

ARTICLE OPEN



On-the-fly active learning of interpretable Bayesian force fields for atomistic rare events

Jonathan Vandermause^{1,2}, Steven B. Torrisi¹, Simon Batzner^{2,3}, Yu Xie², Lixin Sun², Alexie M. Kolpak⁴ and Boris Kozinsky^{1,5}

Machine learned force fields typically require manual construction of training sets consisting of thousands of first principles calculations, which can result in low training efficiency and unpredictable errors when applied to structures not represented in the training set of the model. This severely limits the practical application of these models in systems with dynamics governed by important rare events, such as chemical reactions and diffusion. We present an adaptive Bayesian inference method for automating the training of interpretable, low-dimensional, and multi-element interatomic force fields using structures drawn on the fly from molecular dynamics simulations. Within an active learning framework, the internal uncertainty of a Gaussian process regression model is used to decide whether to accept the model prediction or to perform a first principles calculation to augment the training set of the model. The method is applied to a range of single- and multi-element systems and shown to achieve a favorable balance of accuracy and computational efficiency, while requiring a minimal amount of ab initio training data. We provide a fully open-source implementation of our method, as well as a procedure to map trained models to computationally efficient tabulated force fields.

npj Computational Materials (2020)6:20; <https://doi.org/10.1038/s41524-020-0283-z>

INTRODUCTION

Recent machine learned (ML) force fields have been shown to achieve high accuracy for a number of molecular and solid-state systems^{1–11}. These methods provide a promising path toward long, large-scale molecular dynamics (MD) simulations driven by force predictions that approach the accuracy of quantum mechanical methods like density functional theory (DFT). However, most currently available ML force fields return point estimates of energies, forces, and stresses rather than predictive distributions that reflect model uncertainty, making the incorporation of accurate uncertainty estimates into ML force fields an outstanding challenge^{12–18}. Without model uncertainty, a laborious fitting procedure is required, which usually involves manually or randomly selecting thousands of reference structures from a database of first principles calculations. In production MD runs, a lack of principled means to compute predictive uncertainties makes it difficult to determine when the force field is trustworthy, leading to unreliable results and lack of guidance on how to update the model in the presence of new data.

Here, we show that active learning based on Gaussian process (GP) regression can accelerate and automate the training of high-quality force fields by making use of accurate internal estimates of model error. By combining DFT with low-dimensional GP regression models during molecular dynamics simulations, accurate force fields for a range of single- and multi-element systems are obtained with ~ 100 DFT calculations. Moreover, we demonstrate that the model can be flexibly and automatically updated when the system deviates from previous training data. Such a reduction in the computational cost of training and updating force fields promises to extend ML modeling to a wider class of materials than has been possible to date. The method is shown to successfully model rapid crystal melts and rare diffusive events, and so we call our method FLARE: Fast Learning of

Atomistic Rare Events, and make the open-source software freely available online (<https://github.com/mir-group/flare>).

The key contribution of this work that makes on-the-fly learning possible is the development of a fully interpretable low-dimensional and nonparametric force field that provides trustworthy estimates of model uncertainty. Typical ML force fields involve regression over a high-dimensional descriptor space chosen either on physical grounds^{19,20} or learned directly from ab initio data^{6,10}. These approaches involve highly flexible models with many physically non-interpretable parameters, complicating the task of inferring a posterior distribution over model parameters. We instead bypass the need for a high-dimensional descriptor by imposing a physical prior that constrains the model to n -body interactions, with high accuracy observed in practice with 2- and 3-body models. Because the low-dimensional descriptor space of our models can be sampled with a small amount of training data, our method avoids sparsification, a procedure that is used in Gaussian approximation potentials to make inference tractable with many-body descriptors like SOAP^{20–22}, but that requires approximate treatment of GP uncertainty estimates^{23,24}. The learning task is simplified as a result, making it possible to automatically tune the model's hyperparameters in a data-driven fashion and derive trustworthy estimates of model uncertainty. This opens the door to a practical uncertainty-driven method for selecting training points "on the fly"²⁵, allowing an accurate force field to be trained with a minimal number of relatively expensive first principles calculations.

The resulting GP-based force fields are interpretable in three important respects. First, the underlying energy model of the GP is a physically motivated sum over n -body contributions, such that each cluster of $n - 1$ neighbors in an atom's environment makes a direct contribution to the force on that atom. This establishes a connection to previous physically motivated force fields, most

¹Department of Physics, Harvard University, Cambridge, MA 02138, USA. ²John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA. ³Center for Computational Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA. ⁴Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA. ⁵Bosch Research, Cambridge, MA 02139, USA. ✉email: jonathan_vandermause@g.harvard.edu; bkoz@g.harvard.edu

notably the Stillinger-Weber force field²⁶, which also sums over 2- and 3-body contributions but is limited to a specific analytic form. Our models, by contrast, learn nonparametric 2- and 3-body functions directly from ab initio data, allowing the models to generalize well to complex multi-element systems, as we show in the Results section below. Second, the model does not require a descriptor of the entire local environment of an atom, instead relying on a kernel that directly compares interatomic distances of small clusters of atoms. As a result, the only free parameters in the model are a small set of hyperparameters of the GP kernel function, each of which has a direct interpretation and can be rigorously optimized by maximizing the log marginal likelihood of the training data. Neural network and Gaussian approximation potentials, on the other hand, rely on complex high-dimensional descriptors of an atom's environment, making it less apparent how the force acting on an atom is related to the configuration of its neighbors. Finally, and most importantly for active learning, the uncertainty estimates of our GP models break down into two contributions: the epistemic uncertainty $\sigma_{i\alpha}$, which is assigned to each atom i and force component α and is determined by distance from the training set, and the noise uncertainty σ_n , which characterizes fundamental variability in the training data that cannot be captured by the model. The latter source of error arises from several simplifying approximations that improve computational efficiency, including the exclusion of interactions outside the cutoff radius of the model, the decomposition of global energies into local contributions, and the restriction to 2- and 3-body interactions^{4,22}. By optimizing the noise uncertainty σ_n of the GP, the combined magnitude of these errors can be learned directly from the data (see Methods). The interpretable uncertainties derived from the GP model provide a principled basis for automated training, in which a local environment is added to the training set of the model when the epistemic uncertainty $\sigma_{i\alpha}$ on a force component exceeds a chosen multiple of the noise uncertainty σ_n .

Other GP and active learning based methods for force field training have been proposed in the literature, and we discuss them briefly here to put our method in context. Bartók et al. pioneered the use of GP-based force fields in the Gaussian approximation potential (GAP) framework^{21,22}, with subsequent applications combining 2- and 3-body descriptors with the many-body SOAP kernel to achieve high accuracy for a range of extended systems^{4,7,20}. Recent GAP studies have reported uncertainty estimates on local energy predictions⁸ and introduced self-guided protocols for learning force fields based on random structure searching rather than uncertainty-driven active

learning^{7,27}. Rupp et al.²⁸ and more recently Uteva et al.²⁹ used GP regression to model potential energy surfaces of small molecular systems with active learning, and Smith et al. recently proposed a query-by-committee procedure for actively learning neural network force fields for small molecules³⁰. On-the-fly force field training for extended systems was first proposed by Li, Kermode, and De Vita²⁵, but the method relied on performing DFT calculations to evaluate model error due to a lack of correlation between the internal error of their GP model and true model error³¹. Podryabinkin and Shapeev developed an on-the-fly method for their linear moment tensor potentials³² using the D-optimality criterion, which provides an internal information-theoretic measure of distance from the training set¹³, with subsequent applications to molecules, alloys, and crystal structure prediction^{18,33,34}. The D-optimality criterion is usually restricted to linear models and does not provide direct error estimates on model predictions. More recently, Jinnouchi et al. combined a multi-element variant of the SOAP kernel with Bayesian linear regression to obtain direct Bayesian error estimates on individual force components, which was used to perform on-the-fly training of force fields to study melting points and perovskite phase transitions^{35,36}. This approach relies on a decomposition of the atomic density of each atom into many-body descriptors based on spherical Bessel functions and spherical harmonics, with the number of descriptors growing quadratically with the number of elements in the system³⁷. The machine learned force fields presented here possess four important features that have not been simultaneously achieved before: they are nonparametric, fully Bayesian, explicitly multi-element, and can be mapped to highly efficient tabulated force fields, making our automated method for training these models widely applicable to a range of complex materials.

