

Full-cycle device-scale simulations of memory materials with a tailored atomic-cluster-expansion potential

Received: 12 February 2025

Accepted: 20 August 2025

Published online: 30 September 2025

 Check for updatesYuxing Zhou¹, Daniel F. Thomas du Toit¹, Stephen R. Elliott², Wei Zhang³✉ & Volker L. Deringer¹✉

Computer simulations have long been key to understanding and designing phase-change materials (PCMs) for memory technologies. Machine learning is now increasingly being used to accelerate the modelling of PCMs, and yet it remains challenging to simultaneously reach the length and time scales required to simulate the operation of real-world PCM devices. Here, we show how ultra-fast machine-learned interatomic potentials, based on the atomic cluster expansion (ACE) framework, enable simulations of PCMs reflecting applications in devices with excellent scalability on high-performance computing platforms. We report full-cycle simulations—including the time-consuming crystallisation process (from digital “zeroes” to “ones”)—thus representing the entire programming cycle for cross-point memory devices. We also showcase a simulation of full-cycle operations, relevant to neuro-morphic computing, in a mushroom-type device geometry.

Phase-change materials (PCMs) from the Ge–Sb–Te system have been widely used in emerging electronic devices, including non-volatile memory and neuromorphic in-memory computing technologies^{1–6}. Driven by Joule heating resulting from the application of electric pulses, the SET (crystallisation) and RESET (amorphisation) operations are associated with fast and reversible transitions between the amorphous (low-conductance) and crystalline (high-conductance) states of PCMs. A large property contrast between these states encodes “zeroes” and “ones” in the atomic structure, respectively, for binary memory⁷. Furthermore, finely tuning the conductance of PCM cells between all-amorphous and all-crystalline states enables multi-level programming for neuromorphic in-memory computing⁸.

The switching processes in PCMs can be completed within nanoseconds—that is, within time scales accessible for molecular-dynamics (MD) computer simulations—and this has long made PCMs a prime application area in the field of materials modelling. Density-functional theory (DFT)-driven *ab initio* molecular-dynamics (AIMD) simulations have played a key role in understanding structural features^{9–11}, property contrast^{12–15}, and crystallisation kinetics^{16–18} of

PCMs based on representative small-scale models (typically containing on the order of 1000 atoms or fewer)¹⁸. Building on the long-standing successes of DFT and AIMD, machine-learning (ML)-based interatomic potentials have recently emerged which accelerate first-principles atomistic modelling by many orders of magnitude^{19–21}, and which can therefore provide new insights at much-extended time and length scales.

More than a decade ago already, Sosso et al. reported the first ML potential for modelling PCMs, at that time for the binary compound GeTe²², based on the Behler–Parrinello neural-network framework²³. Since then, ML-driven MD simulations have become gradually more established: for example, they revealed details of the temperature-dependent crystallisation in GeTe²⁴ and Ge₂Sb₂Te₅ alloys²⁵. In time, more ML potentials have begun to be developed for different PCMs^{25–29}. We recently reported a chemically transferable and defect-tolerant ML potential for Ge–Sb–Te (GST) materials along the GeTe–Sb₂Te₃ tie-line, fitted using the Gaussian approximation potential (GAP) framework³⁰ and based on a comprehensive structural dataset and iterative training³¹. A neuro-evolution potential was

¹Inorganic Chemistry Laboratory, Department of Chemistry, University of Oxford, Oxford, UK. ²Physical and Theoretical Chemistry Laboratory, Department of Chemistry, University of Oxford, Oxford, UK. ³Center for Alloy Innovation and Design (CAID), State Key Laboratory for Mechanical Behavior of Materials, Xi'an Jiaotong University, Xi'an, China. ✉ e-mail: wzhang0@mail.xjtu.edu.cn; volker.deringer@chem.ox.ac.uk

recently developed for large-scale crystallisation simulations of Sb_2Te , SbTe , and Sb_2Te_3 , revealing distinct behaviours driven by nucleation and growth³². More generally, graph-based ML methods constitute the current state-of-the-art architectures in terms of accuracy and chemical transferability^{33,34}, and they have begun to attract attention in the PCM community for general-purpose simulation tasks³⁵.

In 2025, computational modelling is now often able to describe “the real thing” thanks to ML-driven potentials³⁶, and yet these models still face a significant obstacle when it comes to describing PCMs in a fully realistic way—e.g., because of the length scales and structural complexity associated with applications in this domain. Complex simulation protocols are therefore required to model PCM devices, such as non-isothermal heating, which we have demonstrated for both cross-point and mushroom-type cells³¹. Using the GST-GAP-22 potential at the time, we simulated a 50-picosecond RESET operation (“1 → 0”), showing non-isothermal melting and rapid cooling in a 532,980-atom model, representing a cell volume of $20 \times 20 \times 40 \text{ nm}^3$ in cross-point memory devices³¹. However, the subsequent SET (“0 → 1”) typically requires tens of nanoseconds, rather than tens of picoseconds, to complete in devices. Performing a crystallisation run over 10 ns for the same structural model, with GST-GAP-22, would have

consumed more than 150 million CPU core hours by our estimate. This type of excessive cost (in terms of time, financial cost, and carbon emissions) would clearly make the use of GST-GAP-22 unfeasible for full-cycle modelling of GST devices.

Herein, we show how one can simultaneously reach both the length and time scales in simulations of switching operations in real-world GST devices, leveraging the atomic-cluster-expansion (ACE) ML framework³⁷. The substantial speed-up by moving from the GAP to the ACE framework³⁸ enables atomistic simulations reflecting device applications on widely available CPU-based computing systems. We have thus outlined an “off-the-shelf”-usable ML approach for the community to study the switching mechanisms of GST-based devices. Beyond PCMs, our work explores the current frontiers of ultra-large-scale all-atom simulations for materials science and engineering.

Results

Fast and CPU-efficient simulations with an optimised ACE potential

We used the ACE framework to develop a computationally efficient ML model for GST. In ACE, the local environment of a given atom is encoded using a many-body expansion (Fig. 1a). The atomic

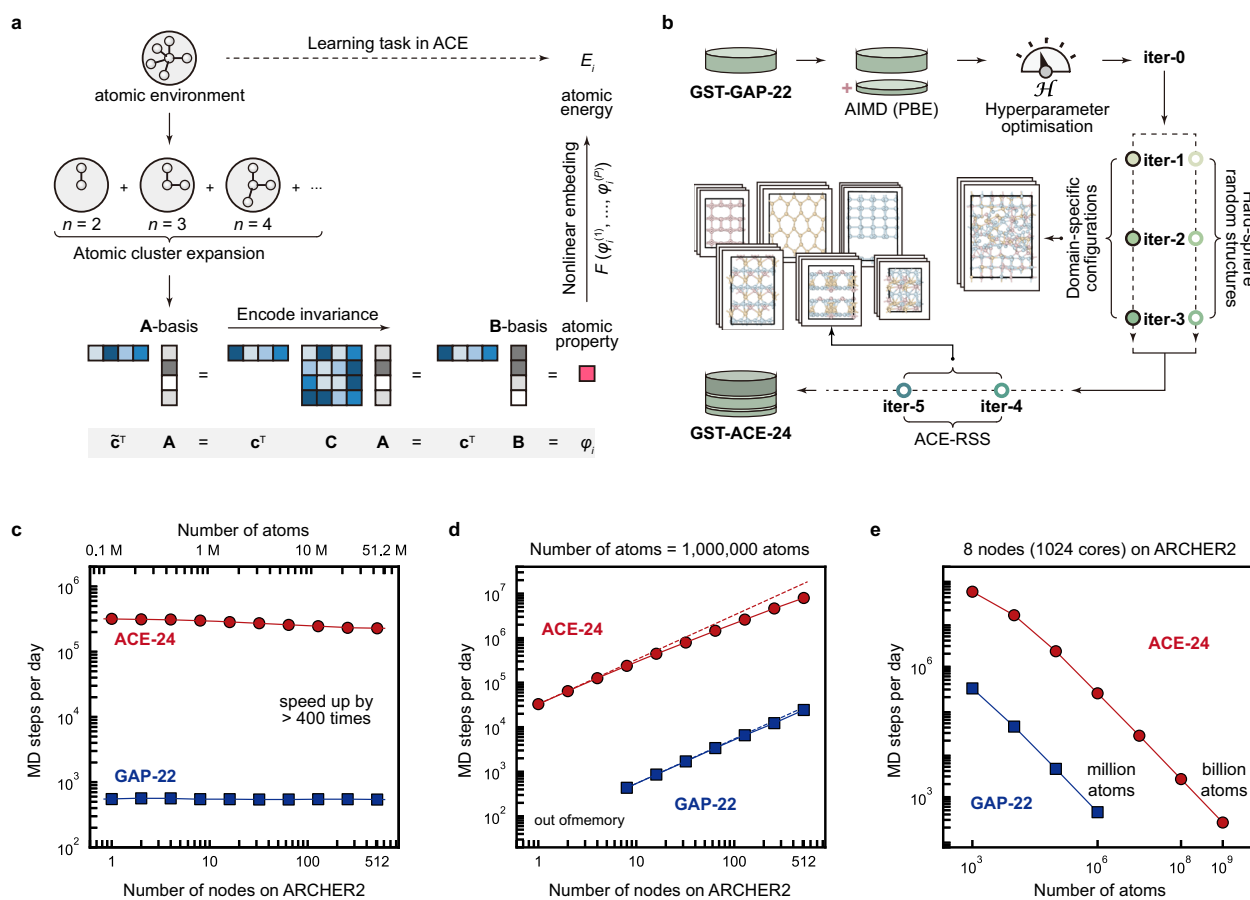


Fig. 1 | An optimised atomic-cluster-expansion (ACE) machine-learning (ML) potential for Ge-Sb-Te (GST) phase-change materials. **a** Schematic overview of the learning task in ACE, showing the matrix equations that define the fitting of ACE models. The atomic property, ϕ , of a given atomic environment is expressed using a linear combination (via the coefficient vector $\tilde{\mathbf{c}}$) of “A-basis” functions, \mathbf{A} , or using a linear combination (via the coefficient vector \mathbf{c}) of invariant “B-basis” functions, \mathbf{B} , by coupling via the generalised Clebsch–Gordan coefficients, \mathbf{C} , i.e., $\tilde{\mathbf{c}} = \mathbf{c} \cdot \mathbf{C}$. The energy of the atom, E , is predicted as a function of various atomic properties, using a linear or nonlinear embedding, i.e., $E = F(\phi^{(1)}, \dots, \phi^{(P)})$. **b** The fitting protocol used for a new ACE potential of GST. Multistep iterations were carried out to expand the reference dataset. Insets show typical configurations for domain-

specific iterations and ACE-driven random-structure searching (ACE-RSS) structures. Ge, Sb, and Te atoms are rendered as light red, light yellow, and light blue, respectively. Details are given in Supplementary Note 1. Different tests of computational efficiency on the ARCHER2 high-performance computing system were carried out for the GST-ACE-24 (this work) and GST-GAP-22³¹ potentials, including: **c** weak scaling (at 100,000 atoms per node); **d** strong scaling up to 512 nodes—the maximum accessible for a single job on ARCHER2—for a 1-million-atom structural model; and **e** molecular-dynamics (MD) steps per day as a function of system size running on 8 computing nodes (1024 CPU cores). The GAP-MD tests in **(d)** failed on configurations with 4 or fewer nodes due to insufficient memory. The dashed lines in **(d)** indicate ideal strong scaling behaviour.

