1 - \frac{r - r_{\text{on}}}{r_{\text{off}} - r_{\text{on}}}\right) + \sigma\left(\frac{r - r_{\text{on}}}{r_{\text{off}} - r_{\text{on}}}\right)}.$$

For simplicity, $r_{\text{on}} = \frac{1}{4} r_{\text{cut}}$ and $r_{\text{off}} = \frac{3}{4} r_{\text{cut}}$, i.e. the switching interval is automatically adjusted depending on the chosen cutoff radius r_{cut} (see Eq. (16)). It is also possible to construct dipole moments $\vec{\mu}$ from the partial charges according to

$$\vec{\mu} = \sum_{i=1}^N q_i \vec{r}_i, \quad (26)$$

which can be useful for calculating infrared spectra from MD simulations and for fitting q_i to ab initio reference data without imposing arbitrary charge decomposition schemes⁹⁸. Long-range dispersion interactions are modeled via the term E_{vdw} . Analytical van der Waals corrections are an active area of research and many different methods, for example the Tkatchenko-Scheffler correction⁷⁷, or many body dispersion⁷⁴, have been proposed⁹⁹. In this work, the two-body term of

the D4 dispersion correction¹⁰⁰ is used for its simplicity and computational efficiency:

$$E_{\text{vdw}} = - \sum_i \sum_{j > i} \sum_{n=6,8} s_n \frac{C_{ij}^{(n)}}{r_{ij}^n} f_{\text{damp}}^{(n)}(r_{ij}). \quad (27)$$

Here s_n are scaling parameters, $f_{\text{damp}}^{(n)}$ is a damping function, and $C_{ij}^{(n)}$ are pairwise dispersion coefficients. They are obtained by interpolating between tabulated reference values based on a (geometry-dependent) fractional coordination number and an atomic partial charge q_i . In the standard D4 scheme, the partial charges are obtained via a charge equilibration scheme¹⁰⁰, in this work, however, the q_i from Eq. (24) are used instead. Note that the D4 method was developed mainly to correct for the lack of dispersion in density functionals, so typically, some of its parameters are adapted to the functional the correction is applied to (optimal values for each functional are determined by fitting to high-quality electronic reference data)¹⁰⁰. In this work, all D4 parameters that vary between different functionals are treated as learnable parameters when SpookyNet is trained, i.e., they are automatically adapted to the reference data. Since Eq. (24) (instead of charge equilibration) is used to determine the partial charges, an additional learnable parameter s_q is introduced to scale the tabulated reference charges used to determine dispersion coefficients $C_{ij}^{(n)}$. For further details on the implementation of the D4 method, the reader is referred to ref. ¹⁰⁰.

Training and hyperparameters. All SpookyNet models in this work use $T = 6$ interaction modules, $F = 128$ features, and a cutoff radius $r_{\text{cut}} = 10 a_0$ ($\approx 5.29177 \text{ \AA}$), unless otherwise specified. Weights are initialized as random (semi-) orthogonal matrices with entries scaled according to the Glorot initialization scheme⁹². An exception are the weights of the second linear layer in residual blocks (linear₂ in Eq. (7)) and the matrix \mathbf{M} used in nuclear embeddings (Eq. (9)), which are initialized with zeros. All bias terms and the $\tilde{\mathbf{k}}$ and $\tilde{\mathbf{v}}$ parameters in the electronic embedding (Eq. (10)) are also initialized with zeros. The parameters for the activation function (Eq. (6)) start as $\alpha = 1.0$ and $\beta = 1.702$, following the recommendations given in Ref. ⁸⁸. The radial decay parameter γ used in Eq. (14) is initialized to $\frac{1}{2} a_0^{-1}$ and constrained to positive values. The parameters of the empirical nuclear repulsion term (Eq. (22)) start from the literature values of the ZBL potential⁸⁰ and are constrained to positive values (coefficients c_k are further constrained such that $\sum_k c_k = 1$ to guarantee the correct asymptotic behavior for short distances). Parameters of the dispersion correction (Eq. (27)) start from the values recommended for Hartree-Fock calculations¹⁰⁰ and the charge scaling parameter s_q is initialized to 1 (and constrained to remain positive).

