

## Visualizing phonon edge states on molybdenum disulfide

Patrick Giese<sup>1</sup>, Mathias Stokkebye Nissen<sup>2</sup>, Stig Helveg<sup>3</sup>, and Jakob Schiøtz<sup>1,\*</sup><sup>1</sup>Department of Physics, *Technical University of Denmark*, DK-2800 Kongens Lyngby, Denmark<sup>2</sup>Department of Energy Conversion and Storage, *Technical University of Denmark*, DK-2800 Kongens Lyngby, Denmark<sup>3</sup>Center for Visualizing Catalytic Processes (VISION), Department of Physics, *Technical University of Denmark*, DK-2800 Kongens Lyngby, Denmark

(Received 10 September 2025; accepted 3 May 2026; published 12 June 2026)

We employ molecular dynamics (MD) simulations to study atom vibrational amplitudes in carbon-supported molybdenum disulfide (MoS<sub>2</sub>) nanoparticles. Enhanced and correlated atom vibrational amplitudes are observed as the nanoparticle edges are approached from the bulk, consistent with recent experimental high-resolution transmission electron microscopy (HRTEM) observations by Chen *et al.* [*Nat. Commun.* **12**, 5007 (2021)]. Analysis of phonon modes in finite systems explains the experimental observation by low-energy phonon modes confined at the nanoparticle edge, underscoring the need of full MD modeling for accurate HRTEM image interpretation. Noticeably, we introduce a workflow for training equivariant neural network-based machine learning potentials using limited density functional theory (DFT) calculations. This approach effectively captures both covalent and van der Waals interactions, enabling accurate extrapolations of DFT calculations to larger systems with built-in error estimation.

DOI: [10.1103/f4y5-bps4](https://doi.org/10.1103/f4y5-bps4)

Electron microscopy has become a powerful tool for visualizing matter at the atomic scale. Ultimately, it now reaches a spatial resolution down to approximately 50 pm in both broad- and focused-beam modes [1,2]. This resolution level is dictated by the atomic potentials and their vibrational smearing rather than the electron optical performance [3,4]. However, while atom vibrations are commonly described using an independent atom model [5,6], there is now evidence that correlations between atomic vibrations may be directly observable. Chen *et al.* [3] measured phase contrast decay from the bulk toward the edge of a single-layer MoS<sub>2</sub> nanoparticle on a few-layer graphene support and attributed this to vibrations that must be correlated.

Here, we examine the role of atomic vibrations in the high-resolution transmission electron microscopy (HRTEM) images using molecular dynamics (MD) simulations. We show that phonon modes localized near the edges of the nanoparticles result in a significant and measurable increase in atomic vibration amplitudes in atoms within a few lattice constants from the edge. This demonstrates the need for modeling vibrations beyond independent atom models when interpreting HRTEM images. Specifically, we follow the frozen phonon approach in which atoms are displaced around their equilibrium sites in the sample and images of a large number of such configurations are generated and averaged to form the actual image. However, rather than drawing random

displacements for each atom from a Gaussian distribution, corresponding to independently vibrating atoms in an Einstein model, we use MD simulations with a machine learning (ML) potential to describe realistic and correlated atomic vibrations.

There are two ways to include vibrations in simulations of HR(S)TEM images. One can replace the electrostatic potential of the sample with a time-averaged potential, for example, by smearing the potentials of the individual atoms with a Debye-Waller factor. Alternatively, one can randomly displace atoms in the sample and generate a large number of images of such configurations before averaging the resulting images. The latter of these, the so-called *frozen phonon method* [5–7], correctly captures the separate timescales of the imaging process. Each electron passes through the sample with a speed close to the speed of light, leading to an interaction time much shorter than the vibrational time of the atoms. On the other hand, the time between two electrons passing through the sample is much larger than the vibrational time; each electron thus images a different snapshot of the atomic configurations [5]. Conversely, applying a Debye-Waller factor corresponds to incorrectly assuming that the interaction time of the electrons is long compared to the vibrational time.

In the usual implementations of the frozen phonon method, each atom is displaced independently; this corresponds to the Einstein model for vibrations in solids, where atoms are uncoupled harmonic oscillators. In reality, however, vibrations are collective waves with vibrational amplitudes depending on the local structure being bulk, surface, interface, or other defective sites.