environment is expressed in terms of radial functions and spherical harmonics, translated into a linear combination of so-called “A-basis” functions, and subsequently into invariant “B-basis” functions by coupling via the generalised Clebsch–Gordan coefficients. A linear combination of B-basis functions is called the “property” of a given atomic environment in the context of ACE. The energy of the atom is predicted as a function of atomic properties, using a linear (depending on just one single property) or nonlinear embedding. The complexity of ACE models is therefore controlled by the numbers of basis and embedding functions; more details of the framework can be found in refs. 37,39–41.

Inference in ACE essentially requires summation operations (Fig. 1a). Hence, ACE models are computationally highly efficient: they can be more than $100 \times$ faster than GAP on CPU cores while achieving the same level of numerical accuracy^{39,42}. In contrast, we note that GAP has a high data efficiency (and “learning capability”), enabling efficient collection of initial training datasets, especially at the early stages of fitting (see ref. 43 for an overview of the GAP framework). We recently developed an initial ML potential for elemental tellurium using GAP, and then re-fitted the reference data using ACE, to study crystallisation and melting in Te-based selector devices⁴⁴. Here, we address the structurally and chemically more complex GST system, starting from the existing GST-GAP-22 dataset—which incorporates relevant domain knowledge³¹—and also making use of the ACE framework. However, simply re-fitting an ACE potential using the GST-GAP-22 training data “out of the box” produced unphysical structural motifs in our tests (e.g., atomic clustering and phase segregation), and lost atoms in MD simulations. We believe that this is due to two reasons: first, the hyperparameters used in ACE are complex and can be difficult to tune manually (cf. Supplementary Table 2); second, an ACE model may require more training structures than GAP⁴⁵. The usefulness of ACE potentials in the field of GST materials has been demonstrated recently³⁸: the authors focused on the use of a previously proposed indirect-learning approach for ML potentials⁴⁶ and built upon the GST-GAP-22 dataset and model³¹, reporting efficient ACE potentials which were applied to the $\text{Ge}_2\text{Sb}_2\text{Te}_5$ compound³⁸.

In the present work, we focus on an alternative route, both optimising the hyperparameters of the model and extending the DFT training dataset (Fig. 1b). The GST-GAP-22 dataset was first re-labelled using DFT with the Perdew–Burke–Ernzerhof (PBE) exchange–correlation functional⁴⁷ that is widely used in simulations of PCMs. We then added further AIMD configurations of disordered GST (taken from ref. 31) and fitted initial ACE models to the combined data, using the XPOT software to optimise hyperparameters (Methods)^{48–50}. Starting from this well-parameterised ACE model (denoted “iter-0”), we carried out three domain-specific iterations (iter-1 to iter-3) to include melt-quenched disordered structures and intermediate configurations during phase-transition processes (Fig. 1b). These stepwise iterations, acting as a “self-correction” process, provide feedback that enables the potential to correct errors and inaccuracies emerging in its own simulations. We also added small-scale hard-sphere random structures (6–40 atoms) with small atomic distances, generated using the `buildcell` code of *ab initio* random structure searching (AIRSS)^{51,52}, for iter-1 to iter-3, to make the ACE models more robust⁴². In addition, we carried out ACE-driven random structure searching (ACE-RSS), akin to previously described GAP-driven RSS^{53–55}, in two further iterations (iter-4 and iter-5). We refer to our final ACE potential model as “GST-ACE-24”. With its training dataset covering multiple GST compositions, we found that our GST-ACE-24 model is chemically transferable along the entire GeTe – Sb_2Te_3 tie-line: it can accurately capture different structural properties of various amorphous GST compounds, as validated against AIMD (Methods).

To evaluate the computational efficiency of GST-ACE-24, we performed weak and strong scaling tests on ARCHER2, a high-performance computing system in the UK (Methods and

Supplementary Note 2). We found that, compared to GST-GAP-22, GST-ACE-24 offers more than $400 \times$ higher efficiency on this large-scale CPU architecture (Fig. 1c). Both ACE and GAP showed reasonable scaling behaviour up to 512 nodes (65,536 CPU cores) in strong scaling tests for a structural model of 1 million atoms (Fig. 1d). An efficiency drop-off from the ideal scaling behaviour occurred for the ACE model when handling “small” structures (e.g., 100,000 atoms) on many computing nodes (Supplementary Fig. 1), because ACE is so fast that the inter-processor communication outweighs the computational cost of predicting energies and forces. For example, $\approx 30\%$ of CPU time was used in inter-processor communication when simulating a 100,000-atom structural model on more than 128 nodes. In addition, we found the system-size limit for a total memory of 512 GB to be $\approx 450,000$ atoms with GAP, whereas for the same hardware the limit was ≈ 650 million atoms with ACE. Hence, ACE is memory-efficient and enables billion-atom MD simulations (Fig. 1e) with only modest computational resources (e.g., 8 nodes on ARCHER2).

Moreover, ACE can also be used on GPU hardware. We tested device-scale ACE-MD simulations on up to four NVIDIA A100 GPU cards (Methods), and found that the compute-time requirement for ACE-MD on one such GPU card is of the same order of magnitude as that on one 128-core CPU node on ARCHER2 (Supplementary Fig. 2). However, a direct comparison between CPU and GPU is not entirely meaningful due to differences in hardware, parallel computing capabilities, and software-level optimisation for computational tasks. For example, the recursive evaluator developed for ACE—enabled via the `pacemaker` package in LAMMPS and designed to further increase computational efficiency—is currently only compatible with CPU. We also found that ACE’s speed compares favourably to state-of-the-art graph-neural-network architectures: our ACE model is about $6 \times$ faster than an equivariant neural-network potential that we directly re-fitted for comparison, using the same training data as for GST-ACE-24 and the MACE architecture^{34,56}, when testing on an NVIDIA A100 GPU card (Methods). We found the system-size limit for a total memory of 80 GB to be ≈ 92 million atoms with ACE on GPU, smaller than the limit on CPU (≈ 650 million atoms). Hence, while ACE-MD can be run on GPU, its excellent scaling behaviour across multiple CPU nodes and the potential large memory capacity of CPU nodes make it particularly suitable for device-scale MD simulations on existing CPU hardware.

For practical MD simulations, we emphasise the importance of the robustness of ACE models: ACE-MD simulations will fail when atoms are lost due to inaccurately predicted energies and forces. This usually stems from insufficient training data for complex atomic environments with small atomic distances. We designed a protocol to quantify the robustness of ACE models via high-temperature annealing: starting with a hard-sphere random structure of 1000 atoms, the model was annealed at 3000 K for 500 ps. (We note that high- T annealing is part of the melt-quench process to generate amorphous GST, allowing the simulation to visit high-energy configurations; we also note that high- T MD has been used previously for stability tests⁵⁷.) We tested the robustness of ACE models on 7 different GST compositions, from GeTe to Sb_2Te_3 , and performed 10 independent high- T annealing runs for each composition. Despite gradually adding hard-sphere random structures to the training dataset, no successful runs were observed from iter-1 to iter-3: very close interatomic contacts were found in the MD simulations (Fig. 2a), which led to large forces and then lost atoms. However, the inclusion of ACE-RSS structures in the training is concomitant with some successful runs in iter-4 and consistently successful runs in iter-5, producing reasonable high- T liquid GST structures (Fig. 2a).

Ablation studies for the ACE model

In ML research, “ablation” studies mean gradually removing aspects of a complex model and testing the effect of that on the performance. Here, we report systematic ablation studies for ACE models, with an

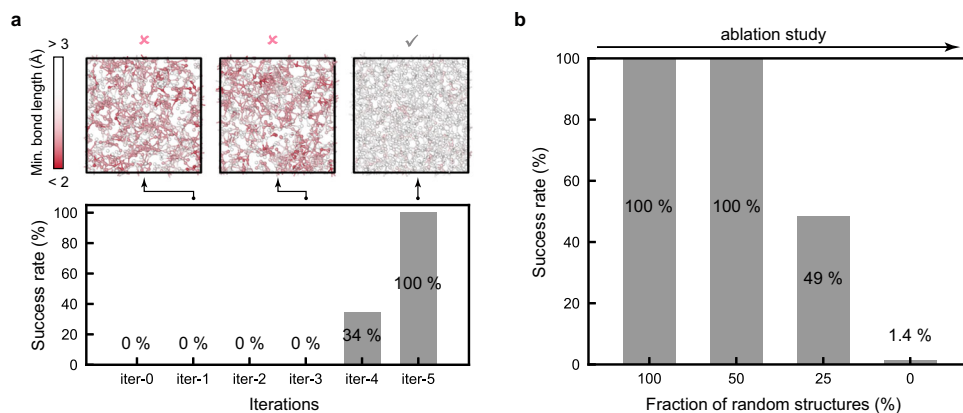


Fig. 2 | Ablation studies for GST-ACE-24. a Fraction of successful runs in high-temperature annealing tests for each iteration of the potential. Three structural models from different tests are shown, with atoms colour-coded by minimum bond length: red colour indicates atoms with very short distances to their nearest neighbour, which might lead to lost atoms in molecular-dynamics simulations. **b** An

ablation study for the number of random structures present in the training data, evaluated based on the fraction of successful high- T annealing runs. The success rate was obtained across 7 different Ge-Sb-Te compositions, from GeTe to Sb₂Te₃, with each composition including 10 independent high- T annealing runs.

aim to more systematically understand the roles of newly added configurations and optimised hyperparameters in ACE models. We first carried out an ablation study for the removal of random structures (including the hard-sphere random and ACE-RSS structures) based on high-temperature annealing tests (Fig. 2b). We observed consistently successful runs when more than half of the random structures were retained. However, removing 75% of them resulted in a marked decrease in stability, and a model where no random structures were present showed almost no successful runs.