RESULTS

FLARE: an on-the-fly learning method

The goal of FLARE is to automate the training of accurate and computationally efficient force fields that can be used for large-scale molecular dynamics simulations of multi-element systems. The low-dimensional GP kernel that we use throughout this work, sketched in Fig. 1a, is calculated by comparing interatomic distances of clusters of two and three atoms, similar to the single-element kernel presented by Glielmo, Zeni, and De Vita³⁸ but here generalized to arbitrarily many chemical species. If the two clusters are not of the same type, as determined by the chemical species of the atoms in the cluster, the kernel is assigned a value

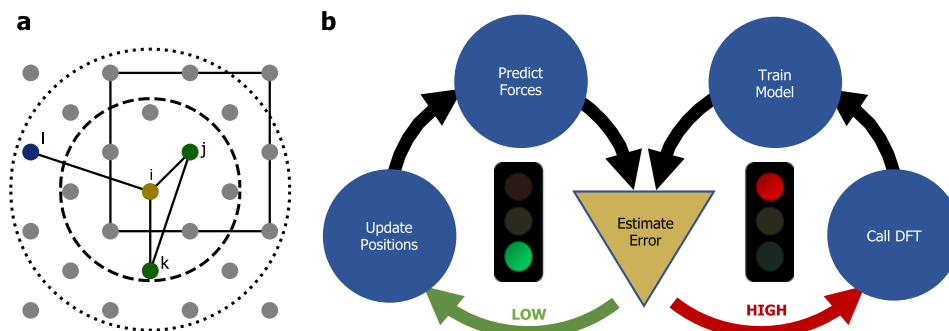


Fig. 1 Fast learning of atomistic rare events (FLARE): an on-the-fly learning method for automatically training force fields during molecular dynamics. **a** The 2- and 3-body multi-element kernels used in this work. The local environment of the central atom (gold) consists of all atoms within the 2- and 3-body cutoff spheres (dotted and dashed lines, respectively), including images of atoms in the primary periodic cell (solid square). The kernel is calculated by comparing clusters of two and three atoms of the same type, as determined by the chemical species of the atoms in the cluster. **b** An overview of the on-the-fly learning algorithm. Left loop: molecular dynamics steps are proposed by the current GP force field, with the epistemic uncertainties $\sigma_{i\alpha}$ on all force components monitored at each step. Right loop: if the epistemic uncertainty on a force component rises above a chosen multiple of the optimized noise uncertainty σ_n of the GP, DFT is called and the training set of the GP is updated with the highest uncertainty local environments.

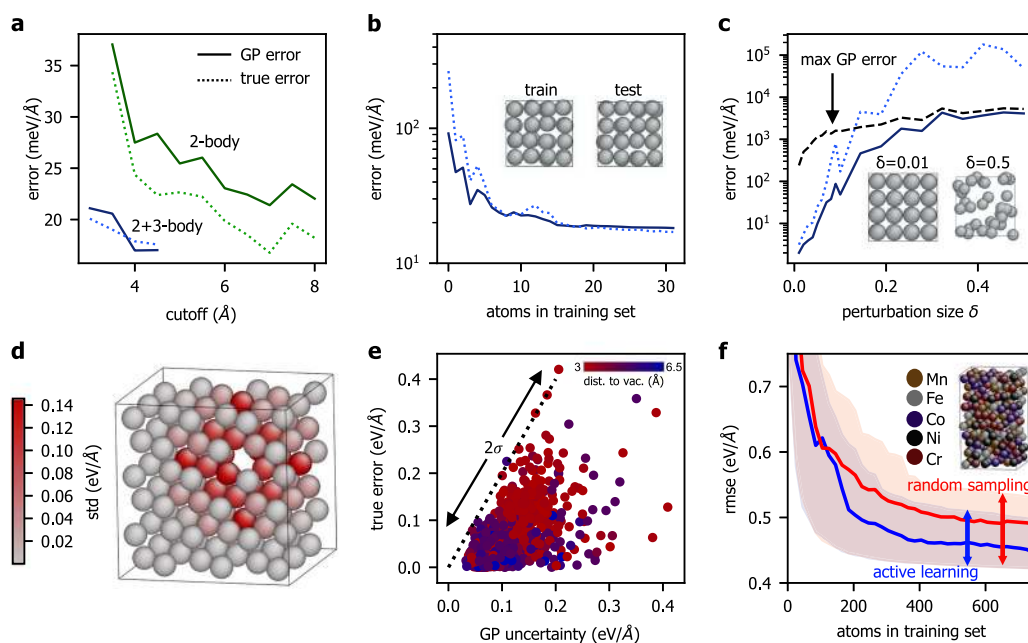


Fig. 2 Tests of Gaussian process (GP) uncertainties. **a** Optimized noise uncertainty σ_n (solid) and root mean squared error (RMSE, dotted) of GPs trained on an aluminum structure as a function of the cutoff radius r_{cut} of the local environment for 2- and 2+3-body GP models (green and blue, respectively). **b** Average GP uncertainty $\sqrt{\sigma_n^2 + \bar{\sigma}_{ia}^2}$ (solid) versus true RMSE (dotted) as a function of the number of training environments. **c** Mean epistemic uncertainty $\bar{\sigma}_{ia}$ (solid) and RMSE (dotted) on test structures with atomic coordinates perturbed from $\delta = 1\text{--}50\%$ of the lattice parameter, with the upper bound on the epistemic uncertainty (dashed) approached for $\delta > 20\%$. **d** Uncertainties of individual force components acting on the atom, with atoms closer to the vacancy having more uncertain forces. **e** Comparison of GP uncertainties and true model error for individual force components predicted on 10 randomly perturbed Al vacancy structures, with most true errors falling within two standard deviations $\sigma = \sqrt{\sigma_n^2 + \bar{\sigma}_{ia}^2}$ of the predictive posterior distribution of the GP (dotted). **f** Learning curves of GP models trained on 5-element high entropy alloy structures, with training environments selected randomly (red) and with active learning (blue). The RMSE on force components of an independent test structure is plotted along with the distribution of uncertainties, shown as a band between the minimum and maximum uncertainties on force components in the structure.

of zero, allowing the GP to differentiate between chemical species while remaining low-dimensional (see Methods). Restricting the model to a sum over two- and three-dimensional contributions reduces the cost of training the model, allowing the descriptor space to be systematically sampled with a relatively small number of DFT calculations, and also reduces the cost of production MD runs with the final trained model, since the GP can be mapped onto efficient cubic spline models that allow the 2- and 3-body contributions to the force on an atom to be directly evaluated³⁸. We have implemented this mapping as a pair style in the molecular dynamics software LAMMPS, allowing us to study multi-element systems containing more than ten thousand atoms over nanosecond timescales (Fig. 5 below).

The low dimensionality of our models also makes it practically feasible to rigorously optimize the hyperparameters of the kernel function, which leads to trustworthy estimates of model uncertainty. The reliability of these uncertainties is the key feature of our approach that enables FLARE, an adaptive method for training force fields on the fly during molecular dynamics. As sketched in Fig. 1b, the algorithm takes an arbitrary structure as input and begins with a call to DFT, which is used to train an initial GP model on the forces acting on an arbitrarily chosen subset of atoms in the structure. The GP then proposes an MD step by predicting the forces on all atoms, at which point a decision is made about whether to accept the predictions of the GP or to perform a DFT calculation. The decision is based on the epistemic uncertainty σ_{ia} of each GP force component prediction (defined in Eq. (5) of Methods), which estimates the error of the prediction due to dissimilarity between the atom's environment and the local environments stored in the training set of the GP. In particular, if

any σ_{ia} exceeds a chosen multiple of the current noise uncertainty σ_n of the model, a call to DFT is made and the training set is augmented with the forces acting on the N_{added} highest uncertainty local environments, the precise number of which can be tuned to increase training efficiency. All hyperparameters, including the noise uncertainty σ_n , are optimized whenever a local environment and its force components are added to the training set, allowing the error threshold to adapt to novel environments encountered during the simulation (see Methods).

Characterization of model uncertainty

To justify an on-the-fly learning algorithm, we first characterize the noise and epistemic uncertainties of GP models constructed with the 2- and 3-body kernels described above, and compare them against test errors on out-of-sample structures. Importantly, the optimized noise uncertainty σ_n and epistemic uncertainties σ_{ia} are found to provide a sensitive probe of true model error, with the noise uncertainty capturing the baseline error level of model predictions on local environments that are well represented in the training set, and the epistemic uncertainties capturing error due to deviation from the training data. In Fig. 2a–c, we test the relationship between GP uncertainties and true error by performing a set of plane-wave DFT calculations on a 32-atom supercell of FCC aluminum with the atoms randomly perturbed from their equilibrium sites. In Fig. 2a, we examine the noise uncertainty σ_n as a function of the cutoff radius of the model, which determines the degree of locality of the trained force field. 2- and 2+3-body GP models were trained on forces acting on atoms in a single structure and then tested on an independently generated structure, with the atomic coordinates in both cases randomly

perturbed by up to 5% of the lattice parameter, $a_{\text{lat}} = 4.046 \text{ \AA}$. For the 2-body models, the cutoff radius was swept from 3.5 to 8 \AA in increments of 0.5 \AA , and for the 2+3-body models, the 2-body cutoff was held fixed at 6 \AA and the 3-body cutoff was swept from 3 to 4.5 \AA . The optimized noise uncertainty σ_n , plotted in Fig. 2a closely tracks the root mean squared error (RMSE) on the test structure for the range of examined cutoff values. The observed correlation provides a principled way to select the cutoff radius of the GP, showing that the expected error of a model with a given cutoff can be directly estimated from the optimized noise uncertainty σ_n when the GP model has been trained on sufficient data.

When the GP model is trained on insufficient data, the epistemic uncertainties σ_{ia} rise above the noise uncertainty σ_n , indicating that the model requires additional training data to make accurate force estimates. The utility of the epistemic uncertainty is illustrated in Fig. 2b, which examines GP uncertainties as a function of the amount of data in the training set. Using the same training and test structures as Fig. 2a, a 2+3-body GP model with a 6- \AA 2-body cutoff and 4- \AA 3-body cutoff was constructed by adding local environments one by one to the training set and evaluating the RMSE and GP uncertainty after each update. The average GP uncertainty $\sqrt{\sigma_n^2 + \bar{\sigma}_{ia}^2}$ closely tracks the RMSE, where $\bar{\sigma}_{ia}$ is the mean epistemic uncertainty over all force components in the test structure.

