Check for updates

Exploring inhomogeneous surfaces: Ti-rich SrTiO₃(110) reconstructions *via* active learning†

Ralf Wanzenböck, [ID] [a] Esther Heid, [ID] [a] Michele Riva, [ID] [b] Giada Franceschi, [ID] [b] Alexander M. Imre, [b] Jesús Carrete, [ID] [c] Ulrike Diebold[b] and Georg K. H. Madsen [ID] *[a]

The investigation of inhomogeneous surfaces, where various local structures coexist, is crucial for understanding interfaces of technological interest, yet it presents significant challenges. Here, we study the atomic configurations of the $(2 \times m)$ Ti-rich surfaces at (110)-oriented SrTiO₃ by bringing together scanning tunneling microscopy and transferable neural-network force fields combined with evolutionary exploration. We leverage an active learning methodology to iteratively extend the training data as needed for different configurations. Training on only small well-known reconstructions, we are able to extrapolate to the complicated and diverse overlayers encountered in different regions of the inhomogeneous SrTiO₃(110)-$(2 \times m)$ surface. Our machine-learning-backed approach generates several new candidate structures, in good agreement with experiment and verified using density functional theory. The approach could be extended to other complex metal oxides featuring large coexisting surface reconstructions.

1 Introduction

The surfaces of metal oxides play key roles in countless natural processes and technological applications. At the atomic level, the properties of metal oxide surfaces, such as reactivity, electronic structure, and defect formation, are intricately linked to their performance. Only through a deep understanding of these surfaces at the atomic scale, can their properties be precisely controlled and optimized.

Machine-learned force fields (MLFFs) are an increasingly popular tool that can be used to explore surface structures, including those of metal oxides. Method development and their applications have been advancing in parallel, with innovative and powerful models synergizing with established and proven methods. For example, moving from neural-network force fields that utilize precomputed invariant descriptors[1–3] to adopting equivariant message passing networks[4–6] has enabled more data-efficient and transferable MLFFs. Modern applications include foundation models trained on a wide range of materials,[7] transferable water potentials,[8] and condensed phase chemistry.[9]

Here, we utilize MLFFs to explore the surface reconstructions of strontium titanate (SrTiO₃), a perovskite oxide used as

a model system for the development of many technological applications, including optoelectronics, catalysis, memory devices, and photovoltaics.[10–13] SrTiO₃ is also exciting from a fundamental standpoint: It exemplifies the richness of bulk, surface, and interface properties that can be accessed within a single perovskite material: donor doping by chemical impurities,[14,15] oxygen vacancies,[14,16,17] or field effects[15,18] can turn it into an insulator, a metal, a superconductor or even induce confined metallic behaviour in the form of 2D electron gases. The diversity extends to the atomistic details of the surface,[19] where a variety of composition-related, polarity-compensating reconstructions have been found for the (001), (110), and (111) orientations.[20–27] Pinpointing the regions of stability of such reconstructions and gaining a deep understanding of the atomic-scale details of their surfaces is essential for designing SrTiO₃-based systems with tailored functionalities.

Many studies exist that explore the atomic details of the surface of single-crystalline SrTiO₃ samples under ultra-high vacuum (UHV) conditions. It is known that specific surface reconstructions are difficult to reproduce and can depend on sample history and preparation conditions.[19] Notably, scanning tunneling microscopy (STM) studies of SrTiO₃ surfaces frequently reveal the coexistence and even intermixing of multiple surface structures,[21,26,28] a feature common to the surfaces of other complex perovskites such as BaTiO₃ and (La, Sr)MnO₃ (ref. 29 and 30). The variety of coexisting surface reconstructions and their dependence on sample history underlines the necessity of changing the framing from identifying a single specific reconstruction to mapping out the range of possible reconstructions. This diversity can serve as an ideal

aInstitute of Materials Chemistry, TU Wien, 1060 Vienna, Austria. E-mail: georg. madsen@tuwien.ac.at

bInstitute of Applied Physics, TU Wien, 1040 Vienna, Austria

cInstituto de Nanociencia y Materiales de Aragón, CSIC-Universidad de Zaragoza, 50009 Zaragoza, Spain

† Electronic supplementary information (ESI) available. See DOI: https://doi.org/10.1039/d4dd00231h

showcase of the power of MLFF-supported stochastic searches, which, in turn, could be applied to the exploration of surface structures of further, technologically relevant materials.

In particular, the enhanced accuracy and reliability of MLFFs facilitate the application of stochastic algorithms for structural exploration of materials.[31–36] Stochastic approaches require a substantial volume of calculations and are impracticable with *ab initio* methods such as density functional theory (DFT) as the backend. This holds especially true for large and complex systems, such as surface reconstructions of multi-element compounds. Such systems usually feature a complex energy surface with too many degrees of freedom to explore exhaustively, as well as many local minima, where local searches for an optimal structure largely depend on the initial geometry of the search. Given that stochastic searches produce more diverse structures than, *e.g.*, molecular dynamics, a transferable, robust, and generalizable force field trained on a diverse dataset is key to their success.

An accurate MLFF is, however, just one part of the toolbox necessary to build a robust and efficient workflow for structure searches. The design of an MLFF can enhance or restrict its transferability and any MLFF has the potential to emit infinitely diverse mispredictions. This is especially significant for stochastic searches, which, by design, tend to move into regions that the model was not trained on. Extrapolation happens almost surely in high-dimensional models,[37] and therefore is not, by itself, an indicator of poor performance. Research has hence focused on estimates of uncertainty as proxies for the error incurred by using a given MLFF.[38–42] When that error is suspected to exceed tolerable margins, retraining with an expanded training set can help extend the applicability of the MLFF. In such scenarios, an efficient algorithm must aim toward issuing only as many *ab initio* calculations as required, while preventing waste of resources on redundant or irrelevant configurations.[43,44] Thus, the need for an uncertainty metric to evaluate the quality of the results dovetails with the usefulness of such metrics for identifying or even generating optimally informative new configurations.[33] This makes stochastic structure searches naturally part of the domain of application of active learning (AL) workflows, where an informed data selection is achieved through uncertainty estimation.[45]