The parameters are trained by minimizing a loss function with mini-batch gradient descent using the AMSGrad optimizer¹⁰¹ with the recommended default momentum hyperparameters and an initial learning rate of 10^{-3} . During training, an exponential moving average of all model parameters is kept using a smoothing factor of 0.999. Every 1k training steps, a model using the averaged parameters is evaluated on the validation set and the learning rate is decayed by a factor of 0.5 whenever the validation loss does not decrease for 25 consecutive evaluations. Training is stopped when the learning rate drops below 10^{-5} and the model that performed best on the validation set is selected. The loss function is given by

$$\mathcal{L} = \alpha_E \mathcal{L}_E + \alpha_F \mathcal{L}_F + \alpha_\mu \mathcal{L}_\mu, \quad (28)$$

where \mathcal{L}_E , \mathcal{L}_F , and \mathcal{L}_μ are separate loss terms for energies, forces and dipole moments and α_E , α_F , and α_μ corresponding weighting hyperparameters that determine the relative influence of each term to the total loss. The energy loss is

given by

$$\mathcal{L}_E = \sqrt{\frac{1}{B} \sum_{b=1}^B (E_{\text{pot},b} - E_{\text{pot},b}^{\text{ref}})^2}, \quad (29)$$

where B is the number of structures in the mini-batch, $E_{\text{pot},b}$ the predicted potential energy (Eq. (4)) for structure b and $E_{\text{pot},b}^{\text{ref}}$ the corresponding reference energy. The batch size B is chosen depending on the available training data: When training sets contain 1k structures or less, $B = 1$, for 10k structures or less, $B = 10$, and for more than 10k structures, $B = 100$. The force loss is given by

$$\mathcal{L}_F = \sqrt{\frac{1}{B} \sum_{b=1}^B \left(\frac{1}{N_b} \sum_{i=1}^{N_b} \left\| -\frac{\partial E_{\text{pot},b}}{\partial \vec{r}_{i,b}} - \vec{F}_{i,b}^{\text{ref}} \right\|^2 \right)}, \quad (30)$$

where N_b is the number of atoms in structure b and $\vec{F}_{i,b}^{\text{ref}}$ the reference force acting on atom i in structure b . The dipole loss

$$\mathcal{L}_\mu = \sqrt{\frac{1}{B} \sum_{b=1}^B \left\| \left(\sum_{i=1}^{N_b} q_{i,b} \vec{r}_{i,b} \right) - \vec{\mu}_b^{\text{ref}} \right\|^2}, \quad (31)$$

allows to learn partial charges (Eq. (24)) from reference dipole moments $\vec{\mu}_b^{\text{ref}}$, which are, in contrast to arbitrary charge decompositions, true quantum mechanical observables⁹⁸. Partial charges learned in this way are typically similar in magnitude to Hirshfeld charges¹⁰² and follow similar overall trends (see Supplementary Fig. 5). Note that for charged molecules, the dipole moment is dependent on the origin of the coordinate system, so consistent conventions must be used. For some datasets or applications, however, reference partial charges $q_{i,b}^{\text{ref}}$ obtained from a charge decomposition scheme might be preferred (or the only data available). In this case, the term $\alpha_\mu \mathcal{L}_\mu$ in Eq. (28) is replaced by $\alpha_q \mathcal{L}_q$ with

$$\mathcal{L}_q = \sqrt{\frac{1}{B} \sum_{b=1}^B \left(\frac{1}{N_b} \sum_{i=1}^{N_b} (q_{i,b} - q_{i,b}^{\text{ref}})^2 \right)}. \quad (32)$$

For simplicity, the relative loss weights are set to $\alpha_E = \alpha_F = \alpha_{\mu,q} = 1$ in this work, with the exception of the MD17 and QM7-X datasets, for which $\alpha_F = 100$ is used following previous work³⁷. Both energy and force prediction errors are significantly reduced when the force weight is increased (see Supplementary Table 4). Note that the relative weight of loss terms also depends on the chosen unit system (atomic units are used here). For datasets that lack the reference data necessary for computing any of the given loss terms (Eqs. (29)–(32)), the corresponding weight is set to zero. In addition, whenever no reference data (neither dipole moments nor reference partial charges) are available to fit partial charges, both E_{ele} and E_{vdw} are omitted when predicting the potential energy E_{pot} (see Eq. (4)).

For the “unknown molecules/unknown conformations” task reported in Table 2, the 25 entries with the following ID numbers (idmol field in the QM7-X file format) were used as a test set: 1771, 1805, 1824, 2020, 2085, 2117, 3019, 3108, 3190, 3217, 3257, 3329, 3531, 4010, 4181, 4319, 4713, 5174, 5370, 5580, 5891, 6315, 6583, 6809, 7020. In addition to energies and forces, SpookyNet uses dipole moments (property D in the QM7-X dataset) to fit atomic partial charges.

Computing and visualizing local chemical potentials and nonlocal contributions.

