

Multifidelity Machine Learning Interatomic Potentials for Charged Point Defects

Xinwei Wang^{1,2,*}, Irea Mosquera-Lois¹, and Aron Walsh¹

¹Thomas Young Centre and Department of Materials, Imperial College London,
London SW7 2AZ, United Kingdom

²Imperial-X, Imperial College London, London W12 0BZ, United Kingdom



(Received 2 March 2026; revised 8 September 2026; accepted 10 September 2026; published 8 October 2026)

Machine learning interatomic potentials (MLIPs) can now reproduce the energy, forces, and stresses of bulk materials with high accuracy compared to first-principles calculations. The description of imperfections, where coordination environments and electron counts deviate from those found in pristine reference structures, remains a challenge. We find that the current generation of foundation MLIPs does not describe the defect physics of the photovoltaic absorber Sb_2Se_3 . We introduce global defect charge embeddings that distinguish the bonding characteristics of different charge states. We further employ a multifidelity approach that combines low-cost (semilocal exchange-correlation functional) reference data with high-quality (nonlocal hybrid functional) energies and forces that describe well the subtleties of the defect energy landscape. The resulting defect-capable force fields can find stable structural configurations and predict defect thermodynamics in quantitative agreement with direct quantum mechanical calculations, at a fraction of the computational cost.

DOI: [10.1103/tzwd-cl3y](https://doi.org/10.1103/tzwd-cl3y)

I. INTRODUCTION

Defects play a crucial role in determining the properties of functional materials, governing key processes in semiconductors, photovoltaics, solid-state electrolytes, and catalytic systems [1]. In energy materials, defects directly impact device performance. Charged defects control carrier concentrations, mediate ion migration, and can introduce deep states that act as recombination centers [2,3]. Accurate predictions of these properties rely fundamentally on identifying the correct ground-state atomic structure. Failure to locate the equilibrium configuration impacts subsequent calculations of formation energies, migration barriers, carrier capture rates, and degradation reactions [4–8]. Ground and metastable structures can be identified through structure-searching methods that explore the energy landscape via symmetry-breaking distortions [9]. However, comprehensive structure searching remains a major bottleneck [10], especially for systems with complex defect landscapes where multiple metastable configurations compete within a narrow energy window [11]. Accurately describing these landscapes requires a high level of electronic-structure theory, such as hybrid

functionals, but their computational cost renders exhaustive structure searching prohibitively expensive. To reduce the cost of the underlying calculations, current approaches rely on low-precision approximations (e.g., coarse sampling of reciprocal space), which do not ensure identification of the true global minimum.

Machine learning interatomic potentials (MLIPs) offer a promising route to address this bottleneck. Recent foundation models [12–15] trained on comprehensive datasets [16–18] have demonstrated generalization across chemical space for pristine systems, with prediction accuracy approaching quantum mechanical calculations [19]. However, their application to defect systems remains an open challenge. First, MLIP-based defect modeling faces a data scarcity problem more severe than that of bulk materials. Defects require training datasets that capture a wide range of local atomic environments, including symmetry-breaking distortions and metastable configurations, all calculated at a high level of theory. Charged defects pose an additional challenge. Different charge states of the same defect can have distinct ground-state geometries because changes in electronic occupation modify the local bonding. However, standard MLIPs are built on the locality assumption [20], using descriptors that capture only the chemical species and local atomic environment within a truncated cutoff [21,22]. This representation contains no explicit information about the charge state, so configurations with similar local environments produce nearly identical descriptors despite corresponding to distinct charge states, thereby mapping the

*Contact author: xwwang20@gmail.com

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

multiple potential energy surfaces (PESs) into a single energy landscape. Consequently, existing studies largely focus on neutral defects [23–28], although these are not intrinsically simpler for MLIP modeling, as they often undergo substantial structural reconstructions driven by excess electrons or holes localized on the defect. Others train separate models for each charge state [29–31], rather than a single model that distinguishes charge states within a unified architecture. Despite growing developments in charge-aware ML models, their application to charged defects in periodic semiconductor supercells remains limited. Methods that explicitly model charge redistribution and long-range electrostatics, such as learned long-range descriptors [32], charge equilibration schemes [33,34], and polarizable electrostatic models [35,36], have been developed primarily for bulk liquids, molecular crystals, and interfaces. However, for charged defects, these approaches can suffer from ill-defined atomic charge partitioning [37], unphysical delocalization into the bulk [34], or the nonlinear scaling of Ewald summation [38]. Global charge conditioning, which encodes the total system charge through a learnable embedding [39], offers a simpler alternative that captures charge information and distinguishes charge states with negligible additional computational cost or complexity. Across both approaches, a single model capable of simultaneously distinguishing multiple charge states and subtle energy differences between competing defect configurations has not yet been demonstrated [39–41].

As a case study, we focus on Sb_2Se_3 , an emerging thin-film photovoltaic absorber composed of earth-abundant, low-toxicity elements, with promising optoelectronic properties [42]. Its point defects limit photovoltaic performance through nonradiative recombination and have complex energy landscapes [43]. The dominant intrinsic vacancies are amphoteric and undergo structural reconstructions across multiple charge states [6]. Together, these features make Sb_2Se_3 a challenging test case for defect-capable interatomic potentials.