We also demonstrate in Fig. 2c that the epistemic uncertainty provides an accurate indicator of model error when the model is forced to extrapolate on local environments that are significantly different from local environments in the training set. To systematically investigate distance from the training set, a 2+3-body GP model was trained on a single aluminum structure with atomic coordinates perturbed by $\delta = 5\%$ of the lattice parameter and tested on structures generated with values of δ ranging from 1 to 50%, with $\delta = 50\%$ giving rise to a highly distorted structure with a mean absolute force component of 28.6 eV/ \AA and a maximum absolute force component of 200.5 eV/ \AA (compared to a mean of 0.50 eV/ \AA and maximum of 1.48 eV/ \AA for the training structure). As shown in Fig. 2c, the mean epistemic uncertainty $\bar{\sigma}_{ia}$ increases with δ and exceeds the optimized noise uncertainty of $\sigma_n = 11.53 \text{ meV/\AA}$ for $\delta > 5\%$, demonstrating the ability of the GP to detect when it is predicting on structures that are outside the training set. This capability is crucial for on-the-fly learning, as the model must be able to flag when additional training data is needed in order to accurately estimate forces. We furthermore observe that the error is substantially underestimated for large values of δ due to an upper bound on the epistemic uncertainty imposed by the signal variance hyperparameters of the kernel function, with the bound nearly saturated for $\delta > 20\%$ (see Methods for the definition of this bound). This emphasizes the importance of re-optimizing the hyperparameters when additional data is introduced to the training set, allowing the model to adapt to novel structures.

In Fig. 2d, e we demonstrate that GP uncertainties on individual force components can also provide valuable information about the expected errors on structures not represented in the training set. Figure 2d shows individual GP uncertainties $\sqrt{\sigma_{ia}^2 + \sigma_n^2}$ on the predicted force components of a relaxed vacancy structure when the GP was trained on bulk local environments only. Each atom is colored according to the maximum uncertainty of the three predicted force components acting on the atom, with atoms closer to the defect tending to have higher uncertainties. This test was repeated for ten randomly perturbed vacancy structures, with the true error plotted in Fig. 2e against the GP uncertainty $\sqrt{\sigma_{ia}^2 + \sigma_n^2}$ of each force component, showing that higher uncertainties coincide with a wider spread in the true error.

We finally demonstrate in Fig. 2f that the GP uncertainties are trustworthy for more complex multi-element systems. In this test, two GP models were trained on the five-element high entropy

alloy (HEA) DFT forces of Zhang et al.¹⁰, with training environments selected randomly for the first GP model and with active learning for the second. Specifically, thirty-nine HEA structures were drawn from the "rand1" portion of this dataset, and for each structure, twenty training environments were selected either at random or by identifying the highest uncertainty environments in the structure. After each update to the training set, both the GP uncertainties and true model error on an independent HEA structure were evaluated (with the test structure taken from the "rand2" portion of the dataset and having a different random allocation of elements). The distribution of total uncertainties $\sqrt{\sigma_{ia}^2 + \sigma_n^2}$ on force components in the test structure is shown for both models in Fig. 2f by plotting a band between the minimum and maximum uncertainties, which encloses the true RMSE. Actively selecting environments based on model uncertainty has the effect of shifting the learning curve downward, with the actively trained GP reaching a RMSE of 0.445 eV/ \AA on the test structure. The GP model obtained with active learning was subsequently mapped to a tabulated force field in order to rapidly evaluate forces on the entire "rand2" test set of Zhang et al.¹⁰, which consisted of 149 HEA structures with elements placed at random lattice sites. The RMSE averaged over all test structures was found to be 0.466 eV/ \AA for the tabulated GP model, comparable to the RMSE of 0.410 eV/ \AA reported for the deep neural network model of Zhang et al.¹⁰ and outperforming the Deep Potential model of Zhang et al.⁹, which achieved a RMSE of 0.576 eV/ \AA on the same test set. We note that both neural network models were trained on 400 HEA structures¹⁰, which exceeds the number of structures the GP was trained on by more than an order of magnitude.

Aluminum crystal melt

As a first demonstration of on-the-fly learning driven by GP uncertainties, we consider a 32-atom bulk aluminum system initialized in the FCC phase at low temperature, with $\mathcal{N}_{\text{added}} = 1$ local environment added to the training set whenever the epistemic uncertainty on a force component exceeds the current noise uncertainty, $\sigma_{\text{thresh}} = \sigma_n$. As shown in Fig. 3a, DFT is called often at the beginning of the simulation as the GP model learns a force field suitable for FCC aluminum. After about 30 time steps, the model needs far fewer new training points, requiring fewer than 50 DFT calls in the first 5 ps of the simulation. To test the model's ability to adapt to changing conditions, the crystal is melted at time $t = 5 \text{ ps}$ by rescaling the velocities of the atoms to give the system an instantaneous temperature of 10^4 K , well above the experimental melting point of aluminum (933 K) due to the strong finite size effects of the $2 \times 2 \times 2$ supercell. The subsequent temperature in the remaining 5 ps of the simulation stabilizes around 5000 K with a radial distribution function consistent with the liquid phase (Fig. 3c). As shown in Fig. 3b, which plots the cumulative number of DFT calls made during the training run, the GP model makes frequent calls to DFT immediately after the crystal melts, as the local environments in the liquid phase of aluminum are significantly different from the previous solid-state training environments. The noise uncertainty σ_n of the model, shown in red in Fig. 3b, sharply increases as the system enters the liquid phase, reflecting the fact that it is more difficult to model, involving more diverse local environments and significantly larger force fluctuations. Because the error threshold σ_{thresh} is set equal to the optimized noise uncertainty σ_n , the threshold in the liquid phase is higher, and as a result the GP model requires a roughly similar number of DFT calls to learn the solid and liquid phases. Fewer than 100 calls are needed in total during the 10 ps of dynamics, with the majority of DFT calls made at the beginning of the simulation and immediately after melting.

The obtained force field is validated by testing the model on two independent 10-ps ab initio molecular dynamics (AIMD)

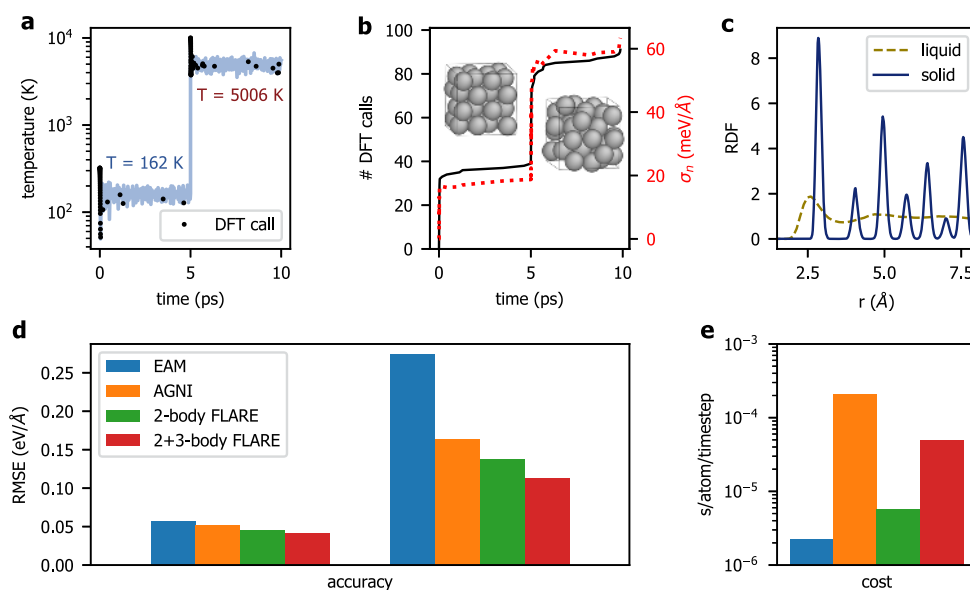


Fig. 3 Active learning of a multi-phase aluminum force field. **a** Instantaneous temperature during a 10-ps on-the-fly MD trajectory generated with the FLARE learning algorithm. The simulation begins in the FCC phase at low temperature and is melted at $t = 5$ ps. When the epistemic uncertainty σ_{ep} on a force component rises above the current noise uncertainty σ_n of the model, DFT is called (black dots). **b** The number of DFT calls (solid) and optimized noise uncertainty (dotted) throughout the simulation. A sharp increase is observed when the crystal is melted, illustrating the model's ability to actively learn the liquid phase. **c** During the first 5 ps of the simulation, the radial distribution function (RDF) is consistent with that of an fcc crystal (solid line), while in the final half of the simulation, the system exhibits an RDF characteristic of the liquid phase (dashed). **d** RMSE on AIMD forces of a tabulated version of the resulting force field compared with EAM, AGNI, and a tabulated 2-body FLARE force field. **e** Computational cost of LAMMPS implementations of these four force fields on a single CPU core in s/atom/timestep.