We note that removing random structures slightly improved the numerical accuracy in describing domain-relevant structures (e.g., in melt-quench and phase-transition processes), in terms of energies (up to 3 meV atom⁻¹) and forces (up to 13 meV Å⁻¹) in our ablation tests (Table 1); however, for practical purposes, we do not expect that this small advantage will typically outweigh the risk of losing atoms described above. Hence, in the context of ongoing research on how datasets for ML potentials are best developed^{21,58–61}, we highlight the key role of such small-scale RSS structures: not only as a starting point for fitting potentials^{54,61–64}, but also as an effective *post-hoc* correction approach that can add substantial MD stability.

We also carried out ablation studies for the complexity of ACE models. We changed the number of atomic properties, P , which controls how the atomic energy is constructed from the local atomic properties in a linear ($P=1$) or non-linear way ($P\geq 2$). We fitted a linear ACE model and a simpler non-linear one with $P=2$, and compared both resulting models against GST-ACE-24 ($P=3$). For the linear model, we found a force RMSE approximately 30 meV Å⁻¹ higher than for GST-ACE-24. We also compared ACE models with gradually reduced basis functions (to 1500, 750, and 300, respectively). Although decreasing the number of basis functions increases the computational efficiency (Table 1), fewer functions lead to larger numerical errors. Hence, we argue that our GST-ACE-24 model offers a favourable combination of robustness and accuracy for practical MD simulations.

Full-cycle operations for cross-point GST memory devices

With the help of ACE, we are now able to simulate full-cycle operations in cross-point GST devices (Fig. 3a). We first reproduced a non-isothermal melt-quench (RESET) simulation that had previously been demonstrated for cross-point memory, using the GST-GAP-22 potential at the time³¹. As in ref. 31, we used a structural model of Ge₁Sb₂Te₄ of 20 × 20 × 40 nm³ (532,980 atoms), which includes a fixed (here, amorphous) slab to prevent unwanted atomic migration across the

periodic cell boundary (Supplementary Note 3)—resembling a thermal barrier in contact with GST in a real device (Fig. 3a).

We used the NVE ensemble (i.e., constant number of particles, volume, and energy) to simulate the RESET process. As shown in Fig. 3b, a 10-ps heating pulse (0.064 pJ) imposed on the model was first simulated by spatially inhomogeneously increasing the kinetic energy of atoms linearly in the direction of the z -axis, corresponding to a large temperature gradient from the bottom (>1000 K) to the top (\approx 300 K) of the 40-nm-long cell. To model the cooling process after removing the heating pulse, the added energy was then gradually removed from each atom over another 40 ps until reaching room temperature. We note that such picosecond-scale data-erasure time is experimentally accessible, evidenced by femtosecond laser experiments⁶⁵ and picosecond-scale optical pulses used in an all-optical calculator⁶⁶. The atoms in Fig. 3 are colour-coded based on the smooth overlap of atomic positions (SOAP) kernel similarity⁶⁷, which was previously used to quantify per-atom crystallinity for GST¹⁸. Based on the ACE-MD trajectories, we found that almost all of the structural model turned into amorphous Ge₁Sb₂Te₄ after heating and cooling (Supplementary Fig. 3). Technical details of these non-isothermal heating and cooling simulations using NVE are given in Supplementary Note 3 and in our previous work³¹. We note that this RESET simulation using GST-ACE-24 and its evolution of temperature gradients is consistent with previous results using GST-GAP-22 (Supplementary Fig. 3)³¹, providing further validation of the present approach.

We next simulated the SET process of the cross-point structural model. Unlike the short, intense RESET pulse (10 ps, 0.064 pJ) that generates a pronounced temperature gradient across the cell, the SET heating pulse has a much longer duration (e.g., tens of nanoseconds) and smaller amplitude, resulting in a lower temperature gradient and smaller fluctuations. Here, we simulated the crystallisation of the device-scale model using the NVT ensemble (i.e., constant number of particles, volume, and temperature). The crystallisation of undoped GST is known to be driven by homogeneous nucleation⁶⁸, in which critical nuclei quickly form during a stochastic incubation process¹⁶. The latter is the bottleneck for crystallisation, which can be bypassed either by applying a low-voltage seeding pre-pulse^{69,70} or by doping with a suitable transition-metal element^{71–75}. We note that in previous nucleation simulations of GST using AIMD, enhanced sampling methods, e.g., meta-dynamics^{76,77} or pre-embedded crystalline seeds^{72,78}, were employed to accelerate, or circumvent, the formation of critical nuclei in small-scale structural models.

Table 1 | Ablation studies exploring varying numbers of quantities in ACE models, including: (i) random structures in the training; (ii) atomic properties (see also ref. 45); and (iii) basis functions, based on computed root-mean-square error (RMSE) and relative molecular-dynamics (MD) speed with reference to GST-ACE-24

	Quantities	Energy RMSE (meV atom ⁻¹)	Force RMSE (meV Å ⁻¹)	MD speed
GST-ACE-24 (reference)				
Number of random structures	50 %	17	129	1.0
	25 %	16	126	1.0
	0 %	15	122	1.0
Number of atomic properties, $\varphi^{(P)}$	$P=2$	21	142	1.0
	$P=1$	23	165	1.0
Number of basis functions	1500	22	144	1.5
	750	23	153	2.1
	300	23	168	3.1

We show computed RMSE values for energy and force predictions on the testing dataset that contains relevant configurations from melt-quench and phase-transition processes. The GST-ACE-24 model developed in this work serves as the reference model in the ablation study, which has 906 random structures (319 hard-sphere random and 587 ACE-driven random-structure searching structures; Supplementary Table 1) in the training dataset, $P=3$ atomic properties, and 3000 basis functions.

GST-ACE-24 is able to describe nucleation in GST without such additional constraints. We annealed the device-scale structural model of amorphous GST at 600 K for 20 ns, which corresponds to typical electrical pulse durations in GST-based devices^{72,79,80}; however, fast crystallisation was observed within several nanoseconds. At 600 K, tens of nucleation centres, with random grain orientations, spontaneously formed after a few nanoseconds. The crystal grains quickly grew at 3 ns, with grain sizes increasing further until 20 ns (Fig. 3c). We analysed this 20-ns SET simulation using the ACE extrapolation grade⁸¹, γ , which allows one to classify atomic configurations into interpolation and extrapolation regimes (Supplementary Fig. 4). Almost all atomic environments fell comfortably within the interpolation regime of our GST-ACE-24 potential, indicating that the nucleation and growth processes are accurately captured by our ACE model. The resulting SET state is a polycrystalline sample of rock-salt-like GST. We counted 277 crystalline grains of different crystal orientations, and the average diameter was ≈ 4.6 nm, consistent with the experimentally measured grain size in GST thin films using in situ transmission electron microscopy⁸².

We next simulated a second RESET process of the device-size model. We imposed a 40-ps heating pulse (0.036 pJ) to melt the recrystallised structure. The evolution of temperature profiles is shown in Supplementary Fig. 5. We note that the energy of this heating pulse (0.036 pJ) is smaller than that (0.064 pJ) initially used to erase the initial state of the cell (trigonal layered GST; cf. Fig. 3b); however, this smaller heating pulse still melted the whole structural model. Both the first and the second RESET pulses led to the formation of amorphous GST, with overall similar local structure compared to the results of small-scale, DFT-accessible models (Supplementary Fig. 6). We note that the overall power consumption in these simulations is much lower than that in real devices, because the input power here is directly assigned to specific atoms to increase their kinetic energy. To programme a device experimentally, the thermal energy is generated by Joule heating via electrical pulsing, which involves thermal dissipation and energy loss. Therefore, our ML-driven MD simulations provide the theoretical minimum energy values for RESET operations³¹. Nevertheless, the reduced RESET energy in our simulations implies that a polycrystal, with numerous rock-salt-like crystal grains, is much more easily melted than the stable trigonal phase of GST. We found that the structural disordering primarily occurred at the disordered grain boundaries, similar to the onset of the melting in simulations of recrystallised, polycrystalline Te⁴⁴. However, our ACE-MD simulations showed that the melting of GST also occurred inside the crystal grains; the latter has been suggested to stem from atomic migration and vacancy diffusion in rock-salt-like crystalline GST⁸³.

We show the evolution of the potential energy and the fraction of crystal-like atoms during the ACE-MD simulated full-cycle operations in Fig. 3d. These properties provide a quantitative measure of the energetics involved in switching, and reveal the degree of structural ordering at different stages of the device operations. We estimate that the full-cycle simulations (i.e., RESET to SET and back to RESET) using ACE consumed $\approx 770,000$ CPU core hours and ≈ 2500 kWh running on ARCHER2 (cf. ref. 84). With more CPU resources available, it is feasible to simulate multiple SET–RESET cycles of GST-based binary memory devices, allowing atomic-scale investigations of structural and compositional variations over repeated full-cycle operations.