In this work, we demonstrate that (i) foundation machine learning force fields, trained on bulk materials data, do not reliably identify defect ground states and often miss low-energy configurations; (ii) bespoke models with global defect charge embeddings can successfully identify ground-state configurations to within $0.05 \text{ \AA}$ and accurately predict defect formation energies and charge transition levels to within 0.02 eV ; (iii) a multifidelity (MF) approach that combines a high volume of low-cost reference (PBE) data with selective high-quality (HSE06) data results in a robust model that can identify global minima that are missed in low-precision searches. The resulting defect-capable force fields act as fast surrogate models that can describe dynamic processes that are beyond the limits of traditional defect modeling approaches. The MF strategy with charge embeddings

also provides a solid foundation toward the development of foundation models for defective crystals.

II. RESULTS AND DISCUSSION

A. Limitation of foundation models for charged defects

Foundation models trained on massive bulk datasets have enabled rapid property prediction for crystalline materials, yet their applicability to the complex landscape of charged point defects remains unexplored. To address this gap, we first benchmarked four state-of-the-art equivariant atomistic force fields (MACE-MH-1 [12,22], MACE-MPA-0 [13], GRACE-2L-OMAT [14], and MatterSim (v1 5M) [15]) against PBE density functional theory (DFT) calculations, a level of theory chosen to match the training data of these foundation models. We use the antimony vacancy in Sb_2Se_3 as a representative system due to its charge-state multiplicity ($q = -3$ to $+1$) and complex defect landscapes [43]. Sb_2Se_3 crystallizes in the $Pnma$ space group with two inequivalent Sb sites and three inequivalent Se sites (Fig. S1 in Supplemental Material [44]). The foundation-model benchmarks reported here were performed for $V_{\text{Sb}(1)}$. To systematically explore these PESs, we employed SHAKENBREAK [9] to generate diverse initial bond-distorted configurations for each charge state. SHAKENBREAK determines the number of nearest-neighbor bonds to distort according to the defect’s effective excess charge, leading to different initial geometries for different charge states. The same initial configurations were used for both foundation model and DFT relaxations in each charge state to ensure a fair comparison. For the fully ionized state ($q = -3$), where the effective excess charge is zero, the initial structures from the neutral state were used to ensure a comprehensive search. These structures were subsequently relaxed using each foundation model to identify the predicted ground-state structures.

Figure 1 illustrates the failure of bulk-trained foundation models when applied to charged point defects. As shown in Fig. 1(a), none of the tested models consistently identifies the DFT-predicted ground-state geometry of $V_{\text{Sb}(1)}$ across all charge states. The structural root mean square deviation (RMSD) generally ranges from 0.2 to $0.4 \text{ \AA}$, well beyond the $0.1 \text{ \AA}$ threshold for reliable structure identification. Although GRACE-2L-OMAT and MACE-MH-1 appear to match the DFT structure at $q = -2$ and $q = -1$, respectively, we show below that these are accidental agreements limited to a single charge state and resulting from error cancellation rather than physically accurate modeling.

We attribute these failures to two key discrepancies between current foundation models and defect physics: a domain gap in the training data and the absence of charge-state descriptors in the model architectures. First, foundation force fields are trained predominantly on closed-shell




FIG. 1. Failure of foundation machine learning interatomic potentials (MLIPs) to capture defect structures. (a) Root mean square deviation (RMSD) of MLIP-predicted ground-state structures relative to density functional theory (DFT) calculations using the PBE functional for $V_{\text{Sb}(1)}$ in Sb_2Se_3 across five charge states. For each charge state, data points are slightly offset along the horizontal axis for clarity. The green-shaded region ($\text{RMSD} \leq 0.1$ Å) indicates successful structural identification. (b),(c) Comparison of defect potential energy surfaces (PESs) and structural predictions for (b) the neutral ($q = 0$) and (c) fully ionized ($q = -3$) charge states of $V_{\text{Sb}(1)}$. The PESs are mapped as a function of the bond distortion percentage, with energies referenced to the minimum of each respective landscape. In these panels, data points for the foundation models and the DFT reference are colored according to the RMSD of the relaxed atomic positions relative to the DFT reference ground state, with circles indicating $\text{RMSD} \leq 0.1$ Å and crosses denoting $\text{RMSD} > 0.1$ Å. The bottom rows compare the local atomic environment of the DFT global minimum with representative MLIP-predicted minima.

bulk structures, so they tend to treat point defects as steric voids instead of electronically active centers that drive local reconstruction. In the neutral state of $V_{\text{Sb}(1)}$ [Fig. 1(b)], removing an Sb atom introduces three holes that drive a covalent reconstruction into a Se trimer with a bond angle of 110° through electron sharing. Since these electronic environments are absent from the bulk training data, none of the tested MLIPs reproduces this reconstruction correctly. Instead, they tend to relax into high-symmetry, bulklike configurations or fall into a spurious local minimum with a bond angle near 99° . This shows that even neutral defects, which dominate the existing MLIP literature, are already challenging cases for current foundation models. Second, standard MLIPs represent the total energy as a function of atomic positions and species without incorporating any charge information. As a result, they assign identical local forces and energies to configurations with the same local geometry, collapsing charge-dependent PESs into a single landscape. Even for the fully ionized $q = -3$ state [Fig. 1(c)], which corresponds to a closed-shell configuration with no localized charge and minimal structural reconstruction, the models still fail to identify the DFT reference minimum. Although they perform slightly better in this regime with lower

RMSD values, all predictions still exceed the 0.1 Å threshold. This improvement arises from the models' preference for low-distortion geometries close to the bulk structure environments. Without explicit descriptors for the charge state, the models cannot distinguish the charge-induced relaxation from the pristine structure.