simulations of the solid and liquid phases of aluminum. 100 structures were sampled from the AIMD trajectories with 0.1-ps spacing between structures. Force predictions on all test structures were obtained with a tabulated version of the GP force field of Fig. 3a and compared against the corresponding DFT values, with the RMSE in eV/Å plotted in Fig. 3d. For reference, the models are compared against recent EAM and AGNI ML force fields, which were also trained on plane-wave DFT calculations with GGA exchange-correlation functionals and PAW pseudopotentials^{12,39}, though we note that they were not trained on exactly the same DFT calculations as our models. Also included for comparison is the performance of a 2-body FLARE model trained on the same local environments as the 2+3-body model. Each force field was tested on the same structures, with the FLARE force field reaching the lowest force errors for both trajectories. This is due in part to the fact that FLARE optimizes the force field for the specific simulation of interest, only augmenting the training set when necessary. This bypasses the need to anticipate all possible phases which a system might explore when creating the force field. To assess computational efficiency, 1000 MD steps were performed with the LAMMPS implementations of these four force fields on a single CPU core for a system of 1372 bulk Al atoms, with the cost of each force field plotted in Fig. 3e in s/atoms/timestep. The cost of the current LAMMPS implementation of the tabulated 2-body FLARE force field is found to be 5.6×10^{-6} s/atom/timestep, which is the same order of magnitude as the EAM cost of 2.2×10^{-6} s/atom/timestep. The 2+3-body model is about an order of magnitude slower at 4.9×10^{-5} s/atom/timestep, but still faster than AGNI, which directly predicts forces with a small neural network. This makes FLARE considerably less expensive than many-body models like GAP, with the cost of the recent GAP silicon model reported as 0.1 s/atom/timestep⁸.

Bulk vacancy and surface adatom diffusion

We next demonstrate that FLARE can be used to train force fields that dramatically accelerate simulations of rare-event dynamics over timescales spanning hundreds of picoseconds by applying the method to aluminum bulk vacancy diffusion and surface adatom diffusion. For bulk vacancy training, a 1-ns simulation was initialized by removing one atom from an equilibrium 32-atom FCC structure and setting the instantaneous initial temperature to 1500 K, giving a mean temperature of 734 K across the simulation. The GP model was constructed with a 2-body kernel with cutoff $r_{\text{cut}}^{(2)} = 5.4 \text{ \AA}$, resulting in a final optimized noise uncertainty of $\sigma_n = 70.2 \text{ meV/\AA}$. Discarding the 3-body contribution was found to significantly accelerate the simulation while still achieving low force errors due to the simplicity of the single-defect bulk crystalline phase, opening up nanosecond timescales during training. As shown in Fig. 4a, most DFT calls are made early on in the simulation, and after the first ~400ps, no additional DFT calls are required. The model predicts vacancy hops every few hundred picoseconds, which appear as sharp jumps in the mean squared displacement plotted in Fig. 4a. To check the accuracy of the underlying energy model of the GP, DFT energies were computed along the high symmetry transition path sketched in the inset of Fig. 4b, with a nearest neighbor migrating into the vacancy while all other atoms in the simulation cell were kept frozen at their fcc lattice sites. GP forces and energies along the transition path were evaluated to give an estimate of the energy barrier, showing close agreement with the ab initio DFT values (Fig. 4c), with the DFT forces lying within one standard deviation of the GP force predictions (Fig. 4b). The entire FLARE training run, including DFT calculations, GP hyperparameter optimization, force evaluations and MD updates, were performed on a 32-core machine in 68.8 h of wall time. Individual DFT calls required over a minute of wall time on average, making FLARE over 300 times faster than an

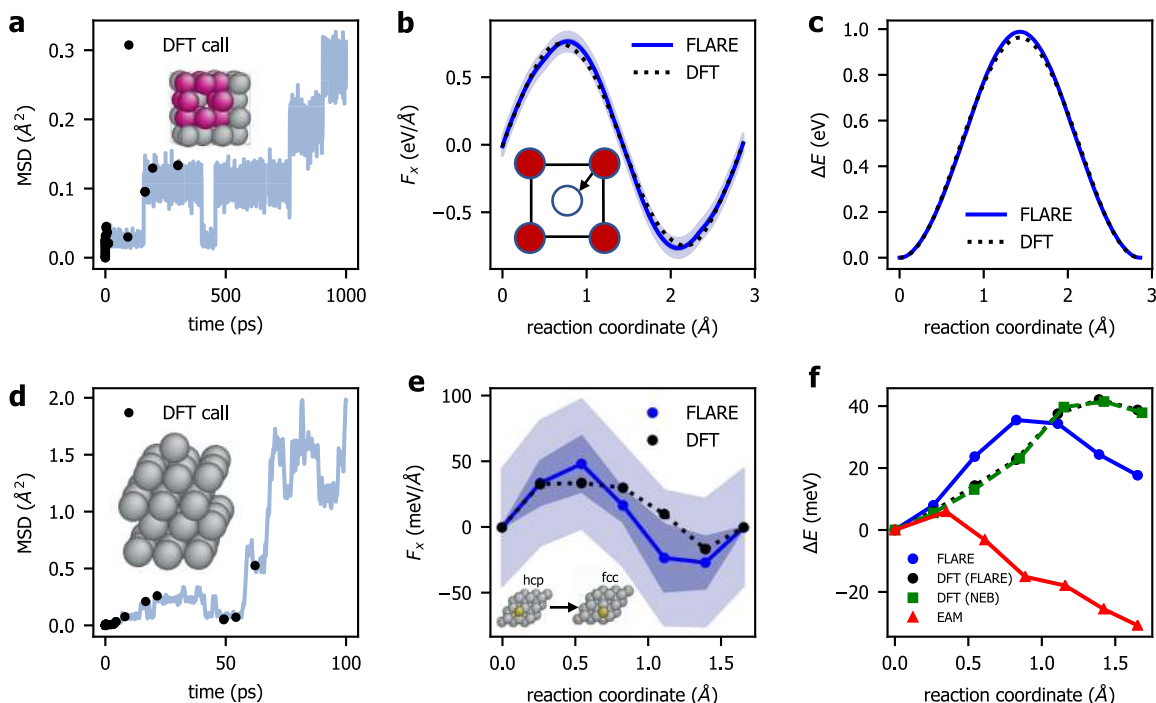


Fig. 4 On-the-fly learning of vacancy and adatom diffusion in aluminum. **a** Mean squared displacement during a FLARE training run of duration 1 ns for a 31-atom fcc aluminum structure with a vacancy (see inset). The majority of DFT calls occur at the beginning of the run, with no additional calls required after the first 400 ps. **b** x-component of the force predicted by the resulting force field for vacancy migration along a high symmetry transition path (see inset), in close agreement with the ab initio barrier (dotted). **c** Predicted energies along the transition path (blue) compared with DFT (black). **d** On-the-fly learning of adatom diffusion on a (111) aluminum surface, with sharp jumps in the MSD signaling movement of the adatom on the surface. **e** x-component of the force on the adatom in nudged elastic band images of the hcp-to-fcc transition computed with the trained GP model (blue line), along with the epistemic uncertainty σ_{ix} (dark blue) and total uncertainty $\sqrt{\sigma_{ix}^2 + \sigma_n^2}$ (light blue). DFT forces are computed for each image (black). **f** Predicted energies of the FLARE-generated NEB images (blue) relative to the first image. DFT energies are computed for each image (black), showing good agreement with the energies of an independent NEB calculation performed with DFT (green). NEB images with the EAM force field from³⁹ are shown for comparison (red).

equivalent AIMD run (see Supplementary Table 2 for a breakdown of GP prediction costs).

To test the accuracy of FLARE on a subtler transition with a significantly lower energy barrier, we consider aluminum adatom diffusion on a four-layer (111) aluminum slab, with a representative structure shown in the inset of Fig. 4d. As revealed in previous ab initio studies, an isolated Al adatom on the (111) Al surface exhibits a small but surprising preference for the hcp site^{40,41}, making this system an interesting and challenging test for a machine learned force field. For this system, 3-body contributions were found to considerably increase the accuracy of the force field, with a 7-Å 2-body cutoff and 4.5 Å 3-body cutoff giving an optimized noise uncertainty of $\sigma_n = 44.2$ meV/Å after the final DFT call at $t = 62.2$ ps (Fig. 4d). To validate the energetics of the force field, a 7-image nudged elastic band (NEB) calculation characterizing the transition from the hcp to fcc adatom sites was performed using the Atomic Simulation Environment⁴² with the GP energy predictions shown in blue in Fig. 4f. The DFT energies of each image of the NEB calculation are shown in black, showing agreement to within ≈ 20 meV for each image and confirming the GP's prediction of a slight energetic preference for the hcp site in equilibrium, which is not reproduced by the EAM model of Sheng et al.³⁹ (red line in Fig. 4f). An independent DFT NEB calculation was performed for the same transition, showing good agreement with the DFT energies of the FLARE NEB images.