Full-cycle operations for in-memory computing

Beyond their application in data-storage devices, GST alloys have also been used in neuromorphic in-memory computing tasks, which aim to process and store data directly within the same memory cell, thereby avoiding frequent data transfer between conventional memory and processing units^{4,5}. In addition to binary ones and zeroes, in-memory computing requires multiple distinct intermediate logic states to represent (near-) continuous weights or values, which are essential for analogue computations (e.g., matrix–vector multiplications). In fact, the electrical-resistance level of GST depends on the ratio of the crystalline to the amorphous volume, making it possible to obtain multiple logic states via appropriate iterative RESET and cumulative SET operations. Such operations can be achieved using small-size bottom electrodes and large programming volumes in mushroom-type devices⁸. As shown in Fig. 4a, given a large programming volume, heating pulses of different amplitudes can thus create mushroom-like active regions with very different crystalline-to-amorphous ratios. Given that the diameter of the bottom electrode can be scaled down to ≈ 3 nm (ref. 85), the dimensions of state-of-the-art mushroom-type devices^{4,5} could be further miniaturised, from hundreds to tens of nanometres—providing a broad, tuneable range of cell dimensions for optimisation.

Here, we demonstrate ACE-driven, full-cycle simulations of such partial programming in mushroom-type cells. We simulated a cross-section of GST of 100×40 nm², which represents the programming in the middle of a mushroom-type cell (Fig. 4b). We set the thickness of the slab model to 5 nm, corresponding to a quasi-two-dimensional periodic box. In total, this structural model contains 794,808 atoms, much larger than the model size used to describe a mushroom-type geometry in our previous work³¹. The initial configuration is a rock-salt-like crystalline phase of Ge₁Sb₂Te₄, corresponding to an idealised single crystal with cation/vacancy disorder but no grain boundaries. A heat barrier (≈ 6 -nm-thick slab of amorphous Ge₁Sb₂Te₄) was added on

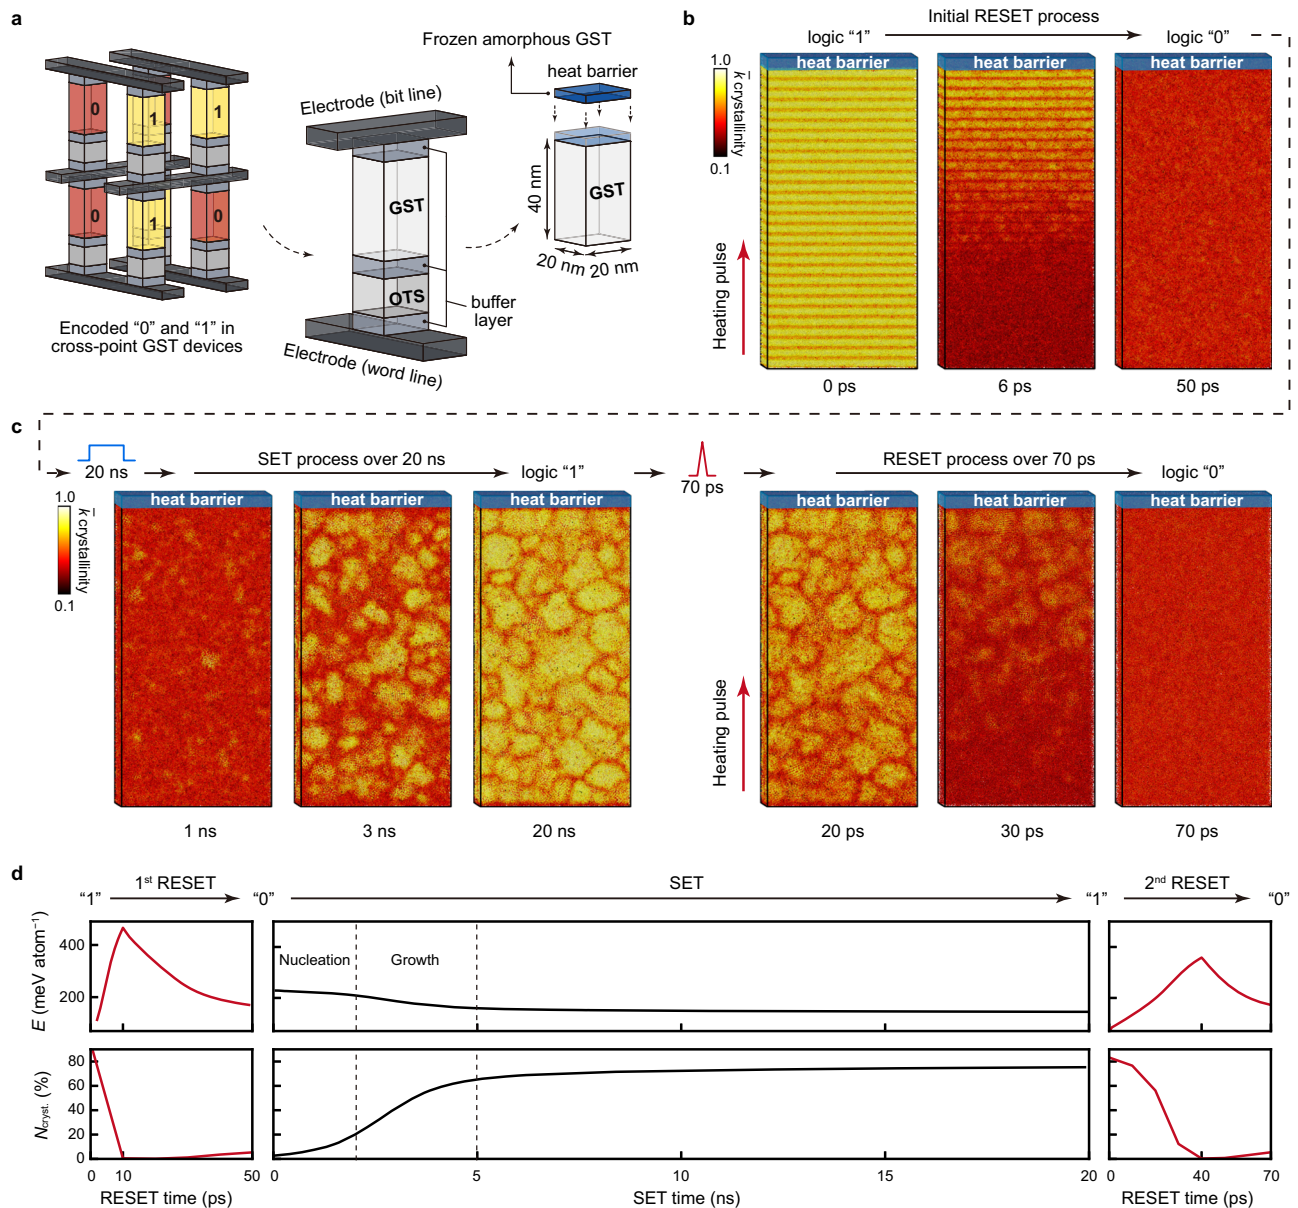


Fig. 3 | Full-cycle device-scale simulations of cross-point Ge-Sb-Te (GST) devices. **a** Schematic of cross-point devices, in which logic “0” and “1” bits are encoded by amorphous and crystalline states of GST. In these devices, GST layers and ovonic threshold switching (OTS) selector layers are sandwiched by buffer layers. A device-size structural model was built as in ref. 31 ($20 \times 20 \times 40 \text{ nm}^3$; 532,980 atoms); note that we here use amorphous GST as a heat barrier. **b** The initial RESET operation, simulated similar to ref. 31 but now using the GST-ACE-24 model, starting from layered trigonal $\text{Ge}_1\text{Sb}_2\text{Te}_4$, triggered by a 10-ps heating pulse (0.064 pJ) from the bottom to the top. After the programming pulse, a 40-ps cooling was performed by removing the added kinetic energy from the structural

model until it reached 300 K. **c** The subsequent SET operation at 600 K, simulated over 20 ns in the NVT ensemble. The resultant recrystallised GST structure contains 277 crystalline grains with an average diameter of $\approx 4.6 \text{ nm}$. A second RESET operation was then simulated, via a 40-ps heating pulse (0.036 pJ) and a 30-ps cooling process (Supplementary Note 3). Colour coding in (b–c) indicates the smooth overlap of atomic positions (SOAP)-based crystallinity measure, \bar{k} (see ref. 18). **d** Computed potential energy and fraction of crystal-like atoms during the full-cycle simulations. A \bar{k} cut-off of 0.57 was used to separate crystal-like and amorphous-like atoms¹⁸. Dashed lines indicate the nucleation and growth processes during crystallisation.

the top of the cell, preventing atomic migration across the periodic boundary (Fig. 4b). To simulate programming operations, heating pulses with different magnitudes were applied to regions of different sizes, representing separate logic states (Fig. 4c). We first added a small heating pulse (0.011 pJ) over 100 ps, resulting in a melted programming region with a diameter of $\approx 50 \text{ nm}$. This structural model was then quenched to 300 K over 200 ps by gradually removing kinetic energy from the structural model. We call the resulting intermediate state “logic state I”. We note that atoms outside the programming domain remained crystal-like after the heating process, leading to a large crystalline–amorphous interface (Fig. 4c).

We then simulated the crystallisation process for the logic state I at 600 K (Fig. 4d). Fast crystal growth proceeded at the crystalline–amorphous interface, leading to an evident shrinkage of the disordered-like region. Meanwhile, multiple nuclei were found inside the programming region. The crystalline seeds quickly grew in size, forming a polycrystalline domain. By distinguishing between atoms recrystallised through growth and those through nucleation, we observed a competition between growth-driven and nucleation-driven crystallisation (qualitatively similar to a recent study⁵⁶ based on a neural-network potential; see Discussion section for details). In our simulation, the growth-driven crystallisation accounts for