Given these limitations, the isolated successes of GRACE-2L-OMAT ($q = -2$) and MACE-MH-1 ($q = -1$) in Fig. 1(a) do not reflect physically meaningful predictions. For $V_{\text{Sb}(1)}$, the $q = -2$ state adopts a geometry similar to the fully ionized $q = -3$ case, while the $q = -1$ state is nonmagnetic with reduced covalent reconstruction compared to the neutral state (Fig. S4 in Supplemental Material [44]). These intermediate charge states coincidentally match foundation models' bias toward bulklike configurations. As shown in Figs. 1(b) and 1(c) (and Fig. S3 in Supplemental Material [44] for the other charge states), the models systematically predict flatter PESs compared to DFT, indicating incorrect forces. Consequently, such coincidental structural agreements resulting from error cancellation are not transferable and cannot be relied upon for predictive defect engineering. Furthermore, subsequent DFT relaxations initiated from foundation-model minima do not consistently converge to




FIG. 2. Schematic diagram for the integration of global charge embeddings for point defects in machine learning interatomic potentials (MACE [22] architecture). A learned vector corresponding to the total charge q_{tot} is added to each atomic species embedding during feature initialization. This charge conditioning propagates through the interaction layers, where node features $h_i^{(t)}$ are iteratively updated via a learned interaction function f_θ that aggregates messages from edge embeddings e_{ij} and neighboring features $h_j^{(t)}$. The interaction energy E_{inter} is computed by summing the readout outputs $g_\theta^{(t)}(h_i^{(t)})$ over all atoms i and message-passing iterations t , while a geometry-independent energy term E_{emb} is extracted from the initial charge-enriched embeddings. The total energy E_{tot} is defined as the sum of the atomic reference energy E_0 , the interaction energy E_{inter} , and the embedding energy E_{emb} , with $E_0 = \sum_i E_{\text{ref}}(Z_i)$ denoting the sum of atomic reference energies. Atomic forces are computed identically to standard MACE, with $\mathbf{F}_i = -\partial E_{\text{tot}}/\partial \mathbf{r}_i = -\partial E_{\text{inter}}/\partial \mathbf{r}_i$, since both the atomic reference energy E_0 and the embedding energy E_{emb} are independent of atomic positions.

the DFT reference ground states across all charge states (Sec. S3.1.2 and Fig. S4 in Supplemental Material [44]).

These results indicate that foundation models trained solely on bulk structures, even when highly accurate for ideal crystals [19], are fundamentally unreliable for defect physics in semiconductors and insulators. This limitation applies across all charge states: not only to charged defects, where charge-state descriptors are absent, but also to neutral defects, where excess carriers can drive reconstructions outside the training domain. Compared to standard MLIPs for pristine materials, reliable descriptions of charged point defects require two additional factors. First, the training data must cover relevant configurations beyond the distribution found in pristine crystals. Second, the model architecture must explicitly encode electronic state information to distinguish between different electronic configurations of the same defect. This can be achieved through approaches such as global embeddings of electronic state variables, including charge and spin, or explicit electrostatic models [32,33,66]. Further benchmarks show that charge-aware foundation models that explicitly impose the total charge but lack relevant defect training data still fail to reproduce the DFT ground states (Sec. S3.2 of Supplemental Material [44]).

B. Charged defect prediction with global embeddings

Building on this insight, we trained a bespoke joint model that satisfies both conditions. The training dataset includes defect configurations across multiple charge

states simulated by DFT, and the architecture combines a global embedding of the total defect charge with atomic features. The modified architecture (Fig. 2) was implemented in MACE v0.3.15 [22], with charge-dependent components highlighted in orange. MACE was chosen for its many-body equivariant message-passing architecture with high-order geometric features. The total charge q_{tot} is mapped to a learnable vector that is added to the atomic species embeddings. The charge information is learned through two pathways: (1) it propagates through the message-passing layers, enabling the network to learn charge-dependent interaction energies E_{inter} that govern the potential energy surfaces and, through their gradients, the forces driving defect reconstruction, and (2) the charge-enriched embedding is read out directly prior to message passing to provide a geometry-independent energy contribution E_{emb} , capturing the charge-state energy shifts that are not represented by local atomic geometry.

While the benchmarks in Fig. 1 used the PBE functional to match the training domain of foundation models, accurate defect physics often requires a higher level of theory. Semilocal functionals like PBE systematically underestimate the band gaps of insulators and suffer from delocalization errors that impact defect charge distribution and energetics [67,68]. We therefore train the model using hybrid HSE06 DFT data, providing a robust reference for defect predictions. The model is trained on raw DFT total energies and forces without finite-size electrostatic corrections. These corrections are determined by the charge

state and supercell geometry, so they do not affect atomic forces or the topology of the potential energy surface (PES). The resulting model therefore reproduces the defect landscape at the HSE06 reference level, with accuracy bounded by the underlying calculations.

While $V_{\text{Sb}(1)}$ provides a clear demonstration of foundation-model failure due to its pronounced covalent reconstruction, we demonstrate the charge-embedded capabilities using $V_{\text{Se}(1)}$ in Sb_2Se_3 as an example due to its challenging PES (Fig. S8 in Supplemental Material [44]). For the neutral charge state, DFT searches using SHAKENBREAK [9] reveal that only 3 out of 14 bond distortions successfully identify the global minimum. Furthermore, the energy difference between the stable and metastable configurations is as small as ~ 0.1 eV for a 59-atom supercell (1–2 meV/atom). We generated training data by relaxing bond-distorted configurations across five charge states, incorporating full relaxation trajectories to sample the PES comprehensively (see Methods in Sec. S2 of Supplemental Material [44]).