This work focuses on the less well-understood Ti-rich surface reconstructions at the (110) orientation of $SrTiO_3$. Literature describes numerous composition-related $SrTiO_3$(110) surface reconstructions that can be broadly grouped into two families, characterized by Ti-poor ($n \times 1$) and Ti-rich ($2 \times m$) overlayers on an otherwise unchanged bulk.[26] Here, n and m denote the number of (1×1) bulk unit cells covered in the [001] and [$\bar{1}$10] directions, respectively. Fig. 1 shows examples of both variations. Ti-rich overlayers pose the particular challenge that, even when the Ti-to-Sr ratio is controlled,[26] numerous reconstructions lacking corresponding DFT models can coexist. Here, we again observe pronounced inhomogeneity in new STM measurements and are able to identify varied reconstructions (see Fig. S1 of the ESI†). We mitigate the lack of suitable atomistic models of such inhomogeneous surfaces by combining an evolutionary search algorithm with a transferable
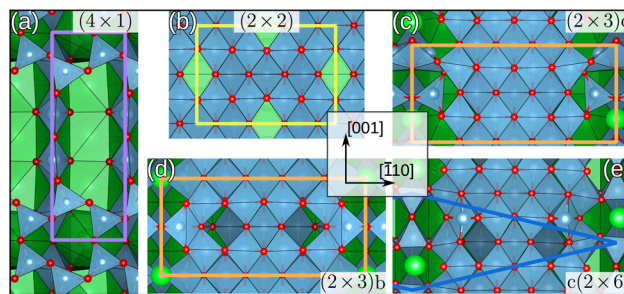

Fig. 1 Top view of $SrTiO_3$(110)-($n \times m$) surface reconstructions that are used for MLFF training: previously explored (4×1),[22] (2×2) and (2×3)b,[25] and newly identified (2×3)c and c(2×6) from this work. O atoms are shown in red, Sr in green and Ti in blue. The blue polyhedra depict TiO_x polyhedra in all panels. Colored lines indicate the borders of unit cells. For (e) only half the unit cell is outlined. See Fig. S10 of the ESI† for the full unit cell and additional side views. ($n \times 1$) overlayers exclusively contain tetrahedrally coordinated TiO_4 units, while $SrTiO_3$(110)-($2 \times m$) surface reconstructions are predominantly composed of octahedrally coordinated TiO_6 units and in experimental observations include at least one Sr atom per unit cell.[19,25]

MLFF to identify valid candidate structures. Transferability, in particular, is an important prerequisite for minimizing the computational cost associated with generating training data. We demonstrate that by utilizing small, well-known reconstructions and implementing a careful data selection routine built on structural and spatially-resolved local uncertainties, we can arrive at an MLFF capable of extrapolating to larger, more complex structures. In the following, we first discuss the active learning approach. We then proceed to show how structural models reproducing the STM images can be systematically obtained for all the coexisting surface structures.

2 Approach

2.1 Active learning workflow for the initial dataset

We set out to iteratively construct a versatile MLFF, which can subsequently be used in combination with an evolutionary search, specifically the covariance matrix adaptation evolution strategy (CMA-ES),[46] to explore the different motifs encountered on the inhomogeneous, Ti-rich $SrTiO_3$(110) surface. In this study, we refer to the suggested atomistic surface reconstruction models as "unit cells" and to the corresponding regions of the STM images as "motifs". Due to the complex nature of the various surface reconstructions, the MLFF must be capable of resolving a wide range of local environments. Additionally, given the nature of the evolutionary search, which produces more diverse and possibly unphysical structures than, *e.g.*, molecular dynamics, intermediate configurations are likely to exhibit unusual properties, such as unphysical bond lengths. For that reason, the MLFF needs to be robust and the underlying training data diverse, making the construction of a suitable training set containing *ab initio* energies and forces far from trivial. Furthermore, we aim to explore (2×5) surface reconstructions with up to 450 atoms per unit cell, where the computational costs associated with DFT evaluation is

prohibitive for constructing diverse training databases. The MLFF must therefore also be able to generalize from smaller training structures to the larger unit cells explored.

We started from the basis of our previous results on a Ti-poor reconstruction of SrTiO$_3$(110), namely the (4×1) (see Fig. 1(a)).[35] There, overlayers were symmetrically set up on opposite sides along the surface normal, attached to bulk-like layers in between. All structures in this study are constructed in the same manner as introduced in ref. 35. We constructed a database of 495 structures, re-evaluated using VASP[47] with the r^{2}SCAN functional.[48] This initial DFT database is illustrated in Fig. 2 together with a 2D projection of the spherical Bessel descriptors of the local environments of Sr atoms obtained using the uniform manifold approximation and projection (UMAP) method for dimension reduction.[49]

Using these data we trained a ten-member committee based on the descriptor-based NeuralIL architecture.[3,39] In committees the uncertainty is approximated by training a set of models that vary by initialization seed, hyperparameters, architecture, or training data, and monitoring their disagreement on a prediction to obtain the model variance. The majority of the computational cost incurred when training a descriptor-based model can be attributed to the calculation of the descriptors and the associated vector-Jacobian product operator. In this study, we vary the initial seed to enable NeuralIL's particularly efficient committee implementation that reuses descriptor encodings for all members, so that committees needed for uncertainty estimation can be trained with a negligible performance penalty.[39]

We then generated CMA-ES trajectories for the (2×2), (2×3)a and (2×3)b reconstructions, see Fig. 1. The CMA-ES[46] samples a population of λ individuals, $x_k^{(g)}$, $k = 1, \ldots, \lambda$, for every generation g from the multivariate normal distribution

$$x_k^{(g)} \sim \mathcal{N}\left(m^{(g-1)}, \left[\sigma^{(g-1)}\right]^2 C^{(g-1)}\right), \tag{1}$$

with distribution mean m, step size σ and covariance C. In the present case, the $x_k^{(g)}$ are the Cartesian coordinates of the atoms as explained in the Technical details section. The population size λ, initial mean $m^{(0)}$, and initial step size $\sigma^{(0)}$, are user-defined hyperparameters. We refer to the initial mean as the founder structure and data generation started from founders that are variations of published structures.[25] The mean is updated to move the average towards individuals/structures with low energy. Similarly, the covariance matrix is updated to let it reflect successful steps according to the CMA-ES algorithm.[46]

We then iteratively added structures from the MLFF backed CMA-ES trajectories using the committee uncertainty estimate aggregated structure-wise[39,42]