To compute the local chemical potentials shown in Fig. 2a and Supplementary Fig. 3, a similar approach as that described in Ref. 53 is followed. To compute the local chemical potential $\Omega_A^M(\vec{r})$ of a molecule M for an atom of type A (here, hydrogen is used), the idea is to introduce a probe atom of type A at position \vec{r} and let it interact with all atoms of M , but not vice versa. In other words, the prediction for M is unperturbed, but the probe atom “feels” the presence of M . Then, the predicted energy contribution of the probe atom is interpreted as its local chemical potential $\Omega_A^M(\vec{r})$. This is achieved as follows: First, the electronic embeddings (Eq. (10)) for all N atoms in M are computed as if the probe atom was not present. Then, the embeddings for the probe atom are computed as if it was part of a larger molecule with $N + 1$ atoms. Similarly, the contributions of local interactions (Eq. (12)) and nonlocal interactions (Eq. (18)) to the features of the probe atom are computed by pretending it is part of a molecule with $N + 1$ atoms, whereas all contributions to the features of the N atoms in molecule M are computed without the presence of the probe atom. For visualization, all chemical potentials are projected onto the $\sum_{i=1}^N \|\vec{r} - \vec{r}_i\|^{-2} = 1 a_0^{-2}$ isosurface, where the sum runs over the positions \vec{r}_i of all atoms i in M .

To obtain the individual contributions for s-, p-, and d-orbital-like interactions shown in Fig. 2, different terms for the computation of \mathbf{I} in Eq. (12) are set to zero. For the s-orbital-like contribution, both \vec{p} and \vec{d} are set to zero. For the p-orbital-like contribution, only \vec{d} is set to zero, and the s-orbital-like contribution is subtracted from the result. Similarly, for the d-orbital-like contribution, the model is evaluated normally and the result from setting only \vec{d} to zero is subtracted.

The nonlocal contributions to the potential energy surface shown in Fig. 2b are obtained by first evaluating the model normally and then subtracting the predictions obtained when setting \mathbf{n} in Eq. (11) to zero.

SchNet and PaiNN training. The SchNet and PaiNN models for the QM7-X experiments use $F = 128$ features, as well as $T = 6$ and $T = 3$ interactions, respectively. Both employ 20 Gaussian radial basis function up to a cutoff of 5 \AA . They were trained with the Adam optimizer¹⁰³ at a learning rate of 10^{-4} and a batch size of 10.

Data generation. For demonstrating the ability of SpookyNet to model different electronic states and nonlocal interactions, energies, forces, and dipoles for three new datasets were computed at the semi-empirical GFN2-xTB level of theory¹⁰⁴. Both the $\text{Ag}_3^+/\text{Ag}_3^-$ (see Fig. 4a) and the singlet/triplet CH_2 (see Fig. 4b) datasets were computed by sampling 550 structures around the minima of both electronic states with normal mode sampling²³ at 1000 K. Then, each sampled structure was re-computed in the other electronic state (e.g., all structures sampled for Ag_3^+ were re-computed with a negative charge), leading to a total of 2200 structures for each dataset (models were trained on a subset of 1000 randomly sampled structures).

The dataset for Fig. 5 was computed by performing bond scans for all nine shown diatomic molecules using 1k points spaced evenly between 1.5 and $20 a_0$, leading to a total of 9k structures. Models were trained on all data with an increased cutoff $r_{\text{cut}} = 18 a_0$ to demonstrate that a model without nonlocal interactions is unable to fit the data, even when it is allowed to overfit and use a large cutoff.

The reference geometries shown in Fig. 2c and Supplementary Fig. 6 were computed at the PBE0+MBD^{73,74} level using the FHI-aims code^{105,106}. All calculations used “tight” settings for basis functions and integration grids and the same convergence criteria as those applied for computing the QM7-X dataset⁷¹.

Data availability

The singlet/triplet carbene and $\text{Ag}_3^+/\text{Ag}_3^-$ datasets generated for this work are available without restrictions from <https://doi.org/10.5281/zenodo.5115732>¹⁰⁸. All other datasets used in this work are publicly available from ref. 109 (see completeness test in the Supplementary Discussion), <http://www.sgdm.org> (MD17), ref. 110 (QM7-X), ref. 111 (datasets used in Table 1), and ref. 68 (QMspin).

Code availability

A reference implementation of SpookyNet using PyTorch¹¹² is available from <https://github.com/OUUnke/SpookyNet>.