To visualize the role of global charge embeddings, we examined the model’s latent representation of the defect. Figure 3(a) maps the descriptor space as the average difference between the MACE feature vectors of each defect configuration and the pristine reference. This metric captures the collective local environment perturbations induced by the vacancy. Without global charge embeddings [Fig. 3(a), left], configurations from different charge

states overlap significantly in the two-dimensional (2D) descriptor space, even though the first two principal components capture 79.9% of the variance. This highlights a fundamental limitation of MLIPs without charge attributes, which rely solely on elemental and structural differences to distinguish between charge states.

Specifically, in bulklike regions far from the defect site, local environments remain nearly identical across charge states, while total energies differ due to different electronic occupation. This leads to a nonunique mapping between local descriptors and energy labels. In contrast, with charge embeddings included (right panel), the descriptors segregate into distinct, charge-specific clusters while retaining sensitivity to local structural variations. This separation allows the model to learn charge-dependent energy offsets for structurally similar configurations, enabling accurate predictions across all five charge states [with test-set $\text{RMSE}_E = 0.48$ meV/atom and $\text{RMSE}_F = 20.15$ meV/Å, Fig. 3(b)]. As shown in Fig. 3(c), the charge-embedded model successfully identifies the HSE06 reference ground states for all cases, with global RMSD values within 0.05 Å and well below the 0.1 Å threshold for reliable structure identification. The model reproduces the local vacancy environments with negligible deviation, capturing the subtle charge-dependent reconstructions that foundation models fail to describe. Beyond ground-state structures, the model provides a quantitatively accurate description of the energetics of the defects. It correctly




FIG. 3. (a) Principal component analysis (PCA) visualization of the atomic descriptor space for $V_{\text{Se}(1)}$ in Sb_2Se_3 across five charge states without (left) and with (right) global charge embeddings. The latent space is constructed using both training and test configurations. Percentages in parentheses indicate the explained variance for each component. (b) Parity plots of total energies and atomic forces predicted by the charge-embedded MACE model compared with HSE06 DFT on the test dataset. Root mean square errors (RMSEs) are reported over the full test set across five charge states. (c) Comparison between the charge-embedded MACE model and HSE06 DFT for ground-state defect geometries (top) and thermodynamic transition levels (bottom). For each charge state, the global minimum was identified by relaxing a set of bond-distorted structures. Key bond lengths are labeled in Å. Thermodynamic transition levels obtained from HSE06 DFT are shown as orange dashed lines, while MLIP predictions are shown as black solid lines.




FIG. 4. Schematic diagram of structure searching for neutral $V_{Se(1)}$ using (a) a standard workflow with SHAKENBREAK [9] and (b) a multifidelity (MF) strategy with an MLIP. The green curves correspond to high-fidelity (HF) PESs calculated using HSE06/DFT with converged k -point sampling. The red star and gray circle denote the global minimum and a metastable minimum on this reference PES, respectively. (a) The black curve represents the coarse PES sampling using HSE06 with Γ -point sampling, and the gray dashed arrow indicates the high-fidelity (HF) relaxation starting from the coarse sampling minimum. (b) The blue curve represents the low-fidelity PES obtained from PBE with converged k -point sampling. Blue crosses mark HF single-point calculations performed on a subset of the PBE-relaxed structures. Orange vertical arrows illustrate the Δ -learning corrections learned by the MF model.

distinguishes metastable configurations with energy differences as small as 1 meV/atom (Fig. S8 in Supplemental Material [44]) and predicts thermodynamic transition levels (Fig. 3, bottom) that closely match the HSE06 reference across all charge transitions, with a maximum deviation of only 0.015 eV (Table S3 in Supplemental Material [44]). These results confirm that combining explicit charge embeddings with targeted defect training data enables a joint model to achieve *ab initio* accuracy across multiple charge states, even for highly complex defect landscapes. Additional benchmarks show that explicitly modeling atomic charge distribution or electrostatics with QET [69], MACE-POLAR [35], and MACELES [32] gives a less accurate description of the DFT defect PESs (Sec. S5 of Supplemental Material [44]).

C. Improved accuracy and efficiency via multifidelity training

While the charge-embedded model discussed above achieves accurate defect predictions, its reliance on full relaxation trajectories at the hybrid-functional level presents a significant bottleneck for large-scale applications. To reduce the cost of expensive hybrid-functional relaxations, commonly used workflows like SHAKENBREAK [9] use a two-stage strategy [70]. The PES is initially explored using lower-cost methods (e.g., Γ -point sampling) to identify the minima, followed by higher-fidelity calculations only for the identified minima

[Fig. 4(a)]. However, this approach assumes that the coarse-level search adequately reproduces the shape of the HF PES and its global minimum. Furthermore, even when limited to Γ -point sampling, the computational cost remains substantial for complex defect landscapes (Sec. S6.3 of Supplemental Material [44]).