Fast-ion diffusion in AgI

As a third and more challenging example of diffusion, we apply FLARE to the fast-ion conductor silver iodide (AgI), which exhibits a structural phase transition at 420 K from the low-temperature

γ/β -phase to a cubic "superionic" α -phase, with silver ions in the α -phase observed to have a liquid-like diffusivity⁴³. A 2+3-body FLARE model was trained in a 15 ps on-the-fly simulation of 48 AgI atoms in the α -phase, with the temperature increased at 5 and 10 ps (Fig. 5a). The uncertainty threshold was set to twice the noise uncertainty, $\sigma_{\text{thresh}} = 2\sigma_n$, making the model slightly less sensitive to changing temperature and contributing to the 1-ps delay observed between the temperature increase at 5 ps and the next call to DFT at $t = 6.121$ ps. Thirty-nine calls to DFT were made in total, with the $N_{\text{added}} = 10$ highest uncertainty local environments added to the training set after each DFT calculation.

After training, the model was mapped to a tabulated cubic spline model in LAMMPS, which was used to perform 1-ns simulations at zero pressure and fixed temperature, with each simulation requiring about three hours of wall time on 32 cpu cores ($\approx 3.2 \times 10^{-5}$ cpu · s/atom/timestep). Ten MD simulations were performed in total with temperatures ranging from 200 to 650 K in intervals of 50 K. In each simulation, the system was initialized in a pristine $14 \times 14 \times 14$ α -phase supercell (10,976 atoms total), with the silver ions placed at the energetically preferred tetragonal interstices of the bcc iodine sublattice. The diffusion coefficients of the Ag ions are plotted in Fig. 5b, showing a sharp increase between 400 K and 450 K, in good agreement with the experimental fast-ion transition temperature of 420 K. The diffusion coefficients are compared with an AIMD study of the α -phase of AgI⁴⁴, which used a similar exchange-correlation functional, showing excellent agreement at 450 K and above. Both FLARE and AIMD show good agreement with experimentally observed α -phase Ag diffusion coefficients⁴⁵, with a slight vertical offset but comparable activation energies of 0.107, 0.114, and

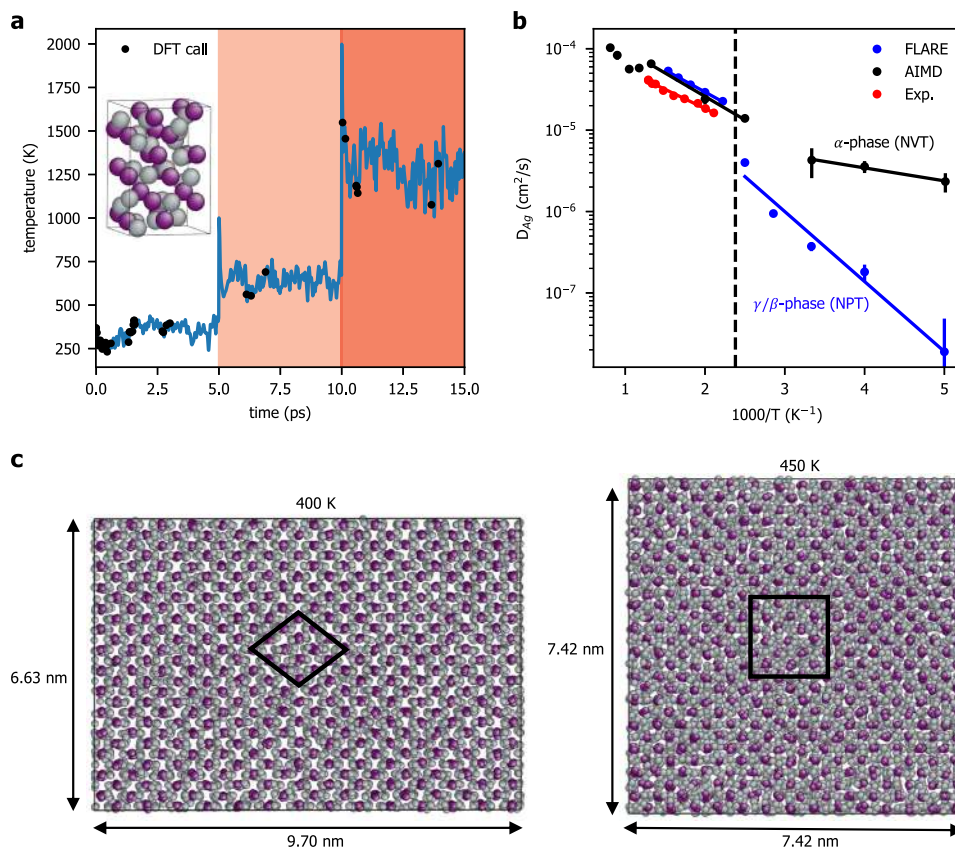


Fig. 5 On-the-fly learning of fast-ion diffusion in silver iodide. **a** Temperature during a FLARE training simulation of duration 15 ps for a 48-atom silver iodide structure in the α -phase (see inset), with the instantaneous temperature of the simulation increased at 5 ps and 10 ps. **b** Silver diffusion coefficients (blue dots) computed with a tabulated version of the resulting force field from 1 ns NPT simulations of 10,976 AgI atoms. The computed coefficients for the α -phase of AgI are in good agreement with the AIMD simulations of Wood and Marzari⁴⁴ (black) and experimental data reported by Kvist and Tärneberg⁴⁵ (red), with the fast-ion phase transition at 420 K (dashed line) correctly modeled. **c** Structures drawn from simulations at 400 K (left) and 450 K (right), illustrating the solid-solid structural phase transition that occurs between these temperatures.

0.093 eV for FLARE, AIMD, and experiment, respectively. Below the transition temperature, the FLARE force field correctly predicts a phase transition to a non-diffusive and non-cubic hcp phase with a nearest neighbor I-I coordination of 12, consistent with the γ and β phases of AgI⁴⁶. This accounts for the discrepancy between the FLARE and AIMD diffusion coefficients in the low-temperature regime, as the latter simulations were conducted in the α -phase with a fixed cubic cell. Example structures from the 400 and 450 K FLARE MD simulations are illustrated in Fig. 5c, with the low-temperature structure giving a c/a ratio of 1.46 and the high-temperature structure having a lattice parameter of $a_{\text{lat}} = 5.30 \text{ \AA}$, in fair agreement with the corresponding experimental values of $c/a = 1.63$ and $a_{\text{lat}} = 5.07 \text{ \AA}$ near these temperatures (for the β - and α -phases, respectively)⁴⁷.

General applicability

Finally, we demonstrate in Fig. 6 that FLARE can be widely applied to diverse systems, including covalently bonded insulators and semiconductors, as well as oxides, alloys, and two-dimensional materials. FLARE training runs were performed for five representative systems—carbon, silicon, aluminum oxide, nickel titanium, and two-dimensional boron nitride—with the instantaneous temperature of each system rescaled at $t = 5 \text{ ps}$ to illustrate the model's ability to detect and adapt to novel local environments (see the left half of Table 1 for training details). To accelerate training of the nickel titanium model, which required expensive DFT calculations, the error threshold was set to twice the noise

uncertainty, $\sigma_{\text{thresh}} = 2\sigma_{nr}$, significantly reducing the total number of DFT calls needed to ~ 20 (as shown in Fig. 6d and Table 1). Adding multiple local environments to the training set after each DFT call also had the effect of reducing the total number of DFT calls needed, as apparent in the aluminum oxide training run, for which $\mathcal{N}_{\text{added}} = 30$ local environments were added after every DFT call and only 16 DFT calls were needed in total to train the model. Each training run was performed on a 32-core machine and took between 11.3 and 64.4 h of wall time (for silicon and carbon, respectively). We emphasize that the training procedure for each material is fully automatic, with the training set and hyperparameters updated on the fly without any human guidance.

To validate the models, independent NVE molecular dynamics trajectories of duration 10 ps were generated with each GP force field, with DFT calculations performed for ten MD frames spaced equally across the simulation and compared against the corresponding GP predictions. We find low root mean squared errors (RMSE) of around 0.1 eV/ \AA for four of the five systems, and for carbon we find a RMSE of 0.42 eV/ \AA due to the much higher temperature of the carbon validation run. The RMSE over all force component predictions in the ten representative frames is reported in Table 1. In order to illustrate the range of force magnitudes present in the simulation, we also report the 95th percentile of the absolute force components in these frames, with the ratio of the two reported in the final column of Table 1. The resulting ratios lie between 3 and 10%, similar to the ratios