Received: 12 July 2021; Accepted: 16 November 2021;

Published online: 14 December 2021

References

1. Warshel, A. Molecular dynamics simulations of biological reactions. *Acc. Chem. Res.* **35**, 385 (2002).
2. Karplus, M. & McCammon, J. A. Molecular dynamics simulations of biomolecules. *Nat. Struct. Biol.* **9**, 646 (2002).
3. Dirac, P. A. M. Quantum mechanics of many-electron systems. *Proc. R. Soc. Lond. A* **123**, 714 (1929).
4. Dykstra, C., Frenking, G., Kim, K., & Scuseria, G. (Eds.). Theory and applications of computational chemistry: the first forty years. (Elsevier 2005).
5. Unke, O. T., Koner, D., Patra, S., Käser, S. & Meuwly, M. High-dimensional potential energy surfaces for molecular simulations: from empiricism to machine learning. *Mach. Learn. Sci. Technol.* **1**, 013001 (2020).
6. Bartók, A. P., Payne, M. C., Kondor, R. & Csányi, G. Gaussian approximation potentials: the accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.* **104**, 136403 (2010).
7. Rupp, M., Tkatchenko, A., Müller, K.-R. & Von Lilienfeld, O. A. Fast and accurate modeling of molecular atomization energies with machine learning. *Phys. Rev. Lett.* **108**, 058301 (2012).
8. Bartók, A. P. et al. Machine learning unifies the modeling of materials and molecules. *Sci. Adv.* **3**, e1701816 (2017).
9. Gastegger, M., Schwiedrzik, L., Bittermann, M., Berzsényi, F. & Marquetand, P. wACSF—weighted atom-centered symmetry functions as descriptors in machine learning potentials. *J. Chem. Phys.* **148**, 241709 (2018).
10. Schütt, K., Gastegger, M., Tkatchenko, A., Müller, K.-R. & Maurer, R. J. Unifying machine learning and quantum chemistry with a deep neural network for molecular wavefunctions. *Nat. Commun.* **10**, 5024 (2019).
11. Unke, O. T. et al. Machine learning force fields. *Chem. Rev.* **121**, 10142–10186 (2021).
12. von Lilienfeld, O. A., Müller, K.-R. & Tkatchenko, A. Exploring chemical compound space with quantum-based machine learning. *Nat. Rev. Chem.* **4**, 347 (2020).
13. Noé, F., Tkatchenko, A., Müller, K.-R. & Clementi, C. Machine learning for molecular simulation. *Annu. Rev. Phys. Chem.* **71**, 361 (2020).
14. Glielmo, A. et al. Unsupervised learning methods for molecular simulation data. *Chem. Rev.* **121**, 9722–9758 (2021).