In this Letter, we explore the vibrational modes in MoS<sub>2</sub> nanoparticles in the context of the observations of Chen *et al.* We present a systematic method for training a machine learning potential based on active learning [8] and equivariant neural networks [9,10] for a supported nanoparticle, in this

\*Contact author: [schiotz@fysik.dtu.dk](mailto:schiotz@fysik.dtu.dk)




FIG. 1. In-plane atomic vibrational amplitudes across  $\text{MoS}_2$  nanoparticles and supporting bilayer graphene. The graphs display the vibrational amplitude of sulfur (blue), molybdenum (red), and carbon (green) atoms as a function of distance from the center of the nanoparticles, for nanoparticles with radii of 2 nm [panel (a)] and 3 nm [panel (b)]. The initial configurations of the nanoparticles are shown as insets. The temperature is 300 K. The vibrational amplitudes of Mo and S atoms increase close to the edge of the nanoparticles. The carbon atom forms two lines, as the lowest layer of the graphite was kept fixed, and the atoms in the layer just above the fixed atoms show lower vibrational amplitude in the transverse direction. We observe no difference in the vibrational amplitudes between the top and bottom sulfur atoms.

case graphite-supported  $\text{MoS}_2$ , including a validity assessment of the potential in subsequent simulations. We then use molecular dynamics (MD) to investigate the vibrational amplitudes of the atoms in the nanoparticle. In agreement with the observations by Chen *et al.*, we find an increased vibrational amplitude not just for the undercoordinated edge atoms, but also for atoms a few lattice constants away from the edge. We use phonon simulations to show that this is due to the existence of edge states where the majority of the vibrational amplitude is localized near the edges.

We optimized equivariant neural network potentials (EN-NPs) describing molybdenum disulfide ( $\text{MoS}_2$ ) supported by graphite or graphene, as described below. The potential was subsequently used for MD simulations of graphite-supported  $\text{MoS}_2$  nanoparticles. Circular nanoparticles with diameters of 30 and 20 Å were created, as shown in Fig. 1. In addition, nanoparticles of the same size but with more sulfur on the edge were created by using a marginally larger cutoff radius for sulfur than for molybdenum.

Molecular dynamics simulations were run in the canonical ensemble at temperatures of 300–500 K. The first 5 ps were thermalization and stabilization, and the following 5 ps were used to collect vibrational amplitudes, defined as the standard deviation of the atomic positions,  $\sigma_n = \sqrt{\langle \mathbf{r}_n^2 \rangle - \langle \mathbf{r}_n \rangle^2}$ , where  $\mathbf{r}_n$  is the position of the  $n$ th atom, and the means are taken over the time series.




FIG. 2. The imaginary values of the peaks in the exit wave simulations for the S-rich sample at different temperatures. The peak values are plotted for all peaks in the exit wave, as a function of the distance from the nanoparticle center.

In Fig. 1, the vibrational amplitudes are plotted versus radial position in the nanoparticle, for the different atomic species. We clearly see an increase in vibrational amplitude for both Mo and S atoms located near the edge of the nanoparticle. The vibrational amplitudes of the outermost sulfur atoms increase slightly when the edge is sulfur rich (see Fig. 9 in the Appendix), but the overall tendency of the vibrational amplitudes increasing in the vicinity of the edge remains the same. It can therefore be expected that other edge terminations and passivations only change this tendency quantitatively; there will most likely still be a tendency for the edge atoms to vibrate more than the bulk atoms, although the actual amplitudes will most likely vary with the termination.

The increase in vibrational amplitudes near the edge of the nanoparticle is consistent, at least qualitatively, with the recently published data of Chen *et al.* [3], where they find an increase in vibrational amplitudes near the edge of  $\text{MoS}_2$  nanoparticles. They recorded a focal series of HRTEM images and reconstructed the wave function of the electron beam exiting a specimen consisting of single-layer  $\text{MoS}_2$  nanoparticles on few-layer graphene support [3]. In order to compare more directly with the results by Chen *et al.*, we simulated the exit wave from our simulation, using the ABTEM simulation software [11] (see Appendix). We modeled atomic vibrations within the frozen phonon approximation, using 100 snapshots of the atomic positions from the last half of our MD simulation. The electron exit wave imaginary part peaks at atomic column positions with a value that represents the projected electrostatic potential averaged over the area of the smeared atoms in the column [3]. Figure 2 shows the exit wave imaginary peak values from the center to the edge of the 3 nm wide  $\text{MoS}_2$  nanoparticle in Fig. 1(b).