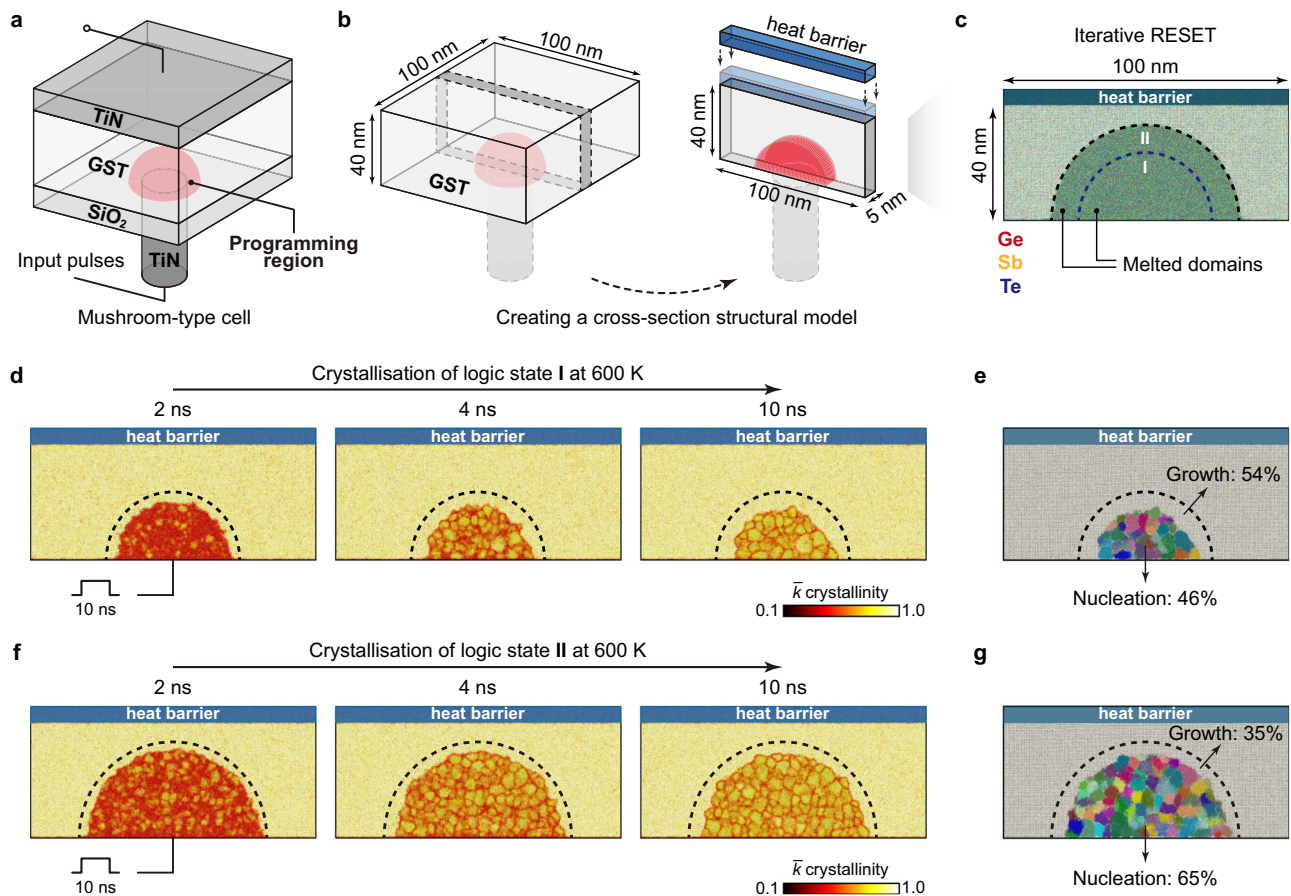


Fig. 4 | Multiple logic states in a Ge-Sb-Te (GST)-based mushroom-type device. **a** Schematic of mushroom-type cells, in which GST layers are sandwiched by top and bottom electrodes (e.g., TiN). **b** A two-dimensional slice model³¹ was built (here, $100 \times 40 \times 5 \text{ nm}^3$; 794,808 atoms), which represents a cross-section of a mushroom-type cell. **c** Two different intermediate RESET states (viz. logic states I and II) were obtained after different 100-ps heating pulses (0.011 and 0.022 pJ, respectively) and the subsequent 200-ps cooling process. The two melted domains have diameters of ≈ 50 and $\approx 70 \text{ nm}$, respectively, indicated by two different dashed lines. Ge, Sb, and Te atoms are rendered as red, yellow, and blue, respectively. **d** Crystallisation of the intermediate state I at 600 K over 10 ns. **e** Grain-segmentation analysis of the

recrystallised intermediate state I. **f–g** As panels (**d–e**), but now for the intermediate state II. Colour-coding in (**d, f**) indicates the smooth overlap of atomic positions (SOAP)-based \bar{k} crystallinity (see ref. 18), illustrating the amorphous-like (red) and crystalline-like domains (yellow). The grain analyses shown in panels **e** and **g** were carried out using polyhedral template matching¹² and grain-segmentation analyses, as implemented in OVITO¹⁰. Different colours indicate crystal grains with different orientations. The “growth” region in **e** and **g** highlight the recrystallisation driven by the growth at the crystalline-amorphous interface. The dashed lines in **d** and **e** represent the initial domain of logic state I, and those in **f** and **g** indicate the initial domain of logic state II.

54% of the recrystallised atoms, whereas nucleation contributed 46% (Fig. 4e).

We next added a larger heating pulse (0.022 pJ) to the recrystallised model and cooled it down to 300 K, which created a larger melt-quenched glassy region with a diameter of $\approx 70 \text{ nm}$. We call this intermediate state “logic state II” (Fig. 4c). In its subsequent crystallisation at 600 K (Fig. 4f), the contributions from the growth and nucleation were 35% and 65%, respectively (Fig. 4g). The increased nucleation contribution stems from the dominant nucleation-driven nature of the crystallisation in GST under these conditions. The larger the amorphous region, the more widespread the occurrence of homogeneous nucleation. This finding also implies that the SET speed in GST-based mushroom-type devices at 600 K is almost independent of the size of amorphised regions and the amplitude of the preceding RESET pulse. Rapid homogeneous nucleation is the key to such fast SET operations.

In fact, GST-based in-memory computing devices exhibit considerable resistance noise and time-dependent drift that erodes the precision and consistency of these devices^{87,88}. On the one hand, the varied recrystallised morphologies, which contained crystal grains of different orientations (Fig. 4e, g), can be the source of stochasticity in cumulative SET operations, leading to cycle-to-cycle and device-to-

device variations. On the other hand, the prominent resistance drift, believed to stem from structural relaxation of amorphous GST (known as ageing), can result in the overlap of two adjacent logic states, causing decoding errors⁸⁹. We show in Supplementary Fig. 7 that our ACE model can well describe the degree of local bond-length asymmetry, sometimes referred to as Peierls distortions, of amorphous GST—a quantitative structural fingerprint of the ageing process⁹⁰. Hence, our ACE model can simulate both stochastic recrystallisation and aged amorphous structures of mushroom-type devices, which provides atomic-scale insights into the programming mechanisms of GST-based mushroom-type devices for in-memory computing tasks.

Discussion

Our ultrafast and chemically transferable ACE potential for GST alloys can serve as a powerful “off-the-shelf” simulation tool with quantum-mechanical accuracy. Its computational efficiency enables full-cycle simulations (multiple RESET to SET operations) of different device architectures at extensive length scales (tens of nanometres) and time scales (tens of nanoseconds). We expect that our ACE model can provide atomic-scale insights into realistic programming conditions of GST-based electronics, including repeated switching for binary memory applications, as well as cumulative SET and iterative RESET

processes for neuromorphic in-memory computing. In the latter case, larger device geometries than in our current proof-of-concept simulation (Fig. 4) could make it possible to more finely tailor the amorphous-to-crystalline volume ratio to accommodate more resistance states. Simulating complex in-memory operations at the atomic scale could provide a more in-depth understanding of phase-change neuromorphic computing, and such simulations would benefit from fast and efficient ACE models. Moreover, our ACE model could also offer useful atomic-scale perspectives for GST-based waveguide memories and other emerging optical technologies^{91–94}. Unlike compact electronic devices, waveguide devices typically feature less confined geometries and require the use of the NPT ensemble (i.e., constant number of particles, pressure, and temperature) to address potential volume changes during switching processes. We note that ACE-MD simulations are well-suited to these scenarios, as they can handle the required time (tens of nanoseconds) and length scales (tens of millions of atoms).

We note that the atomistic modelling of PCMs on large length scales is gaining increasing interest in the community. From a technical perspective, the indirectly-learned ML potential of ref. 38 already illustrated the usefulness of ACE in this domain: the authors reported the simulation of a ≈ 1 -million-atom bulk $\text{Ge}_2\text{Sb}_2\text{Te}_5$ structure over 1 ns on combined CPU and GPU architectures, as well as the repeated switching of a $\approx 100,000$ -atom bulk structure³⁸. The scaling tests described in ref. 38 are qualitatively consistent with our tests on the ARCHER2 high-performance computing system (Fig. 1c–e) where applicable, although (as the authors also note) the details will depend on the specific hardware. In terms of simulation cells and protocols reflecting PCM device geometries, a recent study described the use of a neural-network ML potential for $\text{Ge}_2\text{Sb}_2\text{Te}_5$ (from ref. 25) and multiple GPU cards to perform large-scale simulations (≈ 2.8 million atoms) over several nanoseconds⁸⁶. A structural model was created by embedding an amorphous dome in a crystalline matrix to represent a mushroom-type device, and multiple thermostats were used to simulate SET operations⁸⁶, a competition between nucleation and growth from the interface was identified⁸⁶, which is qualitatively similar to Fig. 4. In our present work, we have combined a carefully optimised ACE potential that makes efficient use of CPU resources with advanced simulation protocols for cell geometries and programming conditions that are relevant to both cross-point and mushroom-type GST devices.

Although we employed a fixed amorphous GST slab as a heat barrier (cf. Figs. 3–4) to approximate the impact of interfaces present in real devices (such as TiN or SiO_2 contacts), incorporating interface effects of the surrounding materials in a realistic way is an important future step—and is expected to be technically feasible, especially given the availability of relevant previously published training datasets, e.g., for the full Si–O binary system⁴⁵. A key challenge in this extension will likely be to construct representative configurations that capture complex interface interactions involving four or more elements. To address this challenge, we highlight the use of GAP in rapidly sampling diverse chemical and structural space with minimal prior knowledge^{61,95}, thereby facilitating the construction of an initial training dataset for interfaces. As demonstrated in the present work, such GAP-based datasets can be directly fed into ACE training, followed by further domain-specific iterations. In addition, uncertainty-based sampling for active learning can be used to obtain a more comprehensive training dataset⁸¹.