$$s = \frac{1}{3N} \sum_j^N \sum_{k \in \{x,y,z\}} s_j^{(k)}, \tag{2}$$

to identify structures that should be added to the training data. Here, N is the number of atoms and $s_j^{(k)}$ the committee standard deviation of the k-th component of the force on atom j. In total, 519 (2×2) and 775 (2×3) training structures were generated during this process. The dataset was then refined by
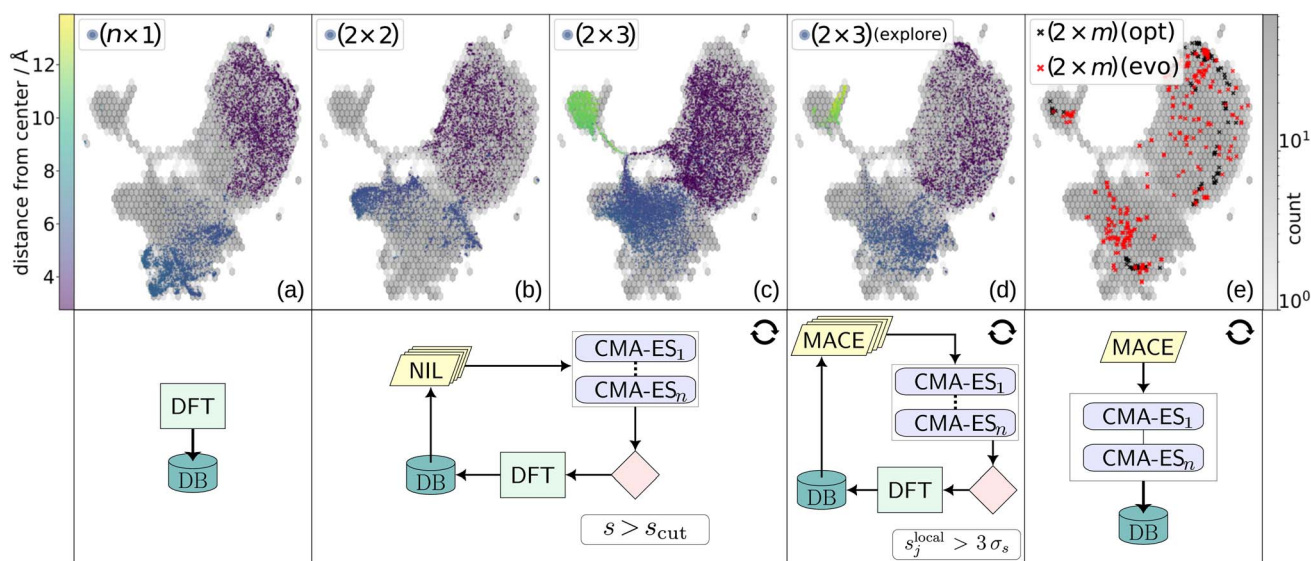
Fig. 2 The upper row shows a two-dimensional UMAP representation of the spherical Bessel descriptors of the local environments of Sr atoms. The background of all plots displays the training data as fitted utilizing a log−log hexbin approach, with darker shading indicating higher data density per bin (right color bar). In the foreground of the first four columns, colored dots depict the local environments of Sr atoms corresponding to the structures labeled $(n \times 1)$, (2×2), (2×3) and "(2×3) (explore)", respectively. Here, the colors indicate each atom's distance from the center of the surface slab (left color bar). Column (e) features the local Sr descriptors of geometry-optimized $(2 \times m)$ results (including c (2×6), (2×4)c, c(2×8), (2×4)d, (2×4)f, (2×5)b, and (2×5)c projected on the same 2D UMAP background as black crosses. The red crosses in the same subplot depict the same for randomly chosen individuals from earlier generations of (2×4)f evolution runs, including generations 10, 25, and 50. In the bottom row, the active learning workflow is schematically illustrated. The circle arrows represent iterations, the stacked parallelograms depict MLFF committees, and the CMA-ES blocks indicate parallel and/or sequential execution of multiple CMA-ES runs.

incorporating 141 (5×1) and a further 296 (4×1) structures from ref. 35 into the training data in order to reinforce performance on $SrTiO_3(110)$-($n \times 1$) structures.

The structures added to the training data are illustrated in Fig. 2 (top row) using the 2D UMAP projection. The background shows the distribution of all $SrTiO_3(110)$ data from the final dataset through hexagonal binning. The foreground of subplots (a)–(c) depicts the local Sr environments of the individual ($n \times 1$), (2×2) and (2×3) data sets, with the colors indicating the distance of each Sr atom from the center of the surface slab. It can be seen how the data from each reconstruction used for the initial dataset contributes their own distinct environments. Not surprisingly, the Sr in the (2×3) overlayer reconstructions, depicted as green points in the top left of Fig. 2(c), is particularly distinctive, since it did not occur in the ($n \times 1$) or (2×2) training data. In the bottom row of Fig. 2, the workflows used for generating the data depicted directly above them are illustrated, with (b) and (c) generated by different iterations of the same workflow.

2.2 Exploration-based active learning

Using the 2226 structures in the training database, we trained a five-member committee using the equivariant message-passing neural network framework MACE.[4] MACE provides significantly improved accuracy and transferability and is more data-efficient than NeuralIL due to its equivariant architecture and custom-learned atomic representations. The mean absolute error in the force components, f_{MAE}, for the (4×1), (2×2) and (2×3) sets decreased by a factor of 2.5 when moving from NeuralIL to MACE with the same training data (see Table S2†). Notably, the aggregated force uncertainty estimates obtained from MACE and NeuralIL committees, eqn (2), exhibit strong correlations for all highly uncertain configurations (see Fig. S3 of the ESI†). This underscores the efficiency of constructing the database by using uncertainties derived from the NeuralIL committee. The initial MACE model demonstrated strong performance on test data and could reliably be applied to investigate ($n \times 1$) and (2×2) structures using the CMA-ES with relatively large population sizes, $\lambda = 100$, and initial step sizes, $\sigma^{(0)} \in [0.1, 0.35]$.

We then performed exploratory CMA-ES searches on the (2×3) surface with population size $\lambda = 100$, and varying the initial step size in the range of $\sigma^{(0)} \in [0.1, 0.5]$. To expand the search space, we developed a more generic founder structure (pictured in Fig. S7 of the ESI†), rather than relying on published findings. From these searches, we identified the new $SrTiO_3(110)$-(2×3)c reconstruction, Fig. 1(c). A key feature of this structure is the alignment of the overlayer Sr atoms relative to the topmost TiO_x rows. Geometry optimizing the (2×3)b and (2×3)c structures using VASP, reveals that the new (2×3)c has a lower energy of $\Delta E = 160$ meV per (1×1) bulk unit cell.