15. Keith, J. A. et al. Combining machine learning and computational chemistry for predictive insights into chemical systems. *Chem. Rev.* **121**, 9816–9872 (2021).
16. Saucedo, H. E., Gastegger, M., Chmiela, S., Müller, K.-R. & Tkatchenko, A. Molecular force fields with gradient-domain machine learning (GDML): Comparison and synergies with classical force fields. *J. Chem. Phys.* **153**, 124109 (2020).
17. Behler, J. & Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **98**, 146401 (2007).
18. Behler, J. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J. Chem. Phys.* **134**, 074106 (2011).
19. McCulloch, W. S. & Pitts, W. A logical calculus of the ideas immanent in nervous activity. *Bull. Math. Biophys.* **5**, 115 (1943).
20. Unke, O. T. & Meuwly, M. Toolkit for the construction of reproducing kernel-based representations of data: application to multidimensional potential energy surfaces. *J. Chem. Inf. Model.* **57**, 1923 (2017).
21. Chmiela, S. et al. Machine learning of accurate energy-conserving molecular force fields. *Sci. Adv.* **3**, e1603015 (2017).
22. Chmiela, S., Saucedo, H. E., Poltavsky, I., Müller, K.-R. & Tkatchenko, A. sGDML: Constructing accurate and data efficient molecular force fields using machine learning. *Computer Phys. Commun.* **240**, 38 (2019).
23. Smith, J. S., Isayev, O. & Roitberg, A. E. ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost. *Chem. Sci.* **8**, 3192 (2017).
24. Zhang, L., Han, J., Wang, H., Car, R. & Weinan, E. Deep potential molecular dynamics: a scalable model with the accuracy of quantum mechanics. *Phys. Rev. Lett.* **120**, 143001 (2018).
25. Unke, O. T. & Meuwly, M. A reactive, scalable, and transferable model for molecular energies from a neural network approach based on local information. *J. Chem. Phys.* **148**, 241708 (2018).
26. Christensen, A. S., Bratholm, L. A., Faber, F. A. & Anatole von Lilienfeld, O. FCHL revisited: faster and more accurate quantum machine learning. *J. Chem. Phys.* **152**, 044107 (2020).
27. Vapnik, V. N. *The Nature of Statistical Learning Theory*. (Springer, 1995).
28. Cortes, C. & Vapnik, V. Support-vector networks. *Mach. Learn.* **20**, 273 (1995).
29. Müller, K.-R., Mika, S., Ratsch, G., Tsuda, K. & Schölkopf, B. An introduction to kernel-based learning algorithms. *IEEE Trans. Neural Netw.* **12**, 181 (2001).
30. Schölkopf, B. & Smola, A. J. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. (MIT press, 2002).
31. Braun, M. L., Buhmann, J. M. & Müller, K.-R. On relevant dimensions in kernel feature spaces. *J. Mach. Learn. Res.* **9**, 1875 (2008).
32. Hansen, K. et al. Assessment and validation of machine learning methods for predicting molecular atomization energies. *J. Chem. Theory Comput.* **9**, 3404 (2013).
33. Schütt, K. T., Arbabzadah, F., Chmiela, S., Müller, K. R. & Tkatchenko, A. Quantum-chemical insights from deep tensor neural networks. *Nat. Commun.* **8**, 13890 (2017).
34. Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O. & Dahl, G. E. Neural message passing for quantum chemistry, in *International Conference on Machine Learning*, 1263–1272 <https://proceedings.mlr.press/v70/gilmer17a.html> (2017).
35. Schütt, K. T., Saucedo, H. E., Kindermans, P.-J., Tkatchenko, A. & Müller, K.-R. SchNet—a deep learning architecture for molecules and materials. *J. Chem. Phys.* **148**, 241722 (2018).
36. Lubbers, N., Smith, J. S. & Barros, K. Hierarchical modeling of molecular energies using a deep neural network. *J. Chem. Phys.* **148**, 241715 (2018).
37. Unke, O. T. & Meuwly, M. PhysNet: A neural network for predicting energies, forces, dipole moments, and partial charges. *J. Chem. Theory Comput.* **15**, 3678 (2019).
38. Friesner, R. A. Ab initio quantum chemistry: methodology and applications. *Proc. Natl Acad. Sci. USA* **102**, 6648 (2005).
39. Born, M. & Einstein, A. *The Born-Einstein Letters 1916–1955* (Macmillan, 2005).
40. Noodleman, L., Peng, C., Case, D. & Mouesca, J.-M. Orbital interactions, electron delocalization and spin coupling in iron-sulfur clusters. *Coord. Chem. Rev.* **144**, 199 (1995).
41. Dreuw, A., Weisman, J. L. & Head-Gordon, M. Long-range charge-transfer excited states in time-dependent density functional theory require non-local exchange. *J. Chem. Phys.* **119**, 2943 (2003).
42. Duda, L.-C. et al. Resonant inelastic X-Ray scattering at the oxygen K resonance of NiO: nonlocal charge transfer and double-singlet excitations. *Phys. Rev. Lett.* **96**, 067402 (2006).
43. Bellec, A. et al. Nonlocal activation of a bistable atom through a surface state charge-transfer process on Si(100)-(2 × 1):H. *Phys. Rev. Lett.* **105**, 048302 (2010).
44. Boström, E. V., Mikkelsen, A., Verdozzi, C., Perfetto, E. & Stefanucci, G. Charge separation in donor-C₆₀ complexes with real-time green functions: the importance of nonlocal correlations. *Nano Lett.* **18**, 785 (2018).