Looking back on the discussion of PCM modelling at the beginning of this paper, we note that ML-driven simulation methods have now been established in the field, allowing for wide-ranging simulation studies of functional materials, and increasingly becoming of relevance to experimental work and practical applications. Our study has exemplified this advance for the field of electronic memories and neuromorphic in-memory computing, and other atomistic ML models have been developed for a wide range of applications across different

disciplines: recently published ML potentials have been used in the search for new stable inorganic crystals (e.g., for layered materials and solid-electrolyte candidates)⁹⁶, in the prediction of supercritical behaviour in high-pressure liquid hydrogen, relevant to the structure and evolution of giant planets⁹⁷, or in biomolecular-dynamics simulations of protein-folding processes and their thermodynamics⁹⁸. The relevance of ML-driven simulations in the computational design of amorphous materials—PCMs and many others—has been pointed out in ref. 99. A very recent preprint discusses the role of ML potentials for device-scale modelling in a wider perspective¹⁰⁰. We expect that our present work will stimulate the further development of efficient ML potentials for exploring structurally and chemically more complex PCM systems and devices, and provide a key approach for investigating scientific questions related to memory and computing applications.

Methods

The GST-ACE-24 potential

All ACE models shown in the present work were fitted using `pace-maker` (version 0.2.7; ref. 40); their optimisation was carried out with XPOT (version 1.1.0; refs. 48,50). The extension of XPOT to ACE specifically, and the physical role of relevant hyperparameters, has been discussed in our more technical study in ref. 49. The latter includes investigations of ACE models for silicon and the binary compound Sb_2Te_3 and provides a basis for the present work.

Using XPOT, we optimised 4 hyperparameters (cf. Supplementary Table 2) based on the iter-0 dataset and performed 32 fitting iterations (cf. Fig. 1b). To guide the target of the XPOT optimisation, we defined a testing dataset consisting of conventional disordered structures (≈ 200 atoms each) and intermediate configurations during phase transitions (1008 atoms each). These two types of structures were taken from AIMD simulations reported in ref. 31 and ref. 18, respectively. This testing dataset was also used in the computation of RMSE values shown in Table 1. In the XPOT optimisations, we first performed 8 exploratory fits using a Hammersley sequence to sample hyperparameters. Next, Bayesian Optimisation (BO) was used to optimise the hyperparameters over the remaining iterations. After XPOT optimisation, we “upfitted” the best potential (with an increased relative weighting of the energy; see ref. 49). In fact, after iter-3, we performed another XPOT run to determine whether the model required further hyperparameter optimisation based on the newly added configurations (i.e., those from iter-1 to iter-3). However, we found no notable improvements in accuracy on the testing dataset, and therefore continued with the existing hyperparameters, as optimised on the iter-0 dataset. We note that the hyperparameters determined here for GST-ACE-24 were also used in fitting a separate ACE model for elemental tellurium, which is described in ref. 44.

The final potential model combines linear and nonlinear embeddings of the atomic neighbour environments over 3000 basis functions and uses a radial cut-off of 8 Å. Training structures were weighted, based on their configuration types. Crystalline structures, melt-quench structures from AIMD, and RSS structures were given custom weightings to guide model accuracy in these regions. The model was fitted using an NVIDIA A100 GPU.

We note that a positive core-repulsion term can be included in ACE models to stop unphysical energies and forces from being produced at short atomic distances and to correct the core-repulsion behaviour, which can mitigate simulation issues with “lost” atoms⁴⁰; such an approach has been taken for the ACE models of ref. 38. By contrast, adding high-energy, small-scale random structures helps to explore a diverse configurational space (see, e.g., refs. 51,55). In particular, such additional configurations improve the ability of potentials to describe two-body interactions at unusual distances, preventing the potentials from predicting the formation of clusters which, when evaluated with DFT, were found to be energetically unfavourable. We

note that the addition of training data to represent short interatomic distances, “rather than relying on the core repulsion completely”, has been suggested by the `pacemaker` developers (see ref. 101). Our GST-ACE-24 model does not use a separate core-repulsion term.

Validation

We computed different structural properties of various amorphous GST compounds along the GeTe–Sb₂Te₃ compositional tie-line, such as radial and angular distribution functions (Supplementary Fig. 7), and found that the predictions of our GST-ACE-24 model agreed very well with the AIMD data of ref. 31. Also, our ACE model faithfully reproduced the fraction of homopolar bonds and tetrahedral motifs, as well as the degree of local bond-length asymmetry (Supplementary Fig. 7), which are important structural factors that have been discussed in the context of ageing phenomena in the amorphous phase⁹⁰. These structural validations demonstrate that our new ACE potential is both structurally and chemically transferable and can accurately describe disordered GST structures across various compositions, consistent with results for the GST-GAP-22 model³¹.

Computational performance

A key point in the present study is how ACE allows for ultra-fast device-scale simulations on a CPU-based high-performance computing system, without requiring GPU hardware at runtime. In addition to the simple and fast summation operations performed in the construction and inference of the ACE model, a recursive evaluation algorithm is used to construct the basis functions, reducing the number of arithmetic operations, and thus improving numerical efficiency³⁹. We measured the performance of our GST-ACE-24 model by comparing against the published GST-GAP-22 potential³¹ on the CPU cores of the ARCHER2 system. The compute nodes each have 128 CPU cores, and the memory per node is either 256 GB (standard nodes) or 512 GB (high-memory nodes); see ref. 102 for details. The comparison between ACE and GAP is shown in Fig. 1c–e.

We note that ACE also supports multi-GPU computation: our test on a 1-million-atom structural model showed good scalability of ACE-MD simulations running on up to four NVIDIA A100 GPUs with 80 GB of memory each (Supplementary Fig. 2). In addition, we compared the computational efficiency of GST-ACE-24 with a directly re-fitted equivariant neural-network potential, based on the MACE architecture^{34,56}, on a GPU. This directly re-fitted MACE model used the same training dataset of GST-ACE-24. We performed the tests for GST-ACE-24 and the MACE model on an NVIDIA A100 GPU with 80 GB of memory, using a 10,000-atom structural model. The computational efficiency of GST-ACE-24 in this setting was ≈ 2 million MD steps per day, whereas the computational efficiency of the MACE model was $\approx 335,000$ MD steps per day.

DFT computations

The AIMD data used for the fitting process (Fig. 1b) and the validation (Supplementary Fig. 7) of our ACE model were taken from our previous work (ref. 31). These AIMD simulations had been carried out using the “second-generation” Car–Parrinello scheme, as implemented in the Quickstep code of CP2K (version 2023.1)¹⁰³, a combination of Gaussian-type and plane-wave basis sets, scalar-relativistic Goedecker pseudopotentials¹⁰⁴, and the Perdew–Burke–Ernzerhof (PBE) functional⁴⁷. Details of the AIMD simulations may be found in ref. 31.

To label the reference dataset, we computed the per-structure energies and per-atom forces by performing single-point DFT computations using the Vienna Ab initio Simulation Package (VASP; version 5.4.4)^{105,106} with projector augmented-wave (PAW) pseudopotentials^{107,108}. We used a 600 eV cut-off for plane waves and an energy tolerance of 10^{-7} eV per cell for SCF convergence. An automatically generated k -point grid with a maximum spacing of 0.2 \AA^{-1} was used to sample reciprocal space.

Molecular-dynamics simulations

MD simulations were carried out with the GST-GAP-22 (ref. 31) and GST-ACE-24 ML potential models, using LAMMPS (version 15 Jun 2023)¹⁰⁹, with interfaces to QUIP and `pacemaker`, respectively. The canonical ensemble (NVT) and the microcanonical ensemble (NVE) were used in this work. A Langevin thermostat was used to control the temperature in the NVT simulations. We simulated non-isothermal heating processes in the NVE ensemble. Additional energy was added to the kinetic energy of the atoms in the programming regions (Supplementary Note 3), with a timestep of 2 ps. The timestep for all ML-driven MD simulations was 2 fs. Structures were visualised using OVITO¹¹⁰.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

The raw data for the figures presented in the main text and Supplementary Information have been provided in a Source Data file with this paper. The rest of the data supporting the present study, including the potential parameters, fitting data and structural models shown in Figs. 3–4, are publicly available via Zenodo at <https://doi.org/10.5281/zenodo.14755074> (ref. 111). Source data are provided with this paper.

Code availability

The XPOT software used for hyperparameter optimisation is available at <https://github.com/dft-dutoit/XPOT> under the GPL-2.0 licence. A copy has been deposited in Zenodo and is available at <https://doi.org/10.5281/zenodo.15853809> (ref. 50). Other software packages were used as provided by their respective authors.

References

- Zhang, W., Mazzarello, R., Wuttig, M. & Ma, E. Designing crystallization in phase-change materials for universal memory and neuro-inspired computing. *Nat. Rev. Mater.* **4**, 150–168 (2019).
- Feldmann, J., Youngblood, N., Wright, C. D., Bhaskaran, H. & Pernice, W. H. P. All-optical spiking neurosynaptic networks with self-learning capabilities. *Nature* **569**, 208–214 (2019).
- Pellizzer, F., Pirovano, A., Bez, R. & Meyer, R. L. Status and perspectives of chalcogenide-based crosspoint memories. In *2023 International Electron Devices Meeting (IEDM) 1–4*. <https://doi.org/10.1109/IEDM45741.2023.10413669> (2023).
- Ambrogio, S. et al. An analog-AI chip for energy-efficient speech recognition and transcription. *Nature* **620**, 768–775 (2023).
- Le Gallo, M. et al. A 64-core mixed-signal in-memory compute chip based on phase-change memory for deep neural network inference. *Nat. Electron.* **6**, 680–693 (2023).
- Zhou, W., Shen, X., Yang, X., Wang, J. & Zhang, W. Fabrication and integration of photonic devices for phase-change memory and neuromorphic computing. *Int. J. Extrem. Manuf.* **6**, 022001 (2024).
- Wuttig, M. & Yamada, N. Phase-change materials for rewriteable data storage. *Nat. Mater.* **6**, 824–832 (2007).
- Tuma, T., Pantazi, A., Le Gallo, M., Sebastian, A. & Eleftheriou, E. Stochastic phase-change neurons. *Nat. Nanotechnol.* **11**, 693–699 (2016).
- Akola, J. & Jones, R. Structural phase transitions on the nanoscale: the crucial pattern in the phase-change materials Ge₂Sb₂Te₅ and GeTe. *Phys. Rev. B* **76**, 235201 (2007).
- Caravati, S., Bernasconi, M., Kühne, T. D., Krack, M. & Parrinello, M. Coexistence of tetrahedral- and octahedral-like sites in amorphous phase change materials. *Appl. Phys. Lett.* **91**, 171906 (2007).
- Xu, M., Cheng, Y., Sheng, H. & Ma, E. Nature of atomic bonding and atomic structure in the phase-change Ge₂Sb₂Te₅ Glass. *Phys. Rev. Lett.* **103**, 195502 (2009).