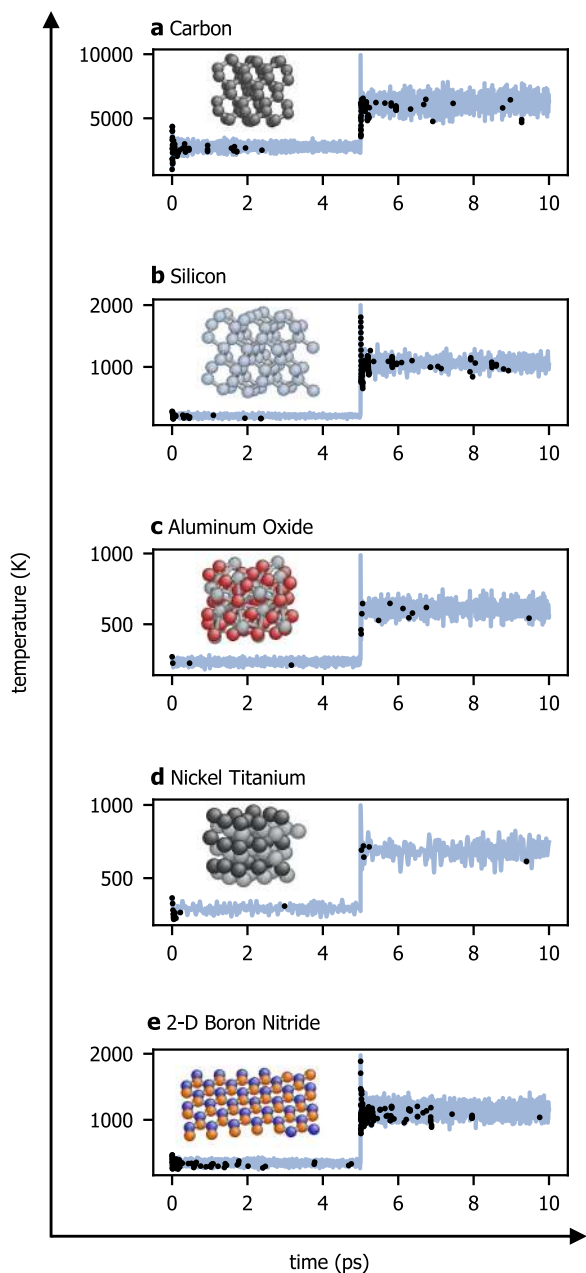


Fig. 6 On-the-fly force field learning applied to a range of single- and multi-element systems. In each training run, the instantaneous temperature (blue) was increased at time $t = 5.0$ ps, triggering DFT calls and updates to the GP model (black dots) caused by model detection of novel local environments. Example structures from each simulation are shown in the insets.

reported in a recent study of amorphous carbon with a Gaussian approximation potential⁴.

DISCUSSION

In summary, we have presented a method for automatically training low-dimensional Gaussian process models that provide accurate force estimates and reliable internal estimates of model uncertainty. The model's uncertainties are shown to correlate well with true out-of-sample error, providing an interpretable, principled basis for active learning of a force field model during molecular dynamics. The nonparametric 2- and 3-body FLARE models described here require fewer training environments than

high-dimensional machine learning approaches, and are therefore well-suited to settings where large databases of ab initio calculations are too expensive to compute. Our models have a simple, accurate, and physically interpretable underlying energy model, which we have shown can be used to map the GP to a faster regression model approaching the speed of a classical force field. This provides a path toward force fields tailored to individual applications that give good agreement with DFT at several orders of magnitude lower computational cost, which we expect to considerably expand the range of materials that can be accurately studied with atomistic simulation. Particularly promising is the application of the FLARE framework to dynamical systems dominated by rare diffusion or reaction events that are very difficult to treat with existing ab initio, classical force field, or machine learning methods.

Extending this active learning method to complex systems like polymers and proteins is an important open challenge. The Bayesian force fields presented here may serve as a useful guide for selecting small, uncertain fragments from these systems that can then be evaluated with DFT to refine the force field, similar to other recent approaches that train on small portions of larger structures^{48,49}. This may provide a path toward accurate machine learned force fields for chemical and biological systems that are currently outside the reach of DFT and other quantum mechanical methods.

METHODS

Gaussian process force fields

As observed by Glielmo et al.^{38,50,51}, the task of fitting a force field can be dramatically simplified by assuming that only small clusters of atoms in the local environment of an atom i contribute to its local energy E_i . We define the n -body local environment $\rho_i^{(n)}$ of atom i to be the set of atoms within a cutoff distance $r_{\text{cut}}^{(n)}$ from atom i , and a cluster of n atoms to be the atom i and $n - 1$ of the atoms in $\rho_i^{(n)}$. The energy $\epsilon_{\mathbf{s}_{i,j_1,\dots,j_{n-1}}}(\mathbf{d}_{i,j_1,\dots,j_{n-1}})$ of each cluster of n atoms is assumed to depend on the species of the atoms in the cluster, $\mathbf{s}_{i,j_1,\dots,j_{n-1}} = (s_i, s_{j_1}, \dots, s_{j_{n-1}})$, and on a corresponding vector of interatomic distances between the atoms, $\mathbf{d}_{i,j_1,\dots,j_{n-1}}$. For example, for clusters of two atoms, this vector consists of a single scalar, $\mathbf{d}_{i,j_1} = (r_{i,j_1})$, where r_{i,j_1} is the distance between the central atom i and atom j_1 , and for clusters of three atoms, $\mathbf{d}_{i,j_1,j_2} = (r_{i,j_1}, r_{i,j_2}, r_{j_1,j_2})$. The local energy assigned to atom i may then be written as

$$E_i = \sum_{n=2}^N \sum_{i_{n-1} > \dots > i_1 \in \rho_i^{(n)}} \epsilon_{\mathbf{s}_{i,i_1,\dots,i_{n-1}}}(\mathbf{d}_{i,i_1,\dots,i_{n-1}}), \quad (1)$$

where the outer sum ranges over each n -body contribution to the energy up to a chosen maximum order N and the inner sum ranges over all clusters of n atoms inside the n -body environment $\rho_i^{(n)}$. The regression task is to learn the functions $\epsilon_{\mathbf{s}_{i,i_1,\dots,i_{n-1}}}(\mathbf{d}_{i,i_1,\dots,i_{n-1}})$, which for small n have much lower dimensionality than the full potential energy surface.

To learn the cluster contributions $\epsilon_{\mathbf{s}_{i,i_1,\dots,i_{n-1}}}$, we use ab initio force data to construct Gaussian process (GP) models, an established Bayesian approach to describing probability distributions over unknown functions²³. In GP regression, the covariance between two outputs of the unknown function is related to the degree of similarity of the inputs as quantified by a kernel function. For our GP force fields, the covariance between n -body energy contributions ($\epsilon_{\mathbf{s}_{i,i_1,\dots,i_{n-1}}}$ in Eq. (1)) is equated to a kernel function k_n that directly compares the interatomic distance vectors while preserving rotational invariance. The local energy kernel between two local environments ρ_i, ρ_j is expressed as a sum over kernels between clusters of atoms,

$$k(\rho_i, \rho_j) = \sum_{n=2}^N \sum_{i_{n-1} > \dots > i_1 \in \rho_i^{(n)}} \sum_{j_{n-1} > \dots > j_1 \in \rho_j^{(n)}} \delta_{\mathbf{s}_{i,i_1,\dots,i_{n-1}}, \mathbf{s}_{j,j_1,\dots,j_{n-1}}} k_n(\mathbf{d}_{i,i_1,\dots,i_{n-1}}, \mathbf{P}_n \mathbf{d}_{j,j_1,\dots,j_{n-1}}). \quad (2)$$

Importantly, this kernel function explicitly distinguishes between distinct species, with the delta function δ evaluating to 1 if the species vectors $\mathbf{s}_{i,i_1,\dots,i_n}$ of the clusters under comparison are equal and 0 otherwise. The

Table 1. Training and validation details of the FLARE models shown in Fig. 6.

	Training							Validation			
	$\mathcal{N}_{\text{atoms}}$	r_2 (Å)	r_3 (Å)	σ_{thresh}	$\mathcal{N}_{\text{added}}$	\mathcal{N}_{DFT}	t_{wall} (hours)	T (K)	RMSE (eV/Å)	P_{95} (eV/Å)	Ratio
C	64	4.0	2.75	σ_n	1	107	11.3	3710	0.42	7.45	0.056
Si	64	6.0	4.2	σ_n	5	133	64.4	620	0.077	1.54	0.050
Al ₂ O ₃	80	4.5	3.5	σ_n	30	16	20.6	533	0.14	1.85	0.076
NiTi	54	4.48	3.2	$2\sigma_n$	10	18	33.3	510	0.10	1.00	0.104
BN	72	5.1	4.0	σ_n	1	237	31.4	677	0.092	3.24	0.029

Training: $\mathcal{N}_{\text{atoms}}$ is the number of atoms in the training simulation, r_2 and r_3 are the 2- and 3-body cutoffs of the GP models, σ_{thresh} is the uncertainty threshold that determines when DFT is called, $\mathcal{N}_{\text{added}}$ is the number of local environments added each time DFT is called, and t_{wall} is the total wall time of the training simulation. Validation: T is the mean temperature during the validation simulation, RMSE is the root mean squared error on ten snapshots from the validation run, and P_{95} is the 95th percentile of force components in these 10 snapshots. The ratio between the RMSE and P_{95} is reported in the final column.

innermost sum of Eq. (2) is over all permutations \mathbf{P}_n of indices of the species and distance vectors of the second cluster, guaranteeing invariance of the model under permutation of atoms of the same species. The resulting force kernel describing the covariance between force components is obtained by differentiating the local energy kernel with respect to the Cartesian coordinates $\vec{r}_{i\alpha}$, $\vec{r}_{j\beta}$ of the central atoms of ρ_i and ρ_j ,

$$k_{\alpha,\beta}(\rho_i, \rho_j) = \frac{\partial^2 k(\rho_i, \rho_j)}{\partial \vec{r}_{i\alpha} \partial \vec{r}_{j\beta}}, \quad (3)$$

giving an exactly rotationally covariant and energy conserving model of interatomic forces^{5,38,50}. For completeness, we provide in Supplementary Table 4 the formulas involved in computing the 3-body derivative kernel described by Eq. (3), along with its derivatives with respect to the hyperparameters of the kernel, which are used to calculate the gradient of the log marginal likelihood during hyperparameter optimization.