45. Ghasemi, S. A., Hofstetter, A., Saha, S. & Goedecker, S. Interatomic potentials for ionic systems with density functional accuracy based on charge densities obtained by a neural network. *Phys. Rev. B* **92**, 045131 (2015).
46. Rappe, A. K. & Goddard III, W. A. Charge equilibration for molecular dynamics simulations. *J. Phys. Chem.* **95**, 3358 (1991).
47. Wilmer, C. E., Kim, K. C. & Snurr, R. Q. An extended charge equilibration method. *J. Phys. Chem. Lett.* **3**, 2506 (2012).
48. Cheng, Y.-T. et al. A charge optimized many-body (comb) potential for titanium and titania. *J. Phys. Condens. Matter* **26**, 315007 (2014).
49. Ko, T. W., Finkler, J. A., Goedecker, S. & Behler, J. A fourth-generation high-dimensional neural network potential with accurate electrostatics including non-local charge transfer. *Nat. Commun.* **12**, 398 (2021).
50. Zubatyuk, R., Smith, J. S., Nebgen, B. T. et al. Teaching a neural network to attach and detach electrons from molecules. *Nat. Commun.* **12**, 4870 (2021).
51. Zubatyuk, R., Smith, J. S., Leszczynski, J. & Isayev, O. Accurate and transferable multitask prediction of chemical properties with an atoms-in-molecules neural network. *Sci. Adv.* **5**, eaav6490 (2019).
52. Xie, X., Persson, K. A. & Small, D. W. Incorporating electronic information into machine learning potential energy surfaces via approaching the ground-state electronic energy as a function of atom-based electronic populations. *J. Chem. Theory Comput.* **16**, 4256 (2020).
53. Qiao, Z., Welborn, M., Anandkumar, A., Manby, F. R. & Miller III, T. F. OrbNet: Deep learning for quantum chemistry using symmetry-adapted atomic-orbital features. *J. Chem. Phys.* **153**, 124111 (2020).
54. Pfau, D., Spencer, J. S., Matthews, A. G. & Foulkes, W. M. C. Ab initio solution of the many-electron Schrödinger equation with deep neural networks. *Phys. Rev. Res.* **2**, 033429 (2020).
55. Hermann, J., Schätzle, Z. & Noé, F. Deep-neural-network solution of the electronic Schrödinger equation. *Nat. Chem.* **12**, 891–897 (2020).
56. Scherbela, M., Reisenhofer, R., Gerard, L., Marquetand, P. & Grohs, P. Solving the electronic Schrödinger equation for multiple nuclear geometries with weight-sharing deep neural networks. *arXiv preprint arXiv:2105.08351* (2021).
57. Klicpera, J., Groß, J. & Günnemann, S. Directional message passing for molecular graphs. *International Conference on Learning Representations (ICLR)* <https://openreview.net/forum?id=B1eWbxStPH> (2020).
58. Bernstein, S. Démonstration du théorème de Weierstrass fondée sur le calcul des probabilités. *Commun. Kharkov Math. Soc.* **13**, 1 (1912).
59. Thomas, N. et al. Tensor field networks: Rotation and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219* (2018).
60. Anderson, Brandon and Hy, Truong Son and Kondor, Risi, Cormorant: Covariant Molecular Neural Networks, *Advances in Neural Information Processing Systems* **32**, <https://papers.nips.cc/paper/2019/hash/03573b32b2746e6e8ca98b9123f2249b-Abstract.html> (2019).
61. Schütt, K. T., Unke, O. T. & Gastegger, M. Equivariant message passing for the prediction of tensorial properties and molecular spectra. *Proceedings of the 38th International Conference on Machine Learning*, 9377–9388 (2021).
62. Batzner, S. et al. SE(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *arXiv preprint arXiv:2101.03164* (2021).
63. Bartók, A. P., Kondor, R. & Csányi, G. On representing chemical environments. *Phys. Rev. B* **87**, 184115 (2013).
64. Thompson, A. P., Swiler, L. P., Trott, C. R., Foiles, S. M. & Tucker, G. J. Spectral neighbor analysis method for automated generation of quantum-accurate interatomic potentials. *J. Comput. Phys.* **285**, 316 (2015).
65. Yao, K., Herr, J. E., Toth, D. W., Mckintyre, R. & Parkhill, J. The TensorMol-0.1 model chemistry: a neural network augmented with long-range physics. *Chem. Sci.* **9**, 2261 (2018).
66. Grisafi, A. & Ceriotti, M. Incorporating long-range physics in atomic-scale machine learning. *J. Chem. Phys.* **151**, 204105 (2019).
67. Bereau, T., DiStasio Jr, R. A., Tkatchenko, A. & Von Lilienfeld, O. A. Non-covalent interactions across organic and biological subsets of chemical space: physics-based potentials parametrized from machine learning. *J. Chem. Phys.* **148**, 241706 (2018).
68. Schwilk, M., Tahchieva, D. N. & von Lilienfeld, O. A. The QMspin data set: Several thousand carbene singlet and triplet state structures and vertical spin gaps computed at MRCISD+Q-F12/cc-pVDZ-F12 level of theory. *Mater. Cloud Arch.* <https://doi.org/10.24435/materialscloud:2020.0051/v1> (2020a).
69. Schwilk, M., Tahchieva, D. N. & von Lilienfeld, O. A. Large yet bounded: Spin gap ranges in carbenes. *arXiv preprint arXiv:2004.10600* (2020b).
70. Deringer, V. L., Caro, M. A. & Csányi, G. A general-purpose machine-learning force field for bulk and nanostructured phosphorus. *Nat. Commun.* **11**, 1 (2020).