12. Huang, B. & Robertson, J. Bonding origin of optical contrast in phase-change memory materials. *Phys. Rev. B* **81**, 081204 (2010).
13. Raty, J.-Y. et al. A quantum-mechanical map for bonding and properties in solids. *Adv. Mater.* **31**, 1806280 (2019).
14. Wang, X.-D. et al. Multiscale simulations of growth-dominated Sb₂Te phase-change material for non-volatile photonic applications. *npj Comput. Mater.* **9**, 136 (2023).
15. Shen, X., Chu, R., Jiang, Y. & Zhang, W. Progress on materials design and multiscale simulations for phase-change memory. *Acta Metall. Sin.* **60**, 1362–1378 (2024).
16. Hegedüs, J. & Elliott, S. R. Microscopic origin of the fast crystallization ability of Ge-Sb-Te phase-change memory materials. *Nat. Mater.* **7**, 399–405 (2008).
17. Kalikka, J., Akola, J. & Jones, R. O. Crystallization processes in the phase change material Ge₂Sb₂Te₅: Unbiased density functional/molecular dynamics simulations. *Phys. Rev. B* **94**, 134105 (2016).
18. Xu, Y. et al. Unraveling crystallization mechanisms and electronic structure of phase-change materials by large-scale ab initio simulations. *Adv. Mater.* **34**, 2109139 (2022).
19. Behler, J. First principles neural network potentials for reactive simulations of large molecular and condensed systems. *Angew. Chem. Int. Ed.* **56**, 12828–12840 (2017).
20. Deringer, V. L., Caro, M. A. & Csányi, G. Machine learning interatomic potentials as emerging tools for materials science. *Adv. Mater.* **31**, 1902765 (2019).
21. Friederich, P., Häse, F., Proppe, J. & Aspuru-Guzik, A. Machine-learned potentials for next-generation matter simulations. *Nat. Mater.* **20**, 750–761 (2021).
22. Sosso, G. C., Miceli, G., Caravati, S., Behler, J. & Bernasconi, M. Neural network interatomic potential for the phase change material GeTe. *Phys. Rev. B* **85**, 174103 (2012).
23. Behler, J. & Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **98**, 146401 (2007).
24. Sosso, G. C., Salvalaglio, M., Behler, J., Bernasconi, M. & Parrinello, M. Heterogeneous crystallization of the phase change material GeTe via atomistic simulations. *J. Phys. Chem. C* **119**, 6428–6434 (2015).
25. Abou El Kheir, O., Bonati, L., Parrinello, M. & Bernasconi, M. Unraveling the crystallization kinetics of the Ge₂Sb₂Te₅ phase change compound with a machine-learned interatomic potential. *npj Comput. Mater.* **10**, 33 (2024).
26. Gabardi, S. et al. Atomistic simulations of the crystallization and aging of GeTe nanowires. *J. Phys. Chem. C* **121**, 23827–23838 (2017).
27. Mocanu, F. C. et al. Modeling the phase-change memory material, Ge₂Sb₂Te₅, with a machine-learned interatomic potential. *J. Phys. Chem. B* **122**, 8998–9006 (2018).
28. Dragoni, D., Behler, J. & Bernasconi, M. Mechanism of amorphous phase stabilization in ultrathin films of monoatomic phase change material. *Nanoscale* **13**, 16146–16155 (2021).
29. Mo, P. et al. Accurate and efficient molecular dynamics based on machine learning and non-Von Neumann architecture. *npj Comput. Mater.* **8**, 107 (2022).
30. Bartók, A. P., Payne, M. C., Kondor, R. & Csányi, G. Gaussian approximation potentials: the accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.* **104**, 136403 (2010).
31. Zhou, Y., Zhang, W., Ma, E. & Deringer, V. L. Device-scale atomistic modelling of phase-change memory materials. *Nat. Electron.* **6**, 746–754 (2023).
32. Li, K., Liu, B., Zhou, J. & Sun, Z. Revealing the crystallization dynamics of Sb–Te phase change materials by large-scale simulations. *J. Mater. Chem. C* **12**, 3897–3906 (2024).
33. Batzner, S. et al. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat. Commun.* **13**, 2453 (2022).
34. Batatia, I., Kovacs, D. P., Simm, G., Ortner, C. & Csányi, G. MACE: higher order equivariant message passing neural networks for fast and accurate force fields. *Advances in Neural Information Processing Systems*. Vol. 35, 11423–11436 (Curran Associates, Inc., 2022).
35. Wang, G., Sun, Y., Zhou, J. & Sun, Z. PotentialMind: graph convolutional machine learning potential for Sb–Te binary compounds of multiple stoichiometries. *J. Phys. Chem. C* **127**, 24724–24733 (2023).
36. Chang, C., Deringer, V. L., Katti, K. S., Van Speybroeck, V. & Wolverton, C. M. Simulations in the era of exascale computing. *Nat. Rev. Mater.* **8**, 309–313 (2023).
37. Drautz, R. Atomic cluster expansion for accurate and transferable interatomic potentials. *Phys. Rev. B* **99**, 014104 (2019).
38. Dunton, O. R., Arbaugh, T. & Starr, F. W. Computationally efficient machine-learned model for GST phase change materials via direct and indirect learning. *J. Chem. Phys.* **162**, 034501 (2025).
39. Lysogorskiy, Y. et al. Performant implementation of the atomic cluster expansion (PACE) and application to copper and silicon. *npj Comput. Mater.* **7**, 97 (2021).
40. Bochkarev, A. et al. Efficient parametrization of the atomic cluster expansion. *Phys. Rev. Mater.* **6**, 013804 (2022).
41. Dusson, G. et al. Atomic cluster expansion: completeness, efficiency and stability. *J. Comput. Phys.* **454**, 110946 (2022).
42. Qamar, M., Mrovec, M., Lysogorskiy, Y., Bochkarev, A. & Drautz, R. Atomic cluster expansion for quantum-accurate large-scale simulations of carbon. *J. Chem. Theory Comput.* **19**, 5151–5167 (2023).
43. Deringer, V. L. et al. Gaussian process regression for materials and molecules. *Chem. Rev.* **121**, 10073–10141 (2021).
44. Zhou, Y., Elliott, S. R., Toit, D. F. T. Du, Z. W. & Deringer, V. L. The pathway to chirality in elemental tellurium. Preprint at arXiv:2409.03860 (2024).
45. Erhard, L. C., Rohrer, J., Albe, K. & Deringer, V. L. Modelling atomic and nanoscale structure in the silicon–oxygen system through active machine learning. *Nat. Commun.* **15**, 1927 (2024).
46. Morrow, J. D. & Deringer, V. L. Indirect learning and physically guided validation of interatomic potential models. *J. Chem. Phys.* **157**, 104105 (2022).
47. Perdew, J. P., Burke, K. & Ernzerhof, M. Generalized gradient approximation made simple. *Phys. Rev. Lett.* **77**, 3865–3868 (1996).
48. Thomas du Toit, D. F. & Deringer, V. L. Cross-platform hyperparameter optimization for machine learning interatomic potentials. *J. Chem. Phys.* **159**, 024803 (2023).
49. Thomas du Toit, D. F., Zhou, Y. & Deringer, V. L. Hyperparameter optimization for atomic cluster expansion potentials. *J. Chem. Theory Comput.* **20**, 10103–10113 (2024).
50. Thomas du Toit, D. F. dft-dutoit/XPOT: ACE Release, Zenodo, <https://doi.org/10.5281/zenodo.15853809> (2025).
51. Pickard, C. J. & Needs, R. J. High-pressure phases of silane. *Phys. Rev. Lett.* **97**, 045504 (2006).
52. Pickard, C. J. & Needs, R. J. Ab initio random structure searching. *J. Phys. Condens. Matter* **23**, 053201 (2011).
53. Deringer, V. L., Proserpio, D. M., Csányi, G. & Pickard, C. J. Data-driven learning and prediction of inorganic crystal structures. *Faraday Discuss.* **211**, 45–59 (2018).
54. Deringer, V. L., Pickard, C. J. & Csányi, G. Data-driven learning of total and local energies in elemental boron. *Phys. Rev. Lett.* **120**, 156001 (2018).
55. Bernstein, N., Csányi, G. & Deringer, V. L. De novo exploration and self-guided learning of potential-energy surfaces. *npj Comput. Mater.* **5**, 99 (2019).
56. Batatia, I. et al. The design space of E(3)-equivariant atom-centred interatomic potentials. *Nat. Mach. Intell.* **7**, 56–67 (2025).
57. Stocker, S., Gasteiger, J., Becker, F., Günnemann, S. & Margraf, J. T. How robust are modern graph neural network potentials in long