While two of the 35 initial CMA-ES searches produced the new stable configuration, a majority of these evolutionary searches were found to be prone to instability, specifically to the expulsion of an Sr atom (see inset structure in Fig. 3). Although problematic structures could be identified manually, an active
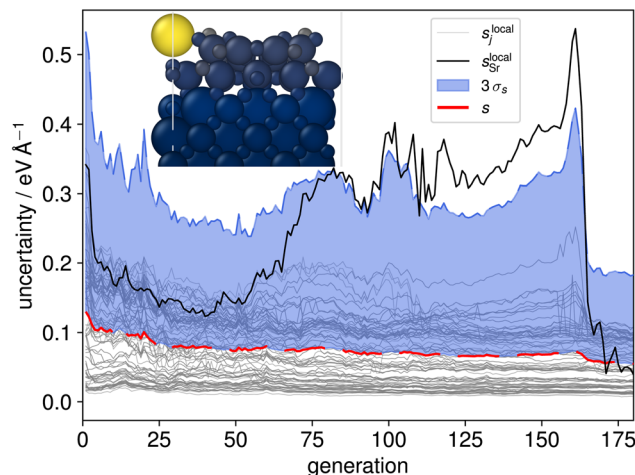

Fig. 3 Spatially resolved uncertainty of a CMA-ES trajectory of a mirror-symmetric $SrTiO_3(110)$-(2×3) structure. The gray lines show the locally aggregated uncertainty s_j^{local} of each atom. The mean of all local uncertainties, *i.e.*, the global structure uncertainty s, is depicted as a dashed red line, with three times their standard deviation σ_s shaded in blue. The local uncertainty associated with the overlayer Sr atom is highlighted in black. The inset shows a structure at generation 160 with the atoms colored according to the local uncertainty estimate. The color scale ranges from dark-blue (lowest) to yellow (highest local uncertainty).

learning procedure needs to be able to identify and incorporate such structures into the training data based on computed quantities such as model uncertainties. Interestingly this behaviour was not reflected in the aggregated structure uncertainty s, eqn (2), as shown with the red line in Fig. 3. We therefore calculated spatially resolved atomic uncertainties by aggregating over neighboring atoms within a cutoff radius[42] (in the following set to 4 Å) instead of over the entire structure

$$s_j^{local} = \frac{1}{3N_j} \sum_{n}^{N_j} \sum_{k \in x,y,z} s_n^{(k)}, \tag{3}$$

for atom j with N_j neighbors. Recently, such locally aggregated uncertainties and errors were shown to have a direct monotonic correlation.[42] This is in contrast to the uncertainties and errors of the individual force components, $s_j^{(k)}$, which only feature an asymmetric relationship, where large errors occur predominantly for large uncertainties. However, large uncertainties do not necessitate large errors, causing many false positives when trying to select high-energy structures.[42] Locally aggregated uncertainties thus allow us to reliably identify high-error sub-regions without false positives.

The local uncertainties, eqn (3), clearly identified the misinterpretation of unphysical local structures which led to escalating errors during the evolution and are thus a reliable indicator for atoms being expelled from the surface. A visual representation of the evolution of local and global uncertainty within a CMA-ES run is given in Fig. 3. The solid black line tracks the local uncertainty associated with the single Sr atom in the overlayer, s_{Sr}^{local}. Notably, s_{Sr}^{local} begins to increase after generation 70 and exceeds three times the standard deviation of

the local uncertainties ($3\sigma_s$) after generation 90. This behavior is further illustrated by the atomic structure shown in the inset: at generation 160 even the configuration with the lowest energy features a bright yellow sphere, indicating the high local uncertainty in the force estimate for the Sr atom. The rise in local uncertainty corresponds to environments that the model is increasingly uncertain about but which are localized enough for them to have a low weight in the globally aggregated uncertainty, eqn (2). Once an atom is separated from the rest by more than the cutoff radius of the underlying MLFF model, its contribution to the energy and forces, and thus also to the local and global uncertainty, becomes zero.

In an additional AL step we then identified trajectories where the local uncertainty associated with at least one atom j exceeded three standard deviations of all local uncertainties ($s_j^{local} > 3\sigma_s$). We randomly sampled structures from 40 CMA-ES evolution trajectories, with 32 of these structures exhibiting similar local uncertainty behaviour as shown in Fig. 3. The sampling was performed uniformly but was restricted to intact surface slabs, meaning that generations following Sr expulsion were excluded. The 2D UMAP of the local Sr descriptors within these structures is shown in Fig. 2(d), highlighting the added diversity that was achieved. With these additional 392 structures, the full training set consisted of 2618 configurations. This was used to train the final MACE model which was utilized for all further CMA-ES runs. The complete database and trained model are made available on Zenodo.[50]

3 Technical details

3.1 CMA-ES

We applied CLINAMEN2,[51] a functional-style Python framework that interfaces to different codes for loss evaluation in a straightforward manner, to perform the covariance matrix adaptation evolution strategy (CMA-ES)[46] for all structure searches in this study.

Surface slabs were set up as illustrated in Fig. S8.† For all system sizes an anchor region of fixed atom positions was defined at the center of each slab, consisting of bulk-like layers that remained unchanged. The Cartesian coordinates of the atoms in the over- and attachment-layers are the degrees of freedom, *i.e.*, the individuals $x_k^{(g)}$, sampled from the multivariate normal distribution. Opposite sides of all slabs were made symmetric. Further symmetry elements (*e.g.*, mirror planes) were leveraged where feasible to drastically reduce the number of degrees of freedom in larger unit cells.

3.2 DFT

We used VASP[47] version 6.2.0 with the r^{2}SCAN functional[48] for all *ab initio* calculations in this study, including single-shot structure evaluations for training and test data and geometry optimization of low-energy candidate structures. The energy cutoff was set to 440 eV, the width of Gaussian smearing to 0.02 eV, and EDIFF to 10^{-5}. To ensure compatibility of DFT energies and forces calculated for different system sizes, we utilized the optimized *k*-point grid generator by Wang *et al.*, with the minimum distance set to 25 in the PRECALC input.[52]

3.3 Machine-learned force fields

All NeuralIL models in this work used $r_{cut} = 4.0$ Å and $n_{max} = 5$, with ResNet core widths set to [256, 128, 64, 32, 32, 32, 16]. Training on forces was run for only 101 epochs due to the replacement of the standard Adam optimizer with the versatile learned optimizer VeLO,[53] which drastically reduced the number of epochs needed for convergence by orders of magnitude and eliminated the need to set up a learning rate schedule.[39]

With the majority of hyperparameters set to default values, MACE trainings were performed with a cutoff radius $r_{max} = 4.0$ Å and two hidden layers set to 128 channels for scalar and vector properties each. The maximum number of epochs was set to 1200 with an early stopping patience of 50, and energy and force weights of 1 and 100, respectively. Afterwards, training was run for an additional 300 epochs with an increased energy weight of 1000 and an unchanged force weight.

Notably, as mentioned in the workflow description, training and test data are generated in a manner that allows the resulting MLFF to serve as the backend for CMA-ES runs, instead of DFT codes. This means that the MLFF needs to be able to reliably evaluate configurations far from the equilibrium, which is achieved by utilizing the CMA-ES itself for data generation. At each step, training and test data belong to the same distribution.

To eliminate unusable or obstructive data, all structures with force components larger than 500 eV Å^{-1}, as well as those that do not converge with the chosen parameters, are excluded.