Motivated by these limitations, we developed an MF training strategy [Fig. 4(b)] that samples the PES densely at the low-cost PBE level while incorporating a sparse subset ($\sim 10\%$) of HF HSE06 corrections selected using trajectory-stratified sampling. The model learns the functional-dependent difference (Δ) from HSE06 single-point energies on selected PBE structures through a joint training framework with a shared latent representation and separate (low fidelity) PBE and (high fidelity) HSE06 readout channels. Using neutral $V_{Se(1)}$ as an example, the MF approach reveals the limitation of standard defect searching workflows. As shown in Fig. 4, conventional coarse search using HSE06 at the Γ -point (Fig. S15 in Supplemental Material [44]) identifies only a metastable minimum (gray circle). In contrast, the MF-trained model discovers a distinct global minimum 0.02 eV lower in energy (red star) by learning the functional-dependent correction required to stabilize this basin [Fig. 4(b)]. This prediction is further confirmed by hybrid-functional DFT with converged k -point sampling, which reproduces the ML-predicted minimum within 9 meV. Beyond improved accuracy, this MF approach substantially reduces the computational cost of comprehensive PES exploration,

with the low-fidelity PBE sampling being approximately 3 orders of magnitude cheaper than HSE06 calculations using the same converged k -point sampling (Sec. S6.3 of Supplemental Material [44]), enabling systematic studies across multiple charge states and defect types.

Moreover, once trained, the MF model serves as an efficient surrogate for the hybrid-functional PES, enabling access to properties that require dense sampling of the defect landscapes that are otherwise prohibitively expensive at the hybrid-functional level. We demonstrate this by calculating configuration coordinate diagrams (CCDs), which describe charge carrier capture processes that limit the power conversion efficiency in optoelectronic devices [71]. Constructing a configuration coordinate diagram (CCD) requires evaluating the energies of two charge states along a continuous reaction coordinate connecting their ground-state geometries, with sufficient sampling to capture their curvature and the crossing region. Our MF model reproduces HSE06 reference CCDs with high precision at negligible inference cost (Fig. S17 in Supplemental Material [44]). This approach can be extended to other thermodynamic and spectroscopic properties, such as finite-temperature free energies [29,72], luminescence lineshapes [28,73], and defect dynamics [74], offering new possibilities for accelerated defect screening and property prediction with hybrid-level accuracy.

III. CONCLUSION

In conclusion, as machine learning force fields have matured in their description of bulk materials, we evaluate their generalizability to the out-of-distribution regime of charged point defects. Using vacancies in Sb_2Se_3 as a case study, we demonstrate that while current foundation models fail to capture the defect energy landscapes, training on defect configurations with global charge embeddings at a fixed defect supercell size can recover ground-state structures to within 0.05 Å, with defect formation energies and thermodynamic transition levels reproduced at hybrid-level DFT accuracy to within 0.02 eV. We further introduce a multifidelity training scheme that not only discovers global minima missed by standard defect screening workflows but also provides a computationally efficient surrogate model that can be used to access larger configurational spaces. Together, these strategies offer a practical framework for bridging the gap between *ab initio* accuracy and large-scale defect modeling, enabling high-throughput prediction of charged defect structures, energetics, and related physical properties. For photovoltaic absorbers such as Sb_2Se_3 , this makes it feasible to search charged defect landscapes with hybrid-level accuracy and to evaluate the configurations that govern carrier recombination. More broadly, our work lays the groundwork for future defect foundation models that generalize across materials, charge states, and levels of

theory. Broad coverage of defect environments, learnable charge conditioning, and multifidelity training together provide a practical basis for extending pristine-material foundation models to charged-defect applications.

ACKNOWLEDGMENTS

We thank Alex M. Ganose for helpful discussions and support. This research received support through Schmidt Sciences, LLC. I. M.-L. thanks Imperial College London (ICL) for funding a President's Ph.D. scholarship. Via our membership of the UK's HEC Materials Chemistry Consortium, which is funded by EPSRC (EP/X035859/1), this work used the ARCHER2 UK National Supercomputing Service [75]. We are grateful to the UK Materials and Molecular Modelling Hub for computational resources, which are partially funded by EPSRC (EP/T022213/1, EP/W032260/1, and EP/P020194/1).

DATA AVAILABILITY

The data that support the findings of this article are publicly available [76]. The MACE software package (v0.3.15) used in this work is open-source and available at [77].

-
- [1] A. Walsh and A. Zunger, *Instilling defect tolerance in new compounds*, *Nat. Mater.* **16**, 964 (2017).
 - [2] A. Van der Ven, Z. Deng, S. Banerjee, and S. P. Ong, *Rechargeable alkali-ion battery materials: Theory and computation*, *Chem. Rev.* **120**, 6977 (2020).
 - [3] C. Freysoldt, B. Grabowski, T. Hickel, J. Neugebauer, G. Kresse, A. Janotti, and C. G. Van de Walle, *First-principles calculations for point defects in solids*, *Rev. Mod. Phys.* **86**, 253 (2014).
 - [4] I. Mosquera-Lois and S. R. Kavanagh, *In search of hidden defects*, *Matter* **4**, 2602 (2021).
 - [5] S. R. Kavanagh, D. O. Scanlon, A. Walsh, and C. Freysoldt, *Impact of metastable defect structures on carrier recombination in solar cells*, *Faraday Discuss.* **239**, 339 (2022).
 - [6] X. Wang, S. R. Kavanagh, D. O. Scanlon, and A. Walsh, *Four-electron negative-U vacancy defects in antimony selenide*, *Phys. Rev. B* **108**, 134102 (2023).
 - [7] F. El-Mellouhi, N. Mousseau, and P. Ordejón, *Sampling the diffusion paths of a neutral vacancy in silicon with quantum mechanical calculations*, *Phys. Rev. B* **70**, 205202 (2004).
 - [8] A. G. Squires, L. Ganeshkumar, C. N. Savory, S. R. Kavanagh, and D. O. Scanlon, *Oxygen dimerization as a defect-driven process in bulk LiNiO_2* , *ACS Energy Lett.* **9**, 4180 (2024).
 - [9] I. Mosquera-Lois, S. R. Kavanagh, A. Walsh, and D. O. Scanlon, *ShakeNBreak: Navigating the defect configurational landscape*, *J. Open Source Softw.* **7**, 4817 (2022).
 - [10] C. J. Pickard and R. J. Needs, *Ab initio random structure searching*, *J. Phys. Condens. Matter* **23**, 053201 (2011).