While it might be expected that the peak values are lowered at the edge, because enhanced atom vibration and thus atom smearing results from the reduced atom coordination at the edge, it is surprising that the decrease in exit wave amplitude is seen up to a nanometer or more from the edge of the nanoparticle. This finding reflects an enhanced vibrational amplitude both in experiment and in the MD simulations, and could indicate softer phonon modes near the edge of the sample.

To investigate this further, we have simulated phonons in  $\text{MoS}_2$  both in the infinite two-dimensional (2D) sheet and




FIG. 3. The phonon state (DOS) in MoS<sub>2</sub> in bulk and near the edges. Panel (a) shows the bulk phonon DOS. Panel (b) shows the phonon DOS of a strip of MoS<sub>2</sub> 20 unit cells wide, projected onto the central unit cell of the strip. It is all but indistinguishable from the bulk DOS. Panels (c) and (d) show the same DOS projected on the formula unit closest to the Mo and S edges, respectively. In both cases, we see a significant increase of the DOS at low energies. Panel (e) shows the bulk band structure corresponding to the DOS in panel (a). Panels (f) and (g) show the phonon states in the strip, colored by their weight (squared amplitudes, see text) on the sulfur and molybdenum edges, respectively.

in a strip of the material, using the same ML potential. The results are summarized in Figs. 3(a)–3(d), where the phonon density of state (DOS) is compared for an infinite 2D sheet of MoS<sub>2</sub> and for a strip 20 formula units wide and infinitely long. In the latter case, the DOS is projected onto the formula unit (one Mo and two S atoms) at the center of the MoS<sub>2</sub> strip, and onto formula units at the two edges of the strip. While the projected DOS near the center of the strip is almost indistinguishable from the bulk DOS, a strong shift toward lower frequencies is seen at both edges. Phonon states with energies below  $k_B T \approx 0.025$  eV will contribute most strongly to the vibrational amplitudes, and this is where we see the increase in DOS. By *projected DOS* we mean the density of states where each state is weighted by the squares of the vibration amplitudes of the atoms we project onto.

The finite width of the strip removes the dispersion of the phonon modes perpendicular to the strip and replaces it with discrete modes resulting in a larger number of bands [12]; these are plotted in Figs. 3(f) and 3(g) colored by the weight of the modes on the two edges of the strip. The weight is defined as the sum of the squared amplitudes on the three atoms (one Mo, two S) nearest the edge, divided by the sum over all atoms. We clearly see that some bands are strongly localized at one of the edges of the strip, resulting in the increase of DOS observed at the edges. The DOS is affected to a lesser degree up to two unit cells from the edge; see Appendix (Fig. 8).

Selected modes at the  $\Gamma$  and M points are shown in Fig. 4. The top three modes at each edge are examples of modes that are strongly localized at the edge. The lower four modes illustrate that there exist modes that extend a few lattice constants away from the edge. We propose that these modes are responsible for the increase in vibrational amplitude observed both experimentally and in our simulations. Animations of the modes in Fig. 4 can be found in the Supplemental Material [13].

It should be noted that the changes to the phonon DOS are mostly located in the low-frequency part of the spectrum, at or below energies corresponding to room temperature ( $kT \approx 0.025$  eV). If that was not the case, the use of MD would be

more dubious as it does not take into account the quantization of the phonons.

Phonons are often described using the Einstein and the Debye models. The Einstein model assumes that atoms vibrate




FIG. 4. Phonon modes localized at either the sulfur edge (left) or the molybdenum edge (right) of the MoS<sub>2</sub> strip. Panels (a)–(f) show the most localized modes. Panels (g)–(n) show modes that extend a few lattice constants away from the edge. Only modes at the center of the Brillouin zone (marked  $\Gamma$ ) or at the edge (marked M) are shown. The color shows the norm-square weight of the mode at each atom.

independently, but still gives a reasonable description of the optical modes. The Debye model describes the phonons as acoustic waves without dispersion and clearly gives a much better description of the low-frequency part of the DOS. However, neither model includes spatial variation and edge effects as observed in this Letter. In analogy to electronic edge modes in topological insulators, which were incidentally first observed in MoS<sub>2</sub> [14,15], phononic edge states may be topological in origin [16]. Indeed, Li *et al.* found that MoS<sub>2</sub> is a topological phononic material [17].

*Generation of the machine learning potentials.* The simulations are based on ENNPs as implemented in the NEQUIP package [10]. NEQUIP is an equivariant graph neural network, where the energy of each atom is calculated from the vectors to the neighboring atoms and their chemical species. Unlike invariant neural network potential, which requires  $10^4$  to  $10^7$  structures in the training set [18–20], ENNPs can typically be trained on a few hundred training structures [10,21]. The main network contains 3 layers with 32 features of each parity and each tensor rank up to  $l = 1$ . The cutoff distance is 5.0 Å. The radial network uses 8 basis functions, 2 layers and 64 neurons in each layer. The full parameter file can be found in the online data repository [22].