71. Hoja, J. et al. QM7-X, a comprehensive dataset of quantum-mechanical properties spanning the chemical space of small organic molecules. *Sci. Data* **8**, 43 (2021).
72. Blum, 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**, 8732 (2009).
73. Adamo, C. & Barone, V. Toward reliable density functional methods without adjustable parameters: the PBE0 model. *J. Chem. Phys.* **110**, 6158 (1999).
74. Tkatchenko, A., DiStasio Jr, R. A., Car, R. & Scheffler, M. Accurate and efficient method for many-body van der Waals interactions. *Phys. Rev. Lett.* **108**, 236402 (2012).
75. Folmsbee, D. & Hutchison, G. Assessing conformer energies using electronic structure and machine learning methods. *Int. J. Quantum Chem.* **121**, e26381 (2021).
76. Perdew, J. P., Burke, K. & Ernzerhof, M. Generalized gradient approximation made simple. *Phys. Rev. Lett.* **77**, 3865 (1996).
77. Tkatchenko, A. & Scheffler, M. Accurate molecular van der Waals interactions from ground-state electron density and free-atom reference data. *Phys. Rev. Lett.* **102**, 073005 (2009).
78. Vaswani, Ashish and Shazeer, Noam and Parmar, Niki and Uszkoreit, Jakob and Jones, Llion and Gomez, Aidan N and Kaiser, Lukasz and Polosukhin, Illia, Attention is All you Need, Advances in Neural Information Processing Systems **30**, <https://papers.nips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html> (2017).
79. Choromanski, K. et al. Rethinking attention with performers. International Conference on Learning Representations <https://openreview.net/forum?id=Ua6zuk0WRH> (2021).
80. Ziegler, J. F., Littmark, U. & Biersack, J. P. *The Stopping and Range of Ions in Solids* (Pergamon, 1985).
81. Artrith, N., Morawietz, T. & Behler, J. High-dimensional neural-network potentials for multicomponent systems: applications to zinc oxide. *Phys. Rev. B* **83**, 153101 (2011).
82. Morawietz, T., Sharma, V. & Behler, J. A neural network potential-energy surface for the water dimer based on environment-dependent atomic energies and charges. *J. Chem. Phys.* **136**, 064103 (2012).
83. Morawietz, T. & Behler, J. A density-functional theory-based neural network potential for water clusters including van der Waals corrections. *J. Phys. Chem. A* **117**, 7356 (2013).
84. Uteva, E., Graham, R. S., Wilkinson, R. D. & Wheatley, R. J. Interpolation of intermolecular potentials using Gaussian processes. *J. Chem. Phys.* **147**, 161706 (2017).
85. Samek, W., Montavon, G., Vedaldi, A., Hansen, L. K. & Müller, K.-R. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, vol. 11700 (Springer Nature, 2019).
86. Samek, W., Montavon, G., Lapuschkin, S., Anders, C. J. & Müller, K.-R. Explaining deep neural networks and beyond: a review of methods and applications. *Proc. IEEE* **109**, 247 (2021).
87. Lapuschkin, S. et al. Unmasking clever Hans predictors and assessing what machines really learn. *Nat. Commun.* **10**, 1096 (2019).
88. Hendrycks, D. & Gimpel, K. Gaussian error linear units (GELUs). *arXiv preprint arXiv:1606.08415* (2016).
89. Elfwing, S., Uchibe, E. & Doya, K. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Netw.* **107**, 3 (2018).
90. Ramachandran, P., Zoph, B. & Le, Q. V. Searching for activation functions. *arXiv preprint arXiv:1710.05941* (2017).
91. Nair, V. & Hinton, G. E. Rectified linear units improve restricted boltzmann machines. in *ICML International Conference on Machine Learning* (2010).
92. Glorot, X. & Bengio, Y. Understanding the difficulty of training deep feedforward neural networks, in *Proc. thirteenth international conference on artificial intelligence and statistics*, 249–256 (JMLR Workshop and Conference Proceedings, 2010).
93. Srivastava, R., Greff, K. & Schmidhuber, J. Highway networks. *arXiv preprint arXiv:1505.00387* (2015).
94. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition, in *Proc. IEEE conference on computer vision and pattern recognition* 770–778 (2016).
95. He, K., Zhang, X., Ren, S. & Sun, J. Identity mappings in deep residual networks. in *European conference on computer vision*, 630–645 (Springer, 2016).
96. Kowalski, G. J. In *Information Retrieval Systems: Theory and Implementation*, Vol. 1 (Springer, 2007).
97. Ewald, P. P. Die Berechnung optischer und elektrostatischer Gitterpotentiale. *Ann. der Phys.* **369**, 253 (1921).
98. Gastegger, M., Behler, J. & Marquetand, P. Machine learning molecular dynamics for the simulation of infrared spectra. *Chem. Sci.* **8**, 6924 (2017).
99. Hermann, J., DiStasio Jr, R. A. & Tkatchenko, A. First-principles models for van der Waals interactions in molecules and materials: concepts, theory, and applications. *Chem. Rev.* **117**, 4714 (2017).
100. Caldeweyher, E. et al. A generally applicable atomic-charge dependent london dispersion correction. *J. Chem. Phys.* **150**, 154122 (2019).
101. Reddi, S. J., Kale, S. & Kumar, S. On the convergence of Adam and beyond. International Conference on Learning Representations (2018).
102. Hirshfeld, F. L. Bonded-atom fragments for describing molecular charge densities. *Theor. Chim. Acta* **44**, 129 (1977).
103. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. International Conference on Learning Representations (2015).
104. Bannwarth, C., Ehlert, S. & Grimme, S. GFN2-xTB—an accurate and broadly parametrized self-consistent tight-binding quantum chemical method with multipole electrostatics and density-dependent dispersion contributions. *J. Chem. Theory Comput.* **15**, 1652 (2019).
105. Blum, V. et al. Ab initio molecular simulations with numeric atom-centered orbitals. *Computer Phys. Commun.* **180**, 2175 (2009).
106. Ren, X. et al. Resolution-of-identity approach to Hartree–Fock, hybrid density functionals, RPA, MP2 and GW with numeric atom-centered orbital basis functions. *N. J. Phys.* **14**, 053020 (2012).
107. Pozdnyakov, S. N. et al. Incompleteness of atomic structure representations. *Phys. Rev. Lett.* **125**, 166001 (2020b).
108. Unke, O. *Singlet/triplet carbene and Ag₃⁺/Ag₃⁻ data*. <https://doi.org/10.5281/zenodo.5115732> (2021).
109. Pozdnyakov, S., Willatt, M. & Ceriotti, M. Randomly-displaced methane configurations. *Mater. Cloud Arch.* <https://doi.org/10.24435/materialscloud:qy-dp> (2020a).
110. Hoja, J. et al. QM7-X: A comprehensive dataset of quantum-mechanical properties spanning the chemical space of small organic molecules, (Version 1.0) [Data set]. *Zenodo* <https://doi.org/10.5281/zenodo.3905361> (2020).
111. Ko, T. W., Finkler, J. A., Goedecker, S. & Behler, J. A fourth-generation high-dimensional neural network potential with accurate electrostatics including non-local charge transfer. *Mater. Cloud Arch.* <https://doi.org/10.24435/materialscloud:f3-yh> (2020).
112. Paszke, A. et al. PyTorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703* (2019).