- [11] X. Wang, S. R. Kavanagh, and A. Walsh, *Sulfur vacancies limit the open-circuit voltage of Sb₂S₃ solar cells*, *ACS Energy Lett.* **10**, 161 (2025).
- [12] I. Batatia, C. Lin, J. Hart, E. Kassoar, A. M. Elena, S. W. Norwood, T. Wolf, and G. Csányi, *Cross learning between electronic structure theories for unifying molecular, surface, and inorganic crystal foundation force fields*, *arXiv:2510.25380*.
- [13] I. Batatia, P. Benner, Y. Chiang, A. M. Elena, D. P. Kovács, J. Riebesell, X. R. Advincula, M. Asta, M. Avaylon, W. J. Baldwin *et al.*, *A foundation model for atomistic materials chemistry*, *J. Chem. Phys.* **163**, 184110 (2025).
- [14] Y. Lysogorskiy, A. Bochkarev, and R. Drautz, *Graph atomic cluster expansion for foundational machine learning interatomic potentials*, *npj Comput. Mater.* **12**, 114 (2026).
- [15] H. Yang, C. Hu, Y. Zhou, X. Liu, Y. Shi, J. Li, G. Li, Z. Chen, S. Chen, C. Zeni *et al.*, *MatterSim: A deep learning atomistic model across elements, temperatures and pressures*, *arXiv:2405.04967*.
- [16] A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder *et al.*, *Commentary: The materials project: A materials genome approach to accelerating materials innovation*, *APL Mater.* **1**, 011002 (2013).
- [17] L. Barros-Luque, M. Shuaibi, X. Fu, B. M. Wood, M. Dzamba, M. Gao, A. Rizvi, M. Uyttendaele, C. L. Zitnick, and Z. W. Ulissi, *The Open Materials 2024 (OMat24) inorganic materials dataset and models*, *Nat. Comput. Sci.* **6**, 642 (2026).
- [18] J. Schmidt, T. F. T. Cerqueira, A. H. Romero, A. Loew, F. Jäger, H.-C. Wang, S. Botti, and M. A. L. Marques, *Improving machine-learning models in materials science through large datasets*, *Mater. Today Phys.* **48**, 101560 (2024).
- [19] A. Loew, D. Sun, H.-C. Wang, S. Botti, and M. A. L. Marques, *Universal machine learning interatomic potentials are ready for phonons*, *npj Comput. Mater.* **11**, 178 (2025).
- [20] E. Prodan and W. Kohn, *Nearsightedness of electronic matter*, *Proc. Natl. Acad. Sci. U.S.A.* **102**, 11635 (2005).
- [21] J. Behler and M. Parrinello, *Generalized neural-network representation of high-dimensional potential-energy surfaces*, *Phys. Rev. Lett.* **98**, 146401 (2007).
- [22] I. Batatia, D. P. Kovács, G. N. C. Simm, C. Ortner, and G. Csányi, *MACE: Higher order equivariant message passing neural networks for fast and accurate force fields*, *Adv. Neural Inf. Process. Syst.* **35**, 11423 (2022).
- [23] A. Mannodi-Kanakkithodi, X. Xiang, L. Jacoby, R. Biegaj, S. T. Dunham, D. R. Gamelin, and M. K. Y. Chan, *Universal machine learning framework for defect predictions in zinc blende semiconductors*, *Patterns* **3**, 100450 (2022).
- [24] I. Mosquera-Lois, S. R. Kavanagh, A. M. Ganose, and A. Walsh, *Machine-learning structural reconstructions for accelerated point defect calculations*, *npj Comput. Mater.* **10**, 121 (2024).
- [25] S. R. Kavanagh, *Identifying split vacancy defects with machine-learned foundation models and electrostatics*, *J. Phys. Energy* **7**, 045002 (2025).
- [26] J. Zhou, X. Li, M. Huang, and S. Chen, *One-defect one-potential strategy for accurate machine learning prediction of phonons in defect-containing supercells*, *Phys. Rev. B* **112**, 235205 (2025).