We use a hierarchical training procedure, where each generation of the neural network is used to identify relevant structures to add to the training set of the next generation. The initial generation 0 is created from four nanoparticles of MoS<sub>2</sub> containing from 9 to 24 atoms, two terminated by sulfur atoms and two by molybdenum atoms. From these 100, different structures were generated, by placing them on one to three graphene layers and randomly removing some atoms; see the Appendix for details.

The 100 generated structures are then energy minimized with density functional theory (DFT). We use the GPAW [23,24] electronic structure code and the Perdew-Burke-Ernzerhof (PBE) exchange-correlation potential [25] with D3 dispersion [26], with parameters described in the Appendix. The energy of the structures was then minimized with the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm as implemented in the Atomic Simulation Environment (ASE) [27] until the force is below  $0.2 \text{ eV Å}^{-1}$ , resulting in a final configuration with small but nonzero forces for the training of the ENNP. The first, the middle, and the last configuration of each trajectory are used for training. Of the 300 configurations thus generated, the last 50 are saved for the test set. The remaining 250 are the training set, which is then split randomly into five parts, and five different ENNPs are trained by leaving out one-fifth of the data for each. These potentials are *generation 0*.

Each successive generation of potentials is then made by generating 10–18 new initial structures. Each is used as an initial configuration for a short MD simulation at constant temperature, using the average energy and forces of the five potentials trained in the previous generation [8]. For each MD trajectory, we selected three configurations where the five potentials disagreed the most on the forces, discarding situations where the forces were unreasonably large, as described in the Appendix.

We repeated this procedure until no further improvement in the potential was seen; this occurred after *generation 5*.

With *generation 6*, we saw a slight increase in the test error. In *generation 5*, the root-mean-square error of the force is  $23 \text{ meV Å}^{-1}$  and the error of the energy is  $7.8 \text{ meV}$  per atom, with the worst system in the test set having an error of  $28 \text{ meV}$  per atom. Importantly, this error is comparable to the error of DFT itself, due to the approximations of the exchange-correlation potential.

An important lesson learned while developing this scheme was the importance of removing configurations with unusually large forces (above  $100 \text{ eV Å}^{-1}$  as described in the Appendix). If such forces were not removed, the error of the forces predicted by the resulting potential would be more than an order of magnitude larger, even on a test set that does include such large-force configurations.

A useful modification to the workflow described above could be to replace the DFT energy minimizations in *generation 0* with either energy minimizations or short molecular dynamics simulations using one of the universal machine learning potentials recently published [28,29]. In this case, three selected configurations will still be picked from each starting system, but a DFT calculation will only have to be done for the selected configurations, not for the entire trajectory. It is, of course, important not to train the new machine learning model on energies and forces calculated by the approximate model, but on forces and energies calculated by DFT in a configuration generated by the approximate model.

The final potential used for the simulations was trained on the entire *generation 5* training set. In addition, the five different potentials trained on 4/5th of the *generation 5* dataset were kept for a *posteriori* error estimation.

*Nanoparticle simulations.* Two circular nanoparticles of MoS<sub>2</sub> were created by cutting out atoms from a flat sheet, with cut-out radii of 20 and 30 Å. As the nanoparticles thus created had rather sulfur-deficient edges, an additional two nanoparticles were made with the same cutoff radius used for the sulfur atoms, but the cutoff radius was reduced by 1 Å for the molybdenum atoms.

The resulting four nanoparticles were placed on three-layer slabs of graphite, and MD was performed at temperatures of 300–500 K, using Langevin dynamics. In all cases, the positions of the atoms in the lowest graphene layer were kept fixed. As simulating using an ensemble of five potentials slows the simulation by a factor 5, we used the single potential trained on the entire *generation 5* dataset and only used the ensemble for validation of shorter sequences.

We chose this temperature range as Chen *et al.* measured at room temperature, but they report significantly enhanced atomic vibrations possibly due to sample heating and beam-sample interactions.

During the simulations, we observed a single case of a chemically unreasonable reaction taking place, where two molybdenum atoms left one of the nanoparticles and migrated onto the graphite substrate (see Fig. 5). In this case, we could confirm that the system had passed through a configuration where five ENNPs disagreed on the force on the affected atoms, confirming that we had encountered a configuration outside the confidence regime of the potential.