In this work, we choose $N=3$, restricting the sum to 2- and 3-body contributions, as we have found the resulting GP models to be sufficiently expressive to describe with high accuracy a range of single- and multi-element systems while remaining computationally efficient. This is consistent with the findings of Glielmo et al.³⁸, which compared the performance of 2-, 3-, and many-body kernels and found that many-body models required substantially more training data while only modestly improving performance for several crystals, nanoclusters, and amorphous systems. Further investigation of model accuracy as a function of the maximum order N of the kernel for different types of materials is an interesting area for future study, as it may provide a systematic data-driven approach to characterizing many-body interactions in complex materials.

For the pair and triplet kernels k_2 and k_3 , we choose the squared exponential kernel multiplied by a smooth quadratic cutoff function f_{cut} that ensures the model is continuous as atoms enter and exit the cutoff sphere,

$$k_2(r_{i,i_1}, r_{j,j_1}) = \sigma_{s,2}^2 \exp\left(-\frac{(r_{i,i_1}-r_{j,j_1})^2}{2\ell_{(2,3)}^2}\right) f_{\text{cut}}(r_{i,i_1}, r_{j,j_1}), \quad (4)$$

$$k_3(\mathbf{d}_{i,i_1,j_2}, \mathbf{d}_{j,j_1,j_2}) = \sigma_{s,3}^2 \exp\left(-\frac{\|\mathbf{d}_{i,i_1,j_2}-\mathbf{d}_{j,j_1,j_2}\|^2}{2\ell_{(2,3)}^2}\right) f_{\text{cut}}(\mathbf{d}_{i,i_1,j_2}, \mathbf{d}_{j,j_1,j_2}),$$

where $\sigma_{s,(2,3)}$ is the signal variance related to the maximum uncertainty of points far from the training set, $\ell_{(2,3)}$ is the length scale of the 2- and 3-body contributions, and $\|\cdot\|$ denotes the vector 2-norm.

The force component $f_{i\alpha}$ on each atom i and the square of the epistemic uncertainty $\sigma_{i\alpha}^2$ assigned to that force component are computed using the standard GP relations²³,

$$f_{i\alpha} = \bar{k}_{i\alpha}^T (K + \sigma_n^2 I)^{-1} \bar{y} \\ \sigma_{i\alpha}^2 = k_{\alpha,\alpha}(\rho_i, \rho_i) - \bar{k}_{i\alpha}^T (K + \sigma_n^2 I)^{-1} \bar{k}_{i\alpha}, \quad (5)$$

where $\bar{k}_{i\alpha}$ is the vector of force kernels between ρ_i and the local environments in the training set, i.e. $\bar{k}_{i\alpha,\beta} = k_{\alpha,\beta}(\rho_i, \rho_j)$, K is the covariance matrix $K_{m\alpha,n\beta} = k_{\alpha,\beta}(\rho_m, \rho_n)$ of the training points, \bar{y} is the vector of forces acting on the atoms in the training set, and σ_n is a hyperparameter that characterizes observation noise. The total uncertainty on the force component, corresponding to the variance of the predictive posterior distribution of the predicted value, is obtained by adding σ_n^2 , the square of the noise uncertainty²³. Notice that the square of the epistemic uncertainty

is bounded above by $k_{\alpha,\alpha}(\rho_i, \rho_i)$, which for our kernel function is determined by the signal variances $\sigma_{s,2}^2$ and $\sigma_{s,3}^2$.

In all models in this work, the hyperparameters $\theta = \{\sigma_2, \sigma_3, \ell_2, \ell_3, \sigma_n\}$ are optimized with SciPy's implementation of the BFGS algorithm⁵² by maximizing the log marginal likelihood of the training data $\rho = \{\rho_1, \rho_2, \dots, \rho_n\}$, which takes the form²³

$$\log p(\bar{y}|\rho, \theta) = -\frac{1}{2} \bar{y}^T (K + \sigma_n^2 I)^{-1} \bar{y} - \frac{1}{2} \log |K + \sigma_n^2 I| - \frac{n}{2} \log 2\pi. \quad (6)$$

To efficiently maximize this quantity with BFGS, the gradient with respect to all hyperparameters is calculated with the analytic expression²³,

$$\frac{\partial}{\partial \theta_i} \log p(\bar{y}|\rho, \theta) = \frac{1}{2} \text{tr} \left((\bar{a} \bar{a}^T - K_y^{-1}) \frac{\partial K_y}{\partial \theta_i} \right), \quad (7)$$

where $\bar{a} = K_y^{-1} \bar{y}$ and $K_y = K + \sigma_n^2 I$. The formulas for the kernel derivatives with respect to the hyperparameters that appear in this expression, $\frac{\partial k_{\alpha,\beta}}{\partial \theta_i}$, can be exactly calculated, and we list them in Supplementary Table 4 for the case of the 3-body kernel. The BFGS algorithm is terminated once the log marginal likelihood gradient falls below a threshold value $\epsilon = 10^{-4}$. Note that computation of the log marginal likelihood and its gradient involves inverting the covariance matrix K and is efficient if the model is trained on fewer than ~ 1000 points. This data-driven approach to selecting model hyperparameters stands in contrast to other GP force fields, in which hyperparameters are chosen heuristically⁴.

Mapping to tabulated spline models

As shown by Glielmo et al.³⁸ for single-element systems, GP models built on n -body kernels can be mapped to efficient cubic spline models, eliminating the expensive loop over training points involved in the calculation of the kernel vector $\bar{k}_{i\alpha}$ in Eq. (5). We have extended this mapping procedure to our multi-element kernels by constructing cubic spline interpolants for each n -body force contribution $-\frac{d}{d\mathbf{r}_i} \epsilon_{s_i, j_1, \dots, j_{n-1}}(\mathbf{d}_{i, j_1, \dots, j_{n-1}})$. The 2- and 3-body contributions require 1- and 3-dimensional cubic splines, respectively. The resulting spline model can be made arbitrarily accurate relative to the original GP model by increasing the number of control points of the spline. In Supplementary Table 3, we report the grid of control points used for each mapped force field in this work.

Computational details

All DFT calculations were performed using Quantum Espresso 6.2.1, with pseudopotentials, k-point meshes, plane-wave energy cutoffs, and charge density energy cutoffs for all calculations reported in Supplementary Table 1. The on-the-fly learning algorithm is implemented with the FLARE package (<https://github.com/mir-group/flare>), which couples our Python-based MD and GP code with Quantum ESPRESSO⁵³. Kernel and distance calculations are accelerated with the open-source just-in-time compiler Numba to enable training simulations spanning hundreds of picoseconds⁵⁴. All on-the-fly molecular dynamics trajectories were performed in the NVE ensemble using the Verlet algorithm. LAMMPS simulations of AgI were performed in the NPT ensemble at zero pressure. Atomistic visualizations were created using Atomeye⁵⁵.

DATA AVAILABILITY

Input and output files for the on-the-fly training simulations discussed in this work as well as the final trained Gaussian process models are available at <https://doi.org/10.24435/materialscloud:2020.0017/v1>.

CODE AVAILABILITY

The FLARE code is available at <https://github.com/mir-group/flare>.

Received: 13 August 2019; Accepted: 13 February 2020;