3.4 Evolution details

In total, approximately 3000 exploratory CMA-ES runs were performed on various system sizes and founder structures. For each founder, runs were started for different random seeds to leverage the stochasticity of the method. Population sizes were varied between $\lambda \in \{25, 35, 50, 100\}$, with the choice of step size σ depending on symmetry. Whenever mirror symmetry was enforced, $\sigma^{(0)}$ was capped at 0.35 Å, while evolutionary searches without symmetry were performed for step sizes up to 0.5 Å. The lower limit for $\sigma^{(0)}$ was 0.1 Å for all cases.

The computation of one CMA-ES trajectory starting from a founder containing 450 atoms and running for up to 1000 generations, with population size $\lambda = 100$, required only between one and three hours on one NVIDIA A40 GPU with 46 GiB memory when utilizing MACE$_{full}$, depending on early stopping. For that reason, it was possible to freely explore various founders to then select highlights for further investigation using DFT.

3.5 Experimental methods

SrTiO$_3$(110) single crystals (CrysTec GmbH, 0.5 wt% Nb$_2$O$_5$, 5 × 5 × 0.5 mm^3, one-side polished, miscut less than 0.3°) were prepared *ex situ* by sonication in heated neutral detergent (3% Extran® MA02, diluted in ultrapure water, 2 × 20 min) and

ultrapure water (milli-Q™, 10 min). Subsequent boiling for 10 min in milli-Q™ water removed commonly observed CaO contamination. The samples were then mounted on flag-style, HNO$_3$-cleaned Nicrofer® sample plates with Nicrofer® clips, and inserted in a UHV setup comprising three interconnected chambers: (i) a preparation chamber (base pressure below 10^{-10} mbar) equipped with sputtering–annealing facilities and an evaporator for Sr deposition; (ii) an analysis chamber (base pressure below 5×10^{-11} mbar) equipped for STM (SPECS Aarhus STM 150), low-energy electron diffraction (LEED) (Omicron SpectaLEED), and X-ray photoelectron spectroscopy (XPS) (nonmonochromatic dual-anode Mg/Al Kα source, SPECS Phoibos 100 analyzer, normal emission); (iii) a pulsed-laser deposition (PLD) chamber (base pressure below 2×10^{-9} mbar).

After a few cycles of sputtering–annealing (6×10^{-6} mbar Ar, 1 keV, 5–10 μA, 10 min, followed by 1 h at 1000 °C, 6×10^{-6} mbar O$_2$), the surface cleanliness was verified through XPS and STM. The surface stoichiometry was then adjusted *via* submonolayer deposition of Sr (*via* molecular-beam epitaxy)[54] or TiO$_2$ (*via* PLD).[55] The resulting surface periodicity was verified by LEED and STM. The surface presented in this work was obtained starting from a mixed $(4 \times 1)/(5 \times 1)$ reconstruction. 1.4 ML Ti was deposited in PLD by keeping the sample at 580 °C in a background oxygen pressure of 6×10^{-6} mbar O$_2$, followed by ramp down at 80 °C min^{-1}.

STM images were acquired in constant-current mode with homemade, electrochemically etched W tips. The tips were prepared *in situ* by Ar sputtering (1 μA, 30 min). Voltage (up to 10 V) or current pulses (up to 10 nA) were applied while in tunneling contact to reshape the tip and improve resolution. Positive bias voltages correspond to tunneling into the empty states of the sample.

4 Results and discussion

4.1 Structure search

The final MACE model described above enabled us to perform a large number of structure searches for SrTiO$_3$(110)-$(2 \times m)$, $m \in \{3, 4, 5\}$, with initial step sizes in the range of $\sigma^{(0)} \in [0.1, 0.5]$. The choice of population size λ strongly influences the stability of the evolution trajectories, especially for such rough loss landscapes. Moreover, a larger population size increases the likelihood of identifying the most stable structure, rather than other stable structures nearby on the loss surface. Because of this trade-off, we performed the same searches with population sizes between 25 and 100, as outlined in the Technical details.

The UMAP in the upper row of Fig. 2(e) illustrates the variety of local Sr environments encountered in randomly selected structures chosen from early generations (10, 25, and 50) along these trajectories (red crosses). In comparison, the local environments in geometry-optimized structures of different unit cell sizes are clearly more uniform (black crosses). Throughout all CMA-ES searches, spatially resolved local uncertainty, along with the loss, served as an indicator for structure stability. Importantly, after adding the structures from the exploration-based active learning in the (2×3) cell, Fig. 2(d), to the

training data, the MLFF learned to avoid regions leading to the previously observed overlayer instability. Fig. S4 of the ESI† shows uncertainty trajectories for (2×3), (2×4), and (2×5) runs, where no local uncertainty exceeds $3\sigma_s$. Of particular interest is the (2×3) trajectory, which still illustrates that the Sr environment is the most uncertain, but does not escalate anymore (compare to Fig. 3). Importantly, this demonstrates the transferability of the model when extrapolating to reconstruction with larger unit cells.

With this approach, we were able to discover the new candidate structures shown in Fig. 4, namely (2×4)d (yellow), (2×4)e (orange), c(2×8) (blue), and (2×5)c (white). All of these structures were then relaxed with VASP and will be compared to experiment in the following. They are available on Zenodo.[50]

4.2 Comparison to experiment

Fig. 5 depicts a high-resolution STM image that illustrates how the preparation of Ti-rich surfaces results in a mixture of various surface structures. In this image, local symmetries with (2×4) (orange and yellow), (2×5) (white), and c(2×8) (blue) motifs can be observed. In the STM images of the Ti-rich SrTiO$_3$(110) surfaces, the Sr atoms and TiO$_x$ rows (henceforth referred to as "TiO ridges") are visible as bright spots. In the cells shown in Fig. 5, the topmost TiO ridges tend to be aligned with the Sr atoms. This feature is incompatible with the STM image obtained from the published (2×5)b structure,[25] where the TiO ridges are offset with respect to the Sr atoms.

In experiments, it is rare to encounter a surface composition that precisely matches the stoichiometry of a given thermodynamically stable surface reconstruction. Rather, the average composition observed in experiments often lies between two specific reconstructions, implying the coexistence of these reconstructions on the surface. Additionally, kinetic limitations may lead to the coexistence of even more surface reconstructions, with local variations in composition, while preserving the overall average surface composition.

To facilitate the investigation of these coexisting reconstructions, founder structures for SrTiO$_3$(110)-(2×4) and $-(2 \times 5)$ cells were created as described in Section S3 of the ESI.† In short, they were generated by varying initial atom positions and adjusting the stoichiometry (addition or removal of TiO$_2$ units). Details regarding the evolutionary searches, including the number of runs, population sizes λ, and step sizes σ are summarized in the Technical details. All newly proposed structures were tested by comparing the corresponding simulated STM images with the experimental data in Fig. 5. Important criteria for matching were the position of the Sr adatoms and their relative alignment to the TiO ridges.