We cut out regions around the problematic atoms, by including atoms in cylinders with radii of 7, 10, and 15 Å centered on the two problematic Mo atoms. These six




FIG. 5. (a) An example of an error produced by the model potential: A few Mo atoms migrate onto the graphite substrate. (b) The variation of the force between five models just as the weird configuration is beginning to appear, clearly demonstrating that the models are outside their training data and essentially guessing. Cutting out the regions around the problematic atoms and adding the two cut-out regions to the training set eliminates the problem.

configurations were added to the training set, and a new potential was fitted. Rerunning the simulations with the new potential confirmed that no problematic configurations were encountered.

**Conclusion.** We have presented a workflow for training machine learning potentials for specific simulation problems and have trained a potential for carbon-supported molybdenum disulfide nanoparticles. The potential is based on equivariant graph neural networks as implemented in the NEQUIP package. The targeted training workflow makes it possible to get well-performing networks with only slightly more than 400 DFT-calculated configurations in the training set. By using an ensemble of five neural network potentials (“query by committee”), the quality of the potential can be monitored even for system sizes where DFT calculations are prohibitively expensive.

The potential is used to investigate thermal vibrations in supported nanoparticles and show that the vibrational amplitudes are significantly enhanced in a region near the edge of the nanoparticle. This is in agreement with recent observations with recent HRTEM observations and can be explained by localized phonon modes near the edges of the two-dimensional nanoparticles.

This illustrates the importance of using molecular dynamics simulations to correctly model the vibrations when modeling HRTEM images, and we provide a procedure for obtaining the machine learning potentials required to perform such simulations.

**Acknowledgments.** The authors acknowledge financial support from the Independent Research Fund Denmark (DFF-FTP) through Grant No. 9041-00161B. The Center for Visualizing Catalytic Processes (VISION) is funded by the Danish National Research Foundation (DNRF146). We acknowledge support from the Novo Nordisk Foundation Data Science Research Infrastructure 2022 Grant: A high-performance computing infrastructure for data-driven research on sustainable energy materials, Grant No. NNF22OC0078009.

**Data availability.** The data used for reproducing this work including Jupyter Notebooks generating the figures are available online [22].




FIG. 6. Initial  $\text{MoS}_2$  clusters used for generations 0–4. Generations 5 and 6 reused the clusters from generations 0 and 1. Some of them are physically unrealistic, but have value to generate a diverse training set.

#### Appendix: Computational details.

**Generation of structures for machine learning.** For each generation, three or four clusters of  $\text{MoS}_2$  are generated (see Fig. 6). From 10 to 100 supported systems are then generated (100, 10, 18, 18, 18, and 18 for generations 0–6) by randomly picking one of the clusters and placing it on one to three layers of graphene. The distances between graphene layers as well as the distance between the nanoparticle and the graphene are varied slightly around the van der Waals minimum distances. The system is periodic in the  $xy$  plane, so the support forms an infinite sheet. The supercell is chosen large enough that the nanoparticles do not interact with themselves through the periodic boundary conditions. All atoms are then displaced by a random amount from a Gaussian distribution with a spread of 0.1 Å. Finally, in one-third of the samples two atoms are randomly removed from the  $\text{MoS}_2$ , in one-third five random carbon atoms are removed, and in the remaining third all atoms are kept.

**Density functional theory calculations.** We use the GPAW electronic structure code [23,24] and the PBE exchange-correlation potential [25] with D3 dispersion [26] added to properly describe the van der Waals interactions between the layers. The Kohn-Sham wave functions were described




FIG. 7. The phonon band structure of the final machine learning potential compared to the DFT band structure. The acoustic phonon modes are well reproduced; the optical modes have somewhat lower frequencies.

in a plane wave basis with a cutoff of 600 eV. As the supercell is large in the periodic directions, the Brillouin zone is sampled with a single  $k$  point, the  $\Gamma$  point. As the forces on the atoms converge more slowly in the Kohn-Sham scheme than the energies and the electron density, and as the forces are explicitly needed to fit the machine learning potential, an additional convergence criterion was added to the Kohn-Sham self-consistency loop, requiring force convergence within  $0.01 \text{ eV}\text{\AA}^{-1}$ . There is typically at least  $10 \text{ \AA}$  of vacuum in the transverse direction, which keeps the interactions between the  $\text{MoS}_2$  nanoparticle and the back side of the substrate in the next periodic image small.