Published online: 18 March 2020

REFERENCES

- Szlachta, W. J., Bartók, A. P. & Csányi, G. Accuracy and transferability of gaussian approximation potential models for tungsten. *Phys. Rev. B* **90**, 104108 (2014).
- Behler, J. Neural network potential-energy surfaces in chemistry: a tool for large-scale simulations. *Phys. Chem. Chem. Phys.* **13**, 17930–17955 (2011).
- 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**, 316–330 (2015).
- Deringer, V. L. & Csányi, G. Machine learning based interatomic potential for amorphous carbon. *Phys. Rev. B* **95**, 094203 (2017).
- Chmiela, S. et al. Machine learning of accurate energy-conserving molecular force fields. *Sci. Adv.* **3**, e1603015 (2017).
- Schütt, K. et al. Schnet: a continuous-filter convolutional neural network for modeling quantum interactions. *Adv. Neural Inf. Process. Syst.* **31**, 991–1001 (2017).
- Deringer, V. L., Pickard, C. J. & Csányi, G. Data-driven learning of total and local energies in elemental boron. *Phys. Rev. Lett.* **120**, 156001 (2018).
- Bartók, A. P., Kermode, J., Bernstein, N. & Csányi, G. Machine learning a general-purpose interatomic potential for silicon. *Phys. Rev. X* **8**, 041048 (2018).
- 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).
- Zhang, L. et al. End-to-end symmetry preserving inter-atomic potential energy model for finite and extended systems. *Adv. Neural Inf. Process. Syst.* **32**, 4441–4451 (2018).
- Smith, J. S. et al. Approaching coupled cluster accuracy with a general-purpose neural network potential through transfer learning. *Nat. Commun.* **10**, 1–8 (2019).
- Botu, V., Batra, R., Chapman, J. & Ramprasad, R. Machine learning force fields: construction, validation, and outlook. *J. Phys. Chem. C* **121**, 511–522 (2016).
- Podryabinkin, E. V. & Shapeev, A. V. Active learning of linearly parametrized interatomic potentials. *Comput. Mater. Sci.* **140**, 171–180 (2017).
- Mishra, A. et al. Multiobjective genetic training and uncertainty quantification of reactive force fields. *npj Comput. Mater.* **4**, 42 (2018).
- Janet, J. P., Duan, C., Yang, T., Nandy, A. & Kulik, H. A quantitative uncertainty metric controls error in neural network-driven chemical discovery. *Chem. Sci.* **10**, 7913–7922 (2019).
- Musil, F., Willatt, M. J., Langovoy, M. A. & Ceriotti, M. Fast and accurate uncertainty estimation in chemical machine learning. *J. Chem. Theory Comput.* **15**, 906–915 (2019).
- Zhang, L., Lin, D.-Y., Wang, H., Car, R. & Weinan, E. Active learning of uniformly accurate interatomic potentials for materials simulation. *Phys. Rev. Mater.* **3**, 023804 (2019).
- Podryabinkin, E. V., Tikhonov, E. V., Shapeev, A. V. & Oganov, A. R. Accelerating crystal structure prediction by machine-learning interatomic potentials with active learning. *Phys. Rev. B* **99**, 064114 (2019).
- Behler, J. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J. Chem. Phys.* **134**, 074106 (2011).
- Bartók, A. P., Kondor, R. & Csányi, G. On representing chemical environments. *Phys. Rev. B* **87**, 184115 (2013).
- 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).
- Bartók, A. P. & Csányi, G. Gaussian approximation potentials: a brief tutorial introduction. *Int. J. Quant. Chem.* **115**, 1051–1057 (2015).
- Williams, C. K. I. & Rasmussen, C. E. *Gaussian Processes For Machine Learning*. Vol. 2. Chapter 2 (MIT Press Cambridge: MA, 2006).
- Quiñonero-Candela, J. & Rasmussen, C. E. A unifying view of sparse approximate gaussian process regression. *J. Mach. Learn. Res.* **6**, 1939–1959 (2005).
- Li, Z., Kermode, J. R. & De Vita, A. Molecular dynamics with on-the-fly machine learning of quantum-mechanical forces. *Phys. Rev. Lett.* **114**, 096405 (2015).
- Stillinger, F. H. & Weber, T. A. Computer simulation of local order in condensed phases of silicon. *Phys. Rev. B* **31**, 5262 (1985).
- Bernstein, N., Csányi, G. & Deringer, V. L. De novo exploration and self-guided learning of potential-energy surfaces. *npj Comput. Mater.* **5**, 1–9 (2019).
- Rupp, M. et al. Machine learning estimates of natural product conformational energies. *PLoS Comput. Biol.* **10**, e1003400 (2014).
- Uteva, E., Graham, R. S., Wilkinson, R. D. & Wheatley, R. J. Active learning in gaussian process interpolation of potential energy surfaces. *J. Chem. Phys.* **149**, 174114 (2018).
- Smith, J. S., Nebgen, B., Lubbers, N., Isayev, O. & Roitberg, A. E. Less is more: sampling chemical space with active learning. *J. Chem. Phys.* **148**, 241733 (2018).
- Li, Z. *On-the-fly Machine Learning of Quantum Mechanical Forces and Its Potential Applications for Large Scale Molecular Dynamics*. Ph.D. thesis, King's College, London (2014).
- Shapeev, A. V. Moment tensor potentials: a class of systematically improvable interatomic potentials. *Multiscale Model. Simul.* **14**, 1153–1173 (2016).
- Gubaev, K., Podryabinkin, E. V., Hart, G. L. W. & Shapeev, A. V. Accelerating high-throughput searches for new alloys with active learning of interatomic potentials. *Comput. Mater. Sci.* **156**, 148–156 (2019).
- Gubaev, K., Podryabinkin, E. V. & Shapeev, A. V. Machine learning of molecular properties: locality and active learning. *J. Chem. Phys.* **148**, 241727 (2018).
- Jinnouchi, R., Karsai, F. & Kresse, G. On-the-fly machine learning force field generation: application to melting points. *Phys. Rev. B* **100**, 014105 (2019).
- Jinnouchi, R., Lahnsteiner, J., Karsai, F., Kresse, G. & Bokdam, M. Phase transitions of hybrid perovskites simulated by machine-learning force fields trained on the fly with bayesian inference. *Phys. Rev. Lett.* **122**, 225701 (2019).
- De, S., Bartók, A. P., Csányi, G. & Ceriotti, M. Comparing molecules and solids across structural and alchemical space. *Phys. Chem. Chem. Phys.* **18**, 13754–13769 (2016).
- Glielmo, A., Zeni, C. & De Vita, A. Efficient nonparametric n-body force fields from machine learning. *Phys. Rev. B* **97**, 184307 (2018).
- Sheng, H. W., Kramer, M. J., Cadien, A., Fujita, T. & Chen, M. W. Highly optimized embedded-atom-method potentials for fourteen fcc metals. *Phys. Rev. B* **83**, 134118 (2011).
- Stumpf, R. & Scheffler, M. Theory of self-diffusion at and growth of Al (111). *Phys. Rev. Lett.* **72**, 254 (1994).
- Stumpf, R. & Scheffler, M. Ab initio calculations of energies and self-diffusion on flat and stepped surfaces of Al and their implications on crystal growth. *Phys. Rev. B* **53**, 4958 (1996).
- Larsen, A. H. et al. The atomic simulation environment—a python library for working with atoms. *J. Phys. Condens. Matter* **29**, 273002 (2017).
- Hull, S. Superionics: crystal structures and conduction processes. *Rep. Prog. Phys.* **67**, 1233 (2004).
- Wood, B. C. & Marzari, N. Dynamical structure, bonding, and thermodynamics of the superionic sublattice in α -AgI. *Phys. Rev. Lett.* **97**, 166401 (2006).
- Kvist, A. & Tärneberg, R. Self-diffusion of silver ions in the cubic high temperature modification of silver iodide. *Z. für Naturforsch.* **A 25**, 257–259 (1970).
- Parrinello, M., Rahman, A. & Vashishta, P. Structural transitions in superionic conductors. *Phys. Rev. Lett.* **50**, 1073 (1983).
- Madelung, O., Rössler, U., Schulz, M. *II-VI and I-VII Compounds; Semimagnetic Compounds*, 1–7 (Springer-Verlag, 1999).
- Gastegger, M., Behler, J. & Marquetand, P. Machine learning molecular dynamics for the simulation of infrared spectra. *Chem. Sci.* **8**, 6924–6935 (2017).
- Mailoa, J. P. et al. Fast neural network approach for direct covariant forces prediction in complex multi-element extended systems. *Nat. Mach. Intell.* **1**, 471–479 (2019).
- Glielmo, A., Sollich, P. & De Vita, A. Accurate interatomic force fields via machine learning with covariant kernels. *Phys. Rev. B* **95**, 214302 (2017).
- Glielmo, A. & Zeni, C. Building nonparametric n-body force fields using gaussian process regression, arXiv preprint arXiv:1905.07626 (2019).
- Jones, E. et al. SciPy. *Open Source Scientific Tools for Python* <http://www.scipy.org/> (2011).
- Giannozzi, P. et al. Quantum espresso: a modular and open-source software project for quantum simulations of materials. *J. Phys. Condens. Matter* **21**, 395502 (2009).
- Lam, S. K., Pitrou, A. & Seibert, S. Numba: a LLVM-based python jit compiler. In *Proc. Second Workshop on the LLVM Compiler Infrastructure in HPC 7* (ACM, 2015).
- Li, J. Atomeye: an efficient atomistic configuration viewer. *Model. Simul. Mater. Sci. Eng.* **11**, 173 (2003).

ACKNOWLEDGEMENTS

We thank Aldo Glielmo, Jin Soo Lim, Nicola Molinari, and Jonathan Mailoa for helpful discussions, and Anders Johansson, Kyle Bystrom, David Clark, and Blake Duschatko for contributions to the FLARE code. B.K. acknowledges generous gift funding support from Bosch Research and partial support from the National Science

Foundation under Grant No. 1808162. L.S. was supported by the Integrated Mesoscale Architectures for Sustainable Catalysis (IMASC), an Energy Frontier Research Center funded by the U.S. Department of Energy, Office of Science, Basic Energy Sciences under Award #DE-SC0012573. A.M.K. and S.B. acknowledge funding from the MIT-Skoltech Center for Electrochemical Energy Storage. S.B.T. is supported by the Department of Energy Computational Science Graduate Fellowship under grant DE-FG02-97ER25308.

AUTHOR CONTRIBUTIONS

J.V. conceived the study and is the primary developer of the FLARE code. Y.X. implemented the mapping of the GP models to cubic splines, and L.S. created the LAMMPS pair style. Y.X., L.S., S.B.T. and S.B. assisted with code development. B.K. supervised the work and contributed to algorithm development. J.V. wrote the manuscript. All authors contributed to manuscript preparation.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information is available for this paper at <https://doi.org/10.1038/s41524-020-0283-z>.

Correspondence and requests for materials should be addressed to J.V. or B.K.

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) 2020