For (2×5) systems, the initial placement of "TiO$_2$ vacancies" resulted in distinct founders, with the vacancies positioned either in-line or out-of-line relative to the overlayer Sr in the $[\bar{1}10]$ direction. All sensible configurations resulting from these founders yielded significantly higher energies than the (2×5)b from ref. 25. However, an alternative candidate structure, (2×5)c (see Fig. 4), could be identified due to its distinct features. There, in contrast to
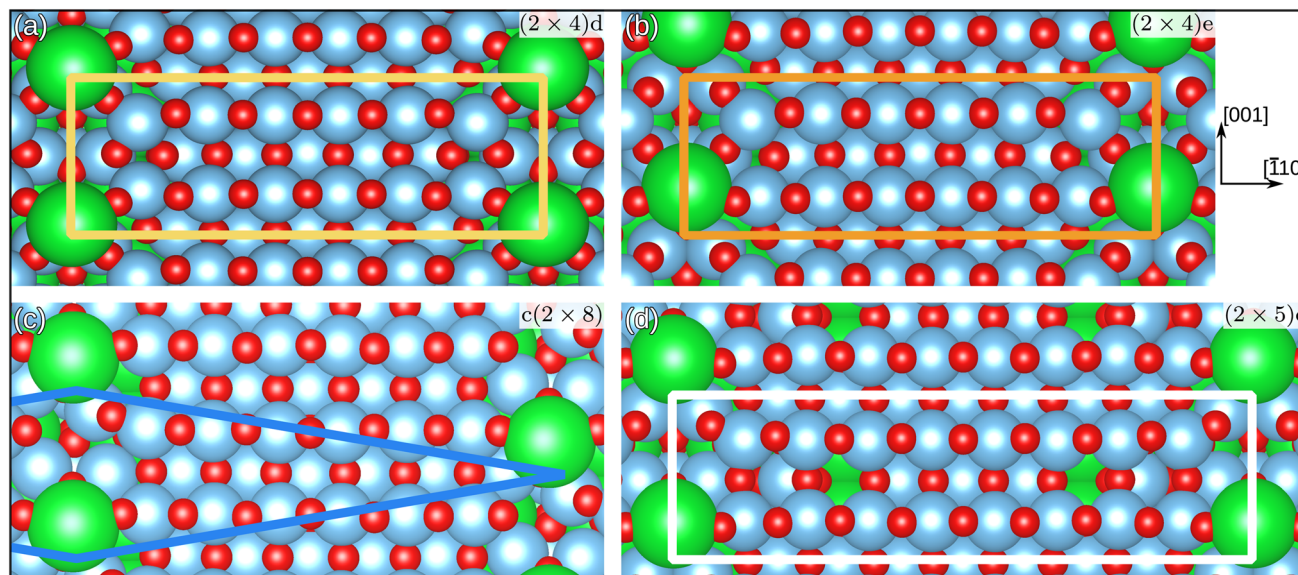
Fig. 4 Top view of SrTiO₃(110) overlayer candidates. (a) and (b) with (2 × 4) bulk periodicity. (c) With centered (2 × 8) periodicity and (d) with (2 × 5) periodicity. The orange and white rectangles and the blue rhombus indicate (parts of) the respective unit cells. Only half of the c(2 × 8) unit cell is outlined. See Fig. S11 of the ESI† for the full picture.
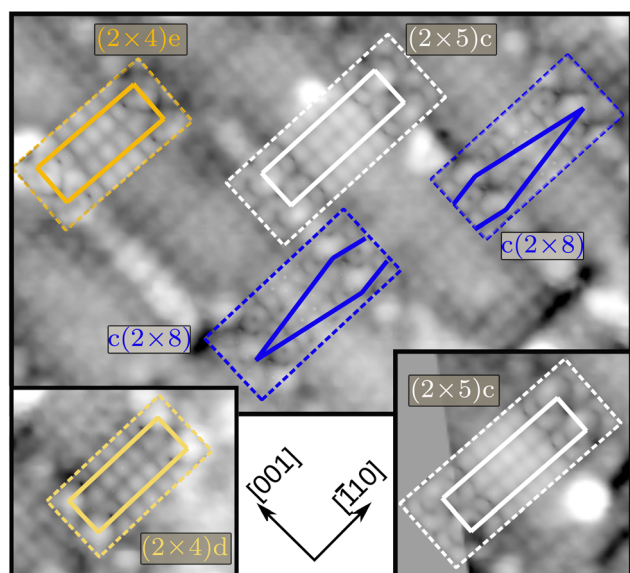
Fig. 5 Different regions of the same STM measurement of Ti-rich SrTiO₃(110), showing (2 × 4) (orange and yellow), c(2 × 8) (blue), and (2 × 5) (white) motifs. The colored dashed lines mark the border of simulated images, overlaid with 50% transparency, colored solid lines represent the model unit cells or parts thereof in the case of c(2 × 8). Imaging parameters: $V_{sample\ bias}$ = +1.8 V, I = 0.04 nA. Simulated images were created using the Tersoff–Hamann approximation.[56] Alternative versions of this figure, including without overlays, with fully opaque simulation images or the topmost atoms overlaid, as well as a larger cutout from the experimental image, are shown in the ESI as Fig. S1.†

(2 × 5)b, the overlayer Sr atom is aligned with a TiO ridge rather than being offset. While the energy difference between the two is ΔE = 208 meV per (1 × 1) bulk unit cell in favor of (2 × 5)b, the

new (2 × 5)c clearly fits regions of the inhomogeneous surface, as shown in Fig. 5.

The investigation of SrTiO₃(110)-(2 × 4) identified two stable surface structures, labeled (2 × 4)d and (2 × 4)e, both shown in Fig. 4. Although the difference in DFT energies between them is vanishingly small – only 2 meV per (1 × 1) bulk unit cell in favor of (2 × 4)d – the arrangement of the overlayer atoms is distinct. The most noticeable differences include the relative position of the overlayer Sr atom with respect to the TiO ridges and the resulting positional changes. Additionally, the centered unit cell c(2 × 8) was found as a candidate structure for explaining regions on the STM measurement showing a shift between Sr positions.

The newly proposed structures (2 × 4)d, (2 × 4)e, c(2 × 8), and (2 × 5)c thus provide previously missing atomistic models, which we successfully matched to the various experimentally observed motifs.

5 Summary and conclusions

We successfully integrated neural-network force fields with the covariance matrix adaptation evolution strategy to develop an accurate and transferable machine-learned force field suitable for the exploration of Ti-rich SrTiO₃(110) surface reconstructions. The required training data were generated through an active learning workflow, which involved repeatedly performing CMA-ES runs on SrTiO₃(110)-(2 × 2) and -(2 × 3) founder structures to gather uncertain and diverse data. During this process, invariant, descriptor-based NeuralIL committees were utilized for energy evaluation and uncertainty estimation. The collected data was then used to train an equivariant MACE model with learned representation suitable for production runs.