In generation 0, the energy of the structures is minimized with the BFGS algorithm as implemented in the ASE [27] until the force is below  $0.2 \text{ eV}\text{\AA}^{-1}$ . This very loose convergence criterion ensures that the energy minimization ends with a structure close to but not at a local energy minimum, such that there are nonzero forces available for the training of the neural network potential.

In generations 1–6, no energy minimization is done as the configurations are realistic MD configurations. The forces and energies are just calculated and used for training.

**Molecular dynamics.** MD was used in the same way when generating structure in the active learning procedure and for

the final simulations. The Langevin algorithm as implemented in ASE [27] is used for temperature control during the molecular dynamics; the time step is set to 1 fs and the Langevin friction was  $0.02 \sqrt{\text{eV/u}}/\text{\AA}$  corresponding to a thermalization time constant of 0.51 ps.

**Active learning.** We use a committee-based uncertainty sampling for improving the machine learning potential. MD was run for 400 time steps using the average of the five potentials trained in the previous generation [8]. The temperature is chosen randomly to be either 300 or 400 K (for generations 1–3, only 400 K was used). For each step in the MD simulation, we calculated the standard deviation of the forces predicted by the ensemble of models. This produced a standard deviation for each atom; the average of these was then used to detect configurations where the five models disagreed.

As the early generations of the potential were not reliable, occasionally absurdly large forces would be predicted, often due to the appearance of physically unreasonable structures. If forces above  $100 \text{ eV}\text{\AA}^{-1}$  appeared in an MD simulation, the configurations from that point on would be discarded and not used in generating training data.

After this filtering, we selected the three configurations of each MD trajectory with the largest standard deviation between the five predicted, but with the additional constraint that no two configurations less than six time steps apart were selected. These new configurations were added to the training set, and a new generation of five models was trained.

This procedure for systematically improving the ML potential will detect and fix errors caused by a lack of similar structures in the training set. However, it will not necessarily detect whether there are contributions to the potential energy that cannot be properly described by the machine learning potential—These could, for example, be long-ranged electrostatic forces [30]. Figure 7 shows how the phonon band structure of the ML potential compares to DFT. The acoustic branches are well reproduced, but the optical modes are somewhat too soft, perhaps due to incomplete description of long-ranged electrostatic forces. However, all results presented here depend on changes induced by the broken symmetry at the edges rather than on detailed frequencies, and the major changes occur in the well-reproduced part of the spectrum (below 25 meV).

**Exit wave simulations.** To compare with the experimental exit waves [3], we simulated thermally averaged exit waves




FIG. 8. The phonon state (DOS) in  $\text{MoS}_2$  is perturbed by the edge state. The perturbation is clearest in the unit cells at the edge, but can be seen up to two unit cells away from it.

from the MD simulations, using the multislice algorithm and the frozen phonon method.

The frozen phonon method leverages the separation of timescales in electron microscopy. The timescale of the camera (microseconds to milliseconds) is very long compared to the atomic vibrations (subpicoseconds), which again are slow compared to the time it takes for an electron to traverse the sample (tens of attoseconds). The final image should thus be an average of individual images calculated for independent snapshots of the atomic vibrations. This is in contrast with the method of applying a Debye-Waller factor, which corresponds to calculating a single image of an averaged atomic potential.

It should be noted that in the frozen phonon approximation, there is not a single exit wave of the sample, but an exit wave of each atomic configuration. Nevertheless, we calculate an *averaged* exit wave, for use in comparing to the exit wave reconstructed from the experimental image sequence.

The exit wave is calculated using the multislice algorithm as implemented in the ABTEM software [11]. As we compare to reconstructed exit waves where the support has been removed by Fourier filtering the images prior to the exit wave reconstruction, we need to simulate an exit wave without the support. We thus remove all carbon atoms from the simulations. We select 100 distributed configurations from the second half of the MD simulation; for each of these, the exit wave is calculated with ABTEM and subsequently averaged. The multislice simulation is performed with a lateral grid size of 0.05 Å and a slice thickness of 0.5 Å. The atomic potentials are parametrized according to Lobato and Van Dyck [31].

Once the averaged exit wave is calculated, we find local peaks in the imaginary part of it. The peak values and their distance from the center of the nanoparticles are then extracted and used to produce Fig. 2.