Acknowledgements

O.T.U. acknowledges funding from the Swiss National Science Foundation (Grant No. P2BSP2_188147). We thank the authors of ref. ¹⁰⁷ for sharing raw data for reproducing the learning curves shown in Supplementary Fig. 2 and the geometries displayed in Supplementary Fig. 1e. K.R.M. was supported in part by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea Government (No. 2019-0-00079, Artificial Intelligence Graduate School Program, Korea University), and was partly supported by the German Ministry for Education and Research (BMBF) under Grants 01IS14013A-E, 01GQ1115, 01GQ0850, 01IS18025A, and 01IS18037A; the German Research Foundation (DFG) under Grant Math+, EXC 2046/1, Project ID 390685689. Correspondence should be addressed to O.T.U. and K.R.M. We thank Alexandre Tkatchenko and Hartmut Maennel for very helpful discussions and feedback on the manuscript.

Author contributions

O.T.U. designed the SpookyNet architecture, wrote the code, and ran experiments. O.T.U., S.C., M.G., K.T.S., and H.E.S. designed the experiments. O.T.U. and K.R.M. designed the figures, all authors contributed to writing the manuscript.

Funding

Open Access funding enabled and organized by Projekt DEAL.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-021-27504-0>.

Correspondence and requests for materials should be addressed to Oliver T. Unke or Klaus-Robert Müller.

Peer review information *Nature Communications* thanks Jan Gerit Brandenburg and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021