To fine-tune the training data in a further AL step, and thereby enhance model performance, we employed spatially resolved uncertainty estimation to identify underrepresented local environments which global uncertainty measures had failed to resolve. The resulting MLFF, MACE$_{full}$, was trained on 2618 structures spanning SrTiO$_3$(110)-($n \times m$), $n \in \{4, 5\}$, $m \in \{2, 3\}$.

We successfully identified two not previously reported candidates for stable (2×3) reconstructions. These structures were then used to extrapolate to (2×4) and (2×5) founder structures for evolutionary exploration. With this approach, we found new stable candidate structures for SrTiO$_3$(110)-(2×4) and -(2×5), explaining different experimentally observed regions of the inhomogeneous Ti-rich surface. This method could be extended to other multi-element oxides featuring complex, composition-related, and possibly coexisting surface reconstructions characterized by large motifs.

Data availability

The trained models, training and test data, and POSCAR files containing founders and results are available on Zenodo. This dataset also includes an example evolution script to be used in combination with CLINAMEN2 and MACE.[50]

Conflicts of interest

There are no conflicts to declare.

Acknowledgements

Notes and references

1 J. Behler and M. Parrinello, *Phys. Rev. Lett.*, 2007, **98**, 146401.

2 K. T. Schütt, H. E. Sauceda, P.-J. Kindermans, A. Tkatchenko and K.-R. Müller, *J. Chem. Phys*, 2018, **148**, 241722.

3 H. Montes-Campos, J. Carrete, S. Bichelmaier, L. M. Varela and G. K. H. Madsen, *J. Chem. Inf. Model.*, 2022, **62**, 88–101.

4 I. Batatia, D. P. Kovacs, G. Simm, C. Ortner and G. Csányi, *Adv. Neural. Inf. Process. Syst.*, 2022, **35**, 11423–11436.

5 S. Batzner, A. Musaelian, L. Sun, M. Geiger, J. P. Mailoa, M. Kornbluth, N. Molinari, T. E. Smidt and B. Kozinsky, *Nat. Commun.*, 2022, **13**, 2453.

6 Y.-L. Liao, B. Wood, A. Das and T. Smidt, EquiformerV2: Improved Equivariant Transformer for Scaling to Higher-Degree Representations, *arXiv*, preprint, arXiv:2306.12059, 2024, DOI: **10.48550/arXiv:2312.04174**.

7 I. Batatia, P. Benner, Y. Chiang, A. M. Elena, D. P. Kovács, J. Riebesell, X. R. Advincula, M. Asta, M. Avaylon, W. J. Baldwin, F. Berger, N. Bernstein, A. Bhowmik, S. M. Blau, V. Cărare, J. P. Darby, S. De, F. Della Pia, V. L. Deringer, R. Elijošius, Z. El-Machachi, F. Falcioni, E. Fako, A. C. Ferrari, A. Genreith-Schriever, J. George, R. E. A. Goodall, C. P. Grey, P. Grigorev, S. Han, W. Handley, H. H. Heenen, K. Hermansson, C. Holm, J. Jaafar, S. Hofmann, K. S. Jakob, H. Jung, V. Kapil, A. D. Kaplan, N. Karimitari, J. R. Kermode, N. Kroupa, J. Kullgren, M. C. Kuner, D. Kuryla, G. Liepuoniute, J. T. Margraf, I.-B. Magdău, A. Michaelides, J. H. Moore, A. A. Naik, S. P. Niblett, S. W. Norwood, N. O'Neill, C. Ortner, K. A. Persson, K. Reuter, A. S. Rosen, L. L. Schaaf, C. Schran, B. X. Shi, E. Sivonxay, T. K. Stenczel, V. Svahn, C. Sutton, T. D. Swinburne, J. Tilly, C. van der Oord, E. Varga-Umbrich, T. Vegge, M. Vondrák, Y. Wang, W. C. Witt, F. Zills and G. Csányi, A Foundation Model for Atomistic Materials Chemistry, *arXiv*, preprint, arXiv:2401.00096, 2024, DOI: **10.48550/arXiv.2401.00096**.

8 T. Maxson and T. Szilvási, *J. Phys. Chem. Lett.*, 2024, **15**, 3740–3747.

9 S. Zhang, M. Z. Makoś, R. B. Jadrich, E. Kraka, K. Barros, B. T. Nebgen, S. Tretiak, O. Isayev, N. Lubbers, R. A. Messerly and J. S. Smith, *Nat. Chem.*, 2024, **16**, 727–734.

10 A. S. Bhalla, R. Guo and R. Roy, *Mat. Res. Innovat.*, 2000, **4**, 3–26.

11 M. A. Peña and J. L. G. Fierro, *Chem. Rev.*, 2001, **101**, 1981–2018.

12 P. Zubko, S. Gariglio, M. Gabay, P. Ghosez and J.-M. Triscone, *Annu. Rev. Condens. Matter Phys.*, 2011, **2**, 141–165.

13 D. P. Kumah, J. H. Ngai and L. Kornblum, *Adv. Funct. Mater.*, 2020, **30**, 1901597.

14 A. Spinelli, M. A. Torija, C. Liu, C. Jan and C. Leighton, *Phys. Rev. B*, 2010, **81**, 155110.

15 K. S. Takahashi, M. Gabay, D. Jaccard, K. Shibuya, T. Ohnishi, M. Lippmaa and J. M. Triscone, *Nature*, 2006, **441**, 195.

16 Z. Wang, Z. Zhong, X. Hao, S. Gerhold, B. Stöger, M. Schmid, J. Sánchez-Barriga, A. Varykhalov, C. Franchini, K. Held and U. Diebold, *Proc. Natl. Acad. Sci. U. S. A.*, 2014, **111**, 3933.

17 L. Dudy, M. Sing, P. Scheiderer, J. D. Denlinger, P. Schütz, J. Gabel, M. Buchwald, C. Schlueter, T.-L. Lee and R. Claessen, *Adv. Mater.*, 2016, **28**, 7443.

18 K. Shibuya, T. Ohnishi, T. Uozumi, T. Sato, M. Lippmaa, M. Kawasaki, K. Nakajima, T. Chikyow and H. Koinuma, *Appl. Phys. Lett.*, 2006, **88**, 212116.

19 T. K. Andersen, D. D. Fong and L. D. Marks, *Surf. Sci. Rep.*, 2018, **73**, 213–232.

20 H. Tanaka, T. Matsumoto, T. K. T. Kawai and S. K. S. Kawai, *Jpn. J. Appl. Phys.*, 1993, **32**, 1405.

21 T. Kubo and H. Nozoye, *Surf. Sci.*, 2003, **542**, 177–191.

22 J. A. Enterkin, A. K. Subramanian, B. C. Russell, M. R. Castell, K. R. Poeppelmeier and L. D. Marks, *Nat. Mater.*, 2010, **9**, 245–248.