- and hot molecular dynamics simulations?. *Mach. Learn. Sci. Technol.* **3**, 045010 (2022).
58. Unke, O. T. et al. Machine learning force fields. *Chem. Rev.* **121**, 10142–10186 (2021).
59. Ben Mahmoud, C., Gardner, J. L. A. & Deringer, V. L. Data as the next challenge in atomistic machine learning. *Nat. Comput. Sci.* **4**, 384–387 (2024).
60. Allen, A. E. A. et al. Learning together: towards foundation models for machine learning interatomic potentials with meta-learning. *npj Comput. Mater.* **10**, 154 (2024).
61. Liu, Y. et al. An automated framework for exploring and learning potential-energy surfaces. *Nat. Commun.* **16**, 7666 (2025).
62. Deringer, V. L., Caro, M. A. & Csányi, G. A general-purpose machine-learning force field for bulk and nanostructured phosphorus. *Nat. Commun.* **11**, 5461 (2020).
63. Pickard, C. J. Ephemeral data derived potentials for random structure search. *Phys. Rev. B* **106**, 014102 (2022).
64. Pickard, C. J. Beyond theory-driven discovery: introducing hot random search and datum-derived structures. *Faraday Discuss.* **256**, 61–84 (2025).
65. Waldecker, L. et al. Time-domain separation of optical properties from structural transitions in resonantly bonded materials. *Nat. Mater.* **14**, 991–995 (2015).
66. Feldmann, J. et al. Calculating with light using a chip-scale all-optical abacus. *Nat. Commun.* **8**, 1256 (2017).
67. Bartók, A. P., Kondor, R. & Csányi, G. On representing chemical environments. *Phys. Rev. B* **87**, 184115 (2013).
68. Welnic, W. & Wuttig, M. Reversible switching in phase-change materials. *Mater. Today* **11**, 20–27 (2008).
69. Loke, D. et al. Breaking the speed limits of phase-change memory. *Science* **336**, 1566–1569 (2012).
70. Loke, D. K. et al. Ultrafast nanoscale phase-change memory enabled by single-pulse conditioning. *ACS Appl. Mater. Interfaces* **10**, 41855–41860 (2018).
71. Li, Z., Si, C., Zhou, J., Xu, H. & Sun, Z. Yttrium-doped Sb_2Te_3 : a promising material for phase-change memory. *ACS Appl. Mater. Interfaces* **8**, 26126–26134 (2016).
72. Rao, F. et al. Reducing the stochasticity of crystal nucleation to enable subnanosecond memory writing. *Science* **358**, 1423–1427 (2017).
73. Wang, Y. et al. Scandium doped $\text{Ge}_2\text{Sb}_2\text{Te}_5$ for high-speed and low-power-consumption phase change memory. *Appl. Phys. Lett.* **112**, 133104 (2018).
74. Hu, S., Xiao, J., Zhou, J., Elliott, S. R. & Sun, Z. Synergy effect of co-doping Sc and Y in Sb_2Te_3 for phase-change memory. *J. Mater. Chem. C* **8**, 6672–6679 (2020).
75. Wang, X.-P. et al. Time-dependent density-functional theory molecular-dynamics study on amorphization of Sc-Sb-Te alloy under optical excitation. *npj Comput. Mater.* **6**, 31 (2020).
76. Ronneberger, I., Zhang, W., Eshet, H. & Mazzarello, R. Crystallization properties of the $\text{Ge}_2\text{Sb}_2\text{Te}_5$ phase-change compound from advanced simulations. *Adv. Funct. Mater.* **25**, 6407–6413 (2015).
77. Laio, A. & Parrinello, M. Escaping free-energy minima. *Proc. Natl. Acad. Sci. USA* **99**, 12562–12566 (2002).
78. Kalikka, J., Akola, J., Larrucea, J. & Jones, R. O. Nucleus-driven crystallization of amorphous $\text{Ge}_2\text{Sb}_2\text{Te}_5$: a density functional study. *Phys. Rev. B* **86**, 144113 (2012).
79. Cheng, H. Y. et al. Atomic-level engineering of phase change material for novel fast-switching and high-endurance PCM for storage class memory application. In *2013 IEEE International Electron Devices Meeting* 30.6.1–30.6.4. <https://doi.org/10.1109/IEDM.2013.6724726> (2013).
80. Cheng, H.-Y., Carta, F., Chien, W.-C., Lung, H.-L. & BrightSky, M. J. 3D cross-point phase-change memory for storage-class memory. *J. Phys. D Appl. Phys.* **52**, 473002 (2019).
81. Lysogorskiy, Y., Bochkarev, A., Mrovec, M. & Drautz, R. Active learning strategies for atomic cluster expansion models. *Phys. Rev. Mater.* **7**, 043801 (2023).
82. Park, Y. J., Lee, J. Y. & Kim, Y. T. In situ transmission electron microscopy study of the nucleation and grain growth of $\text{Ge}_2\text{Sb}_2\text{Te}_5$ thin films. *Appl. Surf. Sci.* **252**, 8102–8106 (2006).
83. Zhang, B. et al. Vacancy structures and melting behavior in rock-salt GeSbTe . *Sci. Rep.* **6**, 25453 (2016).
84. Energy use and emissions - ARCHER2 user guide. <https://docs.archer2.ac.uk/user-guide/energy/>.
85. Song, Z. et al. 12-state multi-level cell storage implemented in a 128 Mb phase change memory chip. *Nanoscale* **13**, 10455–10461 (2021).
86. Abou El Kheir, O. & Bernasconi, M. Million-atom simulation of the set process in phase change memories at the real device scale. *Adv. Electron. Mater.* **11**, e2500110 (2025).
87. Nandakumar, S. R. et al. A phase-change memory model for neuromorphic computing. *J. Appl. Phys.* **124**, 152135 (2018).
88. Sebastian, A. et al. Tutorial: Brain-inspired computing using phase-change memory devices. *J. Appl. Phys.* **124**, 111101 (2018).
89. Zhang, W. & Ma, E. Unveiling the structural origin to control resistance drift in phase-change memory materials. *Mater. Today* **41**, 156–176 (2020).
90. Raty, J. Y. et al. Aging mechanisms in amorphous phase-change materials. *Nat. Commun.* **6**, 7467 (2015).
91. Hosseini, P., Wright, C. D. & Bhaskaran, H. An optoelectronic framework enabled by low-dimensional phase-change films. *Nature* **511**, 206–211 (2014).
92. Du, K.-K. et al. Control over emissivity of zero-static-power thermal emitters based on phase-changing material GST. *Light Sci. Appl.* **6**, e16194 (2017).
93. Dong, W. et al. Tunable mid-infrared phase-change metasurface. *Adv. Opt. Mater.* **6**, 1701346 (2018).
94. Wang, D. et al. Non-volatile tunable optics by design: from chalcogenide phase-change materials to device structures. *Mater. Today* **68**, 334–355 (2023).
95. El-Machachi, Z. et al. Accelerated first-principles exploration of structure and reactivity in graphene oxide. *Angew. Chem. Int. Ed.* **63**, e202410088 (2024).
96. Merchant, A. et al. Scaling deep learning for materials discovery. *Nature* **624**, 80–85 (2023).
97. Cheng, B., Mazzola, G., Pickard, C. J. & Ceriotti, M. Evidence for supercritical behaviour of high-pressure liquid hydrogen. *Nature* **585**, 217–220 (2020).
98. Wang, T. et al. Ab initio characterization of protein molecular dynamics with Al^2BMD . *Nature* **635**, 1019–1027 (2024).
99. Liu, Y., Madanchi, A., Anker, A. S., Simine, L. & Deringer, V. L. The amorphous state as a frontier in computational materials design. *Nat. Rev. Mater.* **10**, 228–241 (2024).
100. Miret, S., Lee, K. L. K., Gonzales, C., Mannan, S. & Krishnan, N. M. A. Energy & Force Regression on DFT Trajectories is Not Enough for Universal Machine Learning Interatomic Potentials. Preprint at arXiv:2502.03660 (2025).
101. Frequently asked questions (FAQ). <https://pacemaker.readthedocs.io/en/latest/pacemaker/faq> (2025).
102. Beckett, G. et al. ARCHER2 Service Description. <https://zenodo.org/records/14507040>, <https://doi.org/10.5281/zenodo.14507040> (2024).
103. Kühne, T. D. et al. CP2K: an electronic structure and molecular dynamics software package - Quickstep: Efficient and accurate electronic structure calculations. *J. Chem. Phys.* **152**, 194103 (2020).
104. Goedecker, S., Teter, M. & Hutter, J. Separable dual-space Gaussian pseudopotentials. *Phys. Rev. B* **54**, 1703 (1996).
105. Kresse, G. & Hafner, J. Ab initio molecular dynamics for liquid metals. *Phys. Rev. B* **47**, 558–561 (1993).

106. Kresse, G. & Furthmüller, J. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B* **54**, 11169–11186 (1996).
 107. Blöchl, P. E. Projector augmented-wave method. *Phys. Rev. B* **50**, 17953–17979 (1994).
 108. Kresse, G. & Joubert, D. From ultrasoft pseudopotentials to the projector augmented-wave method. *Phys. Rev. B* **59**, 1758 (1999).
 109. Thompson, A. P. et al. LAMMPS - a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales. *Comput. Phys. Commun.* **271**, 108171 (2022).
 110. Stukowski, A. Visualization and analysis of atomistic simulation data with OVITO—the Open Visualization Tool. *Model. Simul. Mater. Sci. Eng.* **18**, 015012 (2010).
 111. Zhou, Y., Thomas du Toit, D. F., Elliott, S. R., Zhang, W. & Deringer, V. L. Research data for “Full-cycle device-scale simulations of memory materials with a tailored atomic-cluster-expansion potential”. *Zenodo* <https://doi.org/10.5281/zenodo.14755074> (2025).
 112. Larsen, P. M., Schmidt, S. & Schiøtz, J. Robust structural identification via polyhedral template matching. *Modelling Simul. Mater. Sci. Eng.* **24**, 055007 (2016).
- large-scale molecular-dynamics simulations and visualised the results. All authors (Y.Z., D.F.T.d.T., S.R.E., W.Z., and V.L.D.) contributed to discussions and to the writing of the paper.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-025-63732-4>.

Correspondence and requests for materials should be addressed to Wei Zhang or Volker L. Deringer.

Peer review information *Nature Communications* thanks the anonymous reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

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

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

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

© The Author(s) 2025

Acknowledgements

Y.Z. acknowledges a China Scholarship Council-University of Oxford scholarship. S.R.E. acknowledges the Leverhulme Trust (UK) for a Fellowship. W.Z. thanks support by the National Key Research and Development Programme of China (2023YFB4404500), the National Natural Science Foundation of China (62374131), the Computing Centre in Xi’an and the International Joint Laboratory for Micro/Nano Manufacturing and Measurement Technologies of XJTU. This work was supported by UK Research and Innovation [grant number EP/X016188/1]. We are grateful for computational support from the UK national high-performance computing service, ARCHER2, for which access was obtained via the UKCP consortium and funded by EPSRC grant ref EP/X035891/1, as well as through a separate EPSRC Access to High-Performance Computing award.

Author contributions

Y.Z., W.Z., and V.L.D. designed the study. Y.Z. and D.F.T.d.T. parameterised the ACE potential models. D.F.T.d.T. studied the role of hyperparameters and provided technical advice. Y.Z. carried out the