- [27] Z. Yang, X. Liu, X. Zhang, P. Huang, K. S. Novoselov, and L. Shen, *Modeling crystal defects using defect informed neural networks*, *npj Comput. Mater.* **11**, 229 (2025).
- [28] M. E. Turiansky, J. L. Lyons, and N. Bernstein, *Machine learning phonon spectra for fast and accurate optical lineshapes of defects*, *ACS Nano* **20**, 7454 (2026).
- [29] I. Mosquera-Lois, J. Klarbring, and A. Walsh, *Point defect formation at finite temperatures with machine learning force fields*, *Chem. Sci.* **16**, 8878 (2025).
- [30] I. Mosquera-Lois and A. Walsh, *Dynamic vacancy levels in CsPbCl₃ obey equilibrium defect thermodynamics*, *PRX Energy* **4**, 043008 (2025).
- [31] M. H. Rahman, M. Biswas, and A. Mannodi-Kanakkithodi, *DeFecT-FF: A machine learning force field framework for high throughput defect modeling in CdTe-based solar cells*, *Phys. Chem. Chem. Phys.* **28**, 10718 (2026).
- [32] B. Cheng, *Latent Ewald summation for machine learning of long-range interactions*, *npj Comput. Mater.* **11**, 80 (2025).
- [33] T. W. Ko, J. A. Finkler, S. Goedecker, and J. Behler, *A fourth-generation high-dimensional neural network potential with accurate electrostatics including non-local charge transfer*, *Nat. Commun.* **12**, 398 (2021).
- [34] M. Vondrák, W. J. Baldwin, G. Csányi, K. Reuter, and J. T. Margraf, *Integrating charge equilibration with equivariant machine-learning interatomic potentials*, *J. Chem. Theory Comput.* **22**, 6874 (2026).
- [35] I. Batatia, W. J. Baldwin, D. Kuryla, J. Hart, E. Kassoar, A. M. Elena, H. Moore, M. J. Gawkowski, B. X. Shi, V. Kapil *et al.*, *MACE-POLAR-1: A polarisable electrostatic foundation model for molecular chemistry*, *arXiv:2602.19411*.
- [36] W. J. Baldwin, I. Batatia, M. Vondrák, J. T. Margraf, and G. Csányi, *Design space of self-consistent electrostatic machine learning interatomic potentials*, *arXiv:2603.14700*.
- [37] B. Wang, S. L. Li, and D. G. Truhlar, *Modeling the partial atomic charges in inorganometallic molecules and solids and charge redistribution in lithium-ion cathodes*, *J. Chem. Theory Comput.* **10**, 5640 (2014).
- [38] A. Y. Toukmaji and J. A. Board Jr., *Ewald summation techniques in perspective: A survey*, *Comput. Phys. Commun.* **95**, 73 (1996).
- [39] B. M. Wood, M. Dzamba, X. Fu, M. Gao, M. Shuaibi, L. Barroso-Luque, K. Abdelmaqsoud, V. Gharakhanyan, J. Kitchin, D. S. Levine *et al.*, *UMA: A family of universal models for atoms*, *Adv. Neural Inf. Process. Syst.* **38**, 129391 (2025).
- [40] K. Shimizu, Y. Dou, E. F. Arguëlles, T. Moriya, E. Minamitani, and S. Watanabe, *Using neural network potentials to study defect formation and phonon properties of nitrogen vacancies with multiple charge states in GaN*, *Phys. Rev. B* **106**, 054108 (2022).
- [41] S. Kiyohara, C. Shibui, S. Bae, and Y. Kumagai, *Machine-learning prediction of charged-defect formation energies from crystal structures*, *Phys. Rev. Lett.* **135**, 246101 (2025).
- [42] A. Mavlonov, T. Razykov, F. Raziq, J. Gan, J. Chantana, Y. Kawano, T. Nishimura, H. Wei, A. Zakutayev, T. Minemoto *et al.*, *A review of Sb₂Se₃ photovoltaic absorber materials and thin-film solar cells*, *Sol. Energy* **201**, 227 (2020).