23 D. M. Kienzle, A. E. Becerra-Toledo and L. D. Marks, *Phys. Rev. Lett.*, 2011, **106**, 176102.

24 S. Gerhold, Z. Wang, M. Schmid and U. Diebold, *Surf. Sci.*, 2014, **621**, L1.

25 Z. Wang, A. Loon, A. Subramanian, S. Gerhold, E. McDermott, J. A. Enterkin, M. Hieckel, B. C. Russell, R. J. Green, A. Moewes, J. Guo, P. Blaha, M. R. Castell, U. Diebold and L. D. Marks, *Nano Lett.*, 2016, **16**, 2407–2412.

26 M. Riva, G. Franceschi, Q. Lu, M. Schmid, B. Yildiz and U. Diebold, *Phys. Rev. Mater.*, 2019, **3**, 043802.

27 A. Pancotti, J. Silva, A. de Siervo, R. Landers and P. Nascente, *Surf. Sci.*, 2022, **715**, 121937.

28 L. D. Marks, A. N. Chiaramonti, S. U. Rahman and M. R. Castell, *Phys. Rev. Lett.*, 2015, **114**, 226101.

29 E. H. Morales, J. M. P. Martirez, W. A. Saidi, A. M. Rappe and D. A. Bonnell, *ACS Nano*, 2014, **8**, 4465–4473.

30 G. Franceschi, M. Schmid, U. Diebold and M. Riva, *J. Mater. Chem. A*, 2020, **8**, 22947–22961.

31 M. Andersen, *Nat. Comput. Sci.*, 2023, **3**, 1009–1010.

32 M. K. Bisbo and B. Hammer, *Phys. Rev. B*, 2022, **105**, 245404.

33 X. Du, J. K. Damewood, J. R. Lunger, R. Millan, B. Yildiz, L. Li and R. Gómez-Bombarelli, *Nat. Comput. Sci.*, 2023, **3**, 1034–1044.

34 J. Timmermann, F. Kraushofer, N. Resch, P. Li, Y. Wang, Z. Mao, M. Riva, Y. Lee, C. Staacke, M. Schmid, C. Scheurer, G. S. Parkinson, U. Diebold and K. Reuter, *Phys. Rev. Lett.*, 2020, **125**, 206101.

35 R. Wanzenböck, M. Arrigoni, S. Bichelmaier, F. Buchner, J. Carrete and G. K. H. Madsen, *Digit. Discov.*, 2022, **1**, 703–710.

36 F. Brix, M.-P. Verner Christiansen and B. Hammer, *J. Chem. Phys.*, 2024, **160**, 174107.

37 R. Balestriero, J. Pesenti and Y. LeCun, Learning in High Dimension Always Amounts to Extrapolation, *arXiv*, preprint, arXiv:2110.09485, 2021, DOI: **10.48550/arXiv.2110.09485.**

38 L. Kahle and F. Zipoli, *Phys. Rev. E*, 2022, **105**, 015311.

39 J. Carrete, H. Montes-Campos, R. Wanzenböck, E. Heid and G. K. H. Madsen, *J. Chem. Phys.*, 2023, **158**, 204801.

40 P. B. Jørgensen, J. Busk, O. Winther and M. N. Schmidt, Coherent Energy and Force Uncertainty in Deep Learning Force Fields, *arXiv*, preprint, arXiv:2312.04174, 2023, DOI: **10.48550/arXiv:2312.04174.**

41 M. Kellner and M. Ceriotti, *Mach. learn.: sci. technol.*, 2024, **5**, 035006.

42 E. Heid, J. Schörghuber, R. Wanzenböck and G. K. H. Madsen, *J. Chem. Inf. Model.*, 2024, **64**, 6377–6387.

43 Y. Lee, J. Timmermann, C. Panosetti, C. Scheurer and K. Reuter, *J. Phys. Chem. C*, 2023, **127**, 17599–17608.

44 A. M. Slavensky, M.-P. V. Christiansen and B. Hammer, *J. Chem. Phys.*, 2023, **159**, 024123.

45 J. S. Smith, B. Nebgen, N. Lubbers, O. Isayev and A. E. Roitberg, *J. Chem. Phys.*, 2018, **148**, 241733.

46 N. Hansen, The CMA Evolution Strategy: A Tutorial, *arXiv*, preprint, arXiv:1604.00772, 2016, DOI: **10.48550/arXiv.1604.00772.**

47 G. Kresse and J. Furthmüller, *Phys. Rev. B*, 1996, **54**, 11169–11186.

48 J. W. Furness, A. D. Kaplan, J. Ning, J. P. Perdew and J. Sun, *J. Phys. Chem. Lett.*, 2020, **11**, 8208–8215.

49 L. McInnes, J. Healy and J. Melville, UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction, *arXiv*, preprint, arXiv.1802.03426, 2020, DOI: **10.48550/arXiv.1802.03426.**

50 R. Wanzenböck, E. Heid, M. Riva, G. Franceschi, A. M. Imre, J. Carrete, U. Diebold and G. K. H. Madsen, *Data for Exploring Inhomogeneous Surfaces: Ti-rich SrTiO3(110) Reconstructions via Active Learning*, 2024, DOI: **10.5281/zenodo.12528625.**

51 R. Wanzenböck, F. Buchner, P. Kovács, G. K. H. Madsen and J. Carrete, *Comput. Phys. Commun.*, 2024, **297**, 109065.

52 Y. Wang, P. Wisesa, A. Balasubramanian, S. Dwaraknath and T. Mueller, *Comput. Mater. Sci.*, 2021, **187**, 110100.

53 L. Metz, J. Harrison, C. D. Freeman, A. Merchant, L. Beyer, J. Bradbury, N. Agrawal, B. Poole, I. Mordatch, A. Roberts and J. Sohl-Dickstein, VeLO: Training Versatile Learned Optimizers by Scaling Up, *arXiv*, preprint, arXiv:2211.09760, 2022, DOI: **10.48550/arXiv.2211.09760.**

54 Z. Wang, F. Yang, Z. Zhang, Y. Tang, J. Feng, K. Wu, Q. Guo and J. Guo, *Phys. Rev. B*, 2011, **83**, 155453.

55 M. Riva, M. Kubicek, X. Hao, G. Franceschi, S. Gerhold, M. Schmid, H. Hutter, J. Fleig, C. Franchini, B. Yildiz and U. Diebold, *Nat. Commun.*, 2018, **9**, 3710.

56 J. Tersoff and D. R. Hamann, *Phys. Rev. Lett.*, 1983, **50**, 1998–2001.