**Phonon band structure simulations.** The phonon band structures and densities of states were calculated using the Phonons module of the ASE, using the same ENNP used for the simulations. The phonon spectrum is calculated from a single freestanding layer of MoS<sub>2</sub>; this results in a transverse acoustic mode with quadratic dispersion, corresponding to buckling of the 2D sheet of material. This may distort the lowest parts of the density of states. Unfortunately, calculating the phonon band structure of a supported material is not possible as the unit cell of the material and support do not match.

To investigate the existence of modes localized at the edges, we performed similar phonon calculations with a strip of MoS<sub>2</sub> 20 unit cells wide. This system thus only has




FIG. 9. The vibrational amplitude of sulfur atoms for nanoparticles with (a) Mo-rich and (b) S-rich edges, as a function of the distance from the edge. The red line is a running average. The gray lines show the vibrational level inside the nanoparticle, a level corresponding to 50% increase in amplitude, and at what distance from the edge that level is reached. Panel (c) summarizes these data for Mo- and S-terminated nanoparticles at different temperatures.

periodic boundary conditions along one direction along which the phonons may propagate. In the transverse direction, standing waves will form, splitting each bulk phonon band into 20 new bands. The resulting band structure is therefore quite “messy,” as seen in Fig. 3.

The DOS projected on the formula units at the edge of the strip is significantly altered, as already shown in Fig. 3. As the edge phonon modes extend a few unit cells into the unit cell, so does the perturbation of the projected DOS. This is seen in Fig. 8, where a perturbation can be seen up to two unit cells from the edge.

**Effect of edge termination.** The simulations were made with nanoparticles where the edge was Mo rich or S rich. The difference is minimal; the outermost sulfur atoms are more loosely bound in the S-rich particles, and thus have a higher vibrational amplitude. That is shown in Figs. 9(a) and 9(b). Figure 9(c) shows how far into the nanoparticle the increased vibrational amplitude reaches; it is seen that for the Mo-terminated edge this distance increases with temperature, but it does not for the S-terminated edge.

- [1] C. Kisielowski *et al.*, Detection of single atoms and buried defects in three dimensions by aberration-corrected electron microscope with 0.5-Å information limit, *Microsc. Microanal.* **14**, 469 (2008).
- [2] R. Erni, M. D. Rossell, C. Kisielowski, and U. Dahmen, Atomic-resolution imaging with a sub-50-pm electron probe, *Phys. Rev. Lett.* **102**, 096101 (2009).
- [3] F.-R. Chen, D. Van Dyck, C. Kisielowski, L. P. Hansen, B. Barton, and S. Helveg, Probing atom dynamics of

excited Co-Mo-S nanocrystals in 3D, *Nat. Commun.* **12**, 5007 (2021).

- [4] Z. Chen, Y. Jiang, Y.-T. Shao, M. E. Holtz, M. Odstrčil, M. Guizar-Sicairos, I. Hanke, S. Ganschow, D. G. Schlom, and D. A. Muller, Electron ptychography achieves atomic-resolution limits set by lattice vibrations, *Science* **372**, 826 (2021).
- [5] R. F. Loane, P. Xu, and J. Silcox, Thermal vibrations in convergent-beam electron diffraction, *Acta Crystallogr. A* **47**, 267 (1991).

