

# PRISM: Periodic Representation with multiscale and Similarity graph Modelling for enhanced crystal structure property prediction

Àlex Solé<sup>1,2</sup>, Albert Mosella-Montoro<sup>1</sup>, Joan Cardona<sup>2</sup>, Daniel Aravena<sup>\*3</sup>, Silvia Gómez-Coca<sup