- [43] X. Wang, S. R. Kavanagh, D. O. Scanlon, and A. Walsh, *Upper efficiency limit of Sb₂Se₃ solar cells*, *Joule* **8**, 2105 (2024).
- [44] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/tzwd-cl3y> for detailed methodology and additional results on foundation model benchmarking, charge-embedded models, and multi-fidelity training, which includes Refs. [45–65].
- [45] W. Kohn and L. J. Sham, *Self-consistent equations including exchange and correlation effects*, *Phys. Rev.* **140**, A1133 (1965).
- [46] R. M. Dreizler and E. K. Gross, *Density functional theory of relativistic systems*, in *Density Functional Theory* (Springer, Berlin, Heidelberg, 1990), pp. 245–271.
- [47] G. Kresse and J. Furthmüller, *Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set*, *Phys. Rev. B* **54**, 11169 (1996).
- [48] G. Kresse and D. Joubert, *From ultrasoft pseudopotentials to the projector augmented-wave method*, *Phys. Rev. B* **59**, 1758 (1999).
- [49] S. Grimme, *Accurate description of van der Waals complexes by density functional theory including empirical corrections*, *J. Comput. Chem.* **25**, 1463 (2004).
- [50] X. Wang, Z. Li, S. R. Kavanagh, A. M. Ganose, and A. Walsh, *Lone pair driven anisotropy in antimony chalcogenide semiconductors*, *Phys. Chem. Chem. Phys.* **24**, 7195 (2022).
- [51] J. P. Perdew, K. Burke, and M. Ernzerhof, *Generalized gradient approximation made simple*, *Phys. Rev. Lett.* **77**, 3865 (1996).
- [52] J. Heyd, G. E. Scuseria, and M. Ernzerhof, *Hybrid functionals based on a screened Coulomb potential*, *J. Chem. Phys.* **118**, 8207 (2003).
- [53] X. Liu, X. Xiao, Y. Yang, D. Xue, D. Li, C. Chen, S. Lu, L. Gao, Y. He, M. C. Beard *et al.*, *Enhanced Sb₂Se₃ solar cell performance through theory-guided defect control*, *Prog. Photovolt.* **25**, 861 (2017).
- [54] C. N. Savory and D. O. Scanlon, *The complex defect chemistry of antimony selenide*, *J. Mater. Chem. A* **7**, 10739 (2019).
- [55] A. Stoliaroff, A. Lecomte, O. Rubel, S. Jobic, X. H. Zhang, C. Latouche, and X. Rocquefelte, *Deciphering the role of key defects in Sb₂Se₃, a promising candidate for chalcogenide-based solar cells*, *ACS Appl. Energy Mater.* **3**, 2496 (2020).
- [56] S. R. Kavanagh, A. G. Squires, A. Nicolson, I. Mosquera-Lois, A. M. Ganose, B. Zhu, K. Brlec, A. Walsh, and D. O. Scanlon, *Doped: Python toolkit for robust and repeatable charged defect supercell calculations*, *J. Open Source Softw.* **9**, 6433 (2024).
- [57] I. Batatia, S. Batzner, D. P. Kovács, A. Musaelian, G. N. C. Simm, R. Drautz, C. Ortner, B. Kozinsky, and G. Csányi, *The design space of E(3)-equivariant atom-centred interatomic potentials*, *Nat. Mach. Intell.* **7**, 56 (2025).
- [58] D. P. Kingma and J. Ba, *Adam: A method for stochastic optimization*, in *International Conference on Learning Representations* (Computational and Biological Learning Society, San Diego, CA, USA, 2015).
- [59] P. J. Huber, *Robust estimation of a location parameter, in Breakthroughs in Statistics: Methodology and Distribution* (Springer, New York, 1992), pp. 492–518.
- [60] P. Izmailov, D. Podoprikhin, T. Garipov, D. P. Vetrov, and A. G. Wilson, *Averaging weights leads to wider optima and better generalization*, in *Proceedings of the Conference on Uncertainty in Artificial Intelligence* (AUAI Press, Corvallis, OR, USA, 2018), pp. 876–885.
- [61] R. Fletcher, *Practical Methods of Optimization* (John Wiley & Sons, New York, 1987).
- [62] A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dułak, J. Friis, M. N. Groves, B. Hammer, C. Hargus *et al.*, *The atomic simulation environment—a Python library for working with atoms*, *J. Phys.: Condens. Matter* **29**, 273002 (2017).
- [63] H. W. Kuhn, *The Hungarian method for the assignment problem*, *Nav. Res. Logist. Q.* **2**, 83 (1955).
- [64] Y. Kumagai and F. Oba, *Electrostatics-based finite-size corrections for first-principles point defect calculations*, *Phys. Rev. B* **89**, 195205 (2014).
- [65] D. S. Levine, M. Shuaibi, E. W. C. Spotte-Smith, M. G. Taylor, M. R. Hasyim, K. Michel, I. Batatia, G. Csányi, M. Dzamba, P. Eastman *et al.*, *The open molecules 2025 (OMol25) dataset, evaluations, and models*, *arXiv:2505.08762*.
- [66] Z. Fan, B. Tang, E. Berger, E. Berger, E. Fransson, K. Xu, Z. Yan, Z. Liu, Z. Song, H. Dong *et al.*, *qNEP: A highly efficient neuroevolution potential with dynamic charges for large-scale atomistic simulations*, *J. Chem. Theory Comput.* **22**, 4787 (2026).
- [67] P. Mori-Sánchez, A. J. Cohen, and W. Yang, *Localization and delocalization errors in density functional theory and implications for band-gap prediction*, *Phys. Rev. Lett.* **100**, 146401 (2008).
- [68] A. G. Squires, S. R. Kavanagh, A. Walsh, and D. O. Scanlon, *Guidelines for robust and reproducible point defect simulations in crystals*, *Nat. Rev. Mater.* **11**, 469 (2026).
- [69] T. W. Ko, R. Liu, A. R. Mishra, Z. Yu, J. Qi, and S. P. Ong, *A linear-scaling, charge-aware foundation potential for atomistic simulations*, *arXiv:2511.07249*.
- [70] I. Mosquera-Lois, S. R. Kavanagh, A. Walsh, and D. O. Scanlon, *Identifying the ground state structures of point defects in solids*, *npj Comput. Mater.* **9**, 25 (2023).
- [71] A. Alkauskas, Q. Yan, and C. G. Van de Walle, *First-principles theory of nonradiative carrier capture via multi-phonon emission*, *Phys. Rev. B* **90**, 075202 (2014).
- [72] T. Hainner, E. Berger, E. Berger, O. Hildeberg, P. Erhart, and J. Wiktor, *Thermal stabilization of defect charge states and finite-temperature charge transition levels*, *arXiv:2512.15463*.
- [73] C. Linderålv, N. Österbacka, J. Wiktor, and P. Erhart, *Optical line shapes of color centers in solids from classical autocorrelation functions*, *npj Comput. Mater.* **11**, 101 (2025).
- [74] V. Tyagi, M. Pols, G. Brocks, and S. Tao, *Tracing ion migration in halide perovskites with machine learned force fields*, *J. Phys. Chem. Lett.* **16**, 5153 (2025).
- [75] <http://www.archer2.ac.uk>.
- [76] X. Wang, I. Mosquera-Lois, and A. Walsh, *Supporting data for “multi-fidelity machine learning interatomic potentials for charged point defects”* (2026), [10.5281/zenodo.22756145](https://doi.org/10.5281/zenodo.22756145).
- [77] <https://github.com/ACEsuit/mace>.