- [6] S. Hillyard and J. Silcox, Detector geometry, thermal diffuse scattering and strain effects in ADF STEM imaging, *Ultramicroscopy* **58**, 6 (1995).
- [7] D. Van Dyck, Persistent misconceptions about incoherence in electron microscopy, *Ultramicroscopy* **111**, 894 (2011).
- [8] J. S. Smith, B. Nebgen, N. Lubbers, O. Isayev, and A. E. Roitberg, Less is more: Sampling chemical space with active learning, *J. Chem. Phys.* **148**, 241733 (2018).
- [9] N. Thomas, T. Smidt, S. Kearnes, L. Yang, L. Li, K. Kohlhoff, and P. Riley, Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds, [arXiv:1802.08219](https://arxiv.org/abs/1802.08219).
- [10] S. Batzner, A. Musaelian, L. Sun, M. Geiger, J. P. Mailoa, M. Kornbluth, N. Molinari, T. E. Smidt, and B. Kozinsky, E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials, *Nat. Commun.* **13**, 2453 (2022).
- [11] J. Madsen and T. Susi, The ABTEM code: Transmission electron microscopy from first principles, *Open Res. Eur.* **1**, 24 (2021).
- [12] M. Igami, M. Fujita, and S. Mizuno, Phonon dispersion of nanographite ribbons, *Appl. Surf. Sci.* **130**, 870 (1998).
- [13] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/f4y5-bps4> for animated renderings of the phonon modes in Fig. 4 and the files are named by the edge, the  $k$  point, and the energy in meV.
- [14] S. Helveg, J. V. Lauritsen, E. Lægsgaard, I. Stensgaard, J. K. Nørskov, B. S. Clausen, H. Topsøe, and F. Besenbacher, Atomic-scale structure of single-layer MoS<sub>2</sub> nanoclusters, *Phys. Rev. Lett.* **84**, 951 (2000).
- [15] M. V. Bollinger, J. V. Lauritsen, K. W. Jacobsen, J. K. Nørskov, S. Helveg, and F. Besenbacher, One-dimensional metallic edge states in MoS<sub>2</sub>, *Phys. Rev. Lett.* **87**, 196803 (2001).
- [16] W. Liu, Z. K. Ding, N. Luo, J. Zeng, L. M. Tang, and K. Q. Chen, Phononic hybrid-order topology in semihydrogenated graphene, *Phys. Rev. B* **109**, 115422 (2024).
- [17] J. Li, J. Liu, S. A. Baronett, M. Liu, L. Wang, R. Li, Y. Chen, D. Li, Q. Zhu, and X. Q. Chen, Computation and data driven discovery of topological phononic materials, *Nat. Commun.* **12**, 1204 (2021).
- [18] K. T. Schütt, P.-J. Kindermans, H. E. Sauceda, S. Chmiela, A. Tkatchenko, and K.-R. Müller, SchNet: A continuous-filter convolutional neural network for modeling quantum interactions, *Adv. Neural Inf. Process. Syst.* **30**, 991 (2017).
- [19] J. S. Smith, O. Isayev, and A. E. Roitberg, ANI-1: An extensible neural network potential with DFT accuracy at force field computational cost, *Chem. Sci.* **8**, 3192 (2017).
- [20] A. E. Mikkelsen, J. Schiøtz, T. Vegge, and K. W. Jacobsen, Is the water/Pt(111) interface ordered at room temperature? *J. Chem. Phys.* **155**, 224701 (2021).
- [21] C. Nuñez Valencia, W. B. Lomholdt, M. H. Leth Larsen, T. W. Hansen, and J. Schiøtz, Beam induced heating in electron microscopy modeled with machine learning interatomic potentials, *Nanoscale* **16**, 5750 (2024).
- [22] P. Giese, M. S. Nissen, S. Helveg and J. Schiøtz, Data for “Visualizing phonon edge states on molybdenum disulphide”, Technical University of Denmark, doi:[10.11583/DTU.30084814](https://doi.org/10.11583/DTU.30084814).
- [23] J. Enkovaara *et al.*, Electronic structure calculations with GPAW: A real-space implementation of the projector augmented-wave method, *J. Phys.: Condens. Matter* **22**, 253202 (2010).
- [24] J. J. Mortensen *et al.*, GPAW: An open Python package for electronic structure calculations, *J. Chem. Phys.* **160**, 092503 (2024).
- [25] J. P. Perdew, K. Burke, and M. Ernzerhof, Generalized gradient approximation made simple, *Phys. Rev. Lett.* **77**, 3865 (1996).
- [26] S. Grimme, J. Antony, S. Ehrlich, and H. Krieg, A consistent and accurate *ab initio* parametrization of density functional dispersion correction (DFT-D) for the 94 elements H-Pu, *J. Chem. Phys.* **132**, 154104 (2010).
- [27] A. H. Larsen *et al.*, The atomic simulation environment—A Python library for working with atoms, *J. Phys.: Condens. Matter* **29**, 273002 (2017).
- [28] B. Deng, P. Zhong, K. J. Jun, J. Riebesell, K. Han, C. J. Bartel, and G. Ceder, CHGNet as a pretrained universal neural network potential for charge-informed atomistic modelling, *Nat. Mach. Intell.* **5**, 1031 (2023).
- [29] I. Batatia *et al.*, A foundation model for atomistic materials chemistry, *J. Chem. Phys.* **163**, 184110 (2025).
- [30] D. Kim, X. Wang, S. Vargas, P. Zhong, D. S. King, T. J. Inizan, and B. Cheng, A universal augmentation framework for long-range electrostatics in machine learning interatomic potentials, *J. Chem. Theory Comput.* **21**, 12709 (2025).
- [31] I. Lobato and D. Van Dyck, An accurate parameterization for scattering factors, electron densities and electrostatic potentials for neutral atoms that obey all physical constraints, *Acta Crystallogr. A* **70**, 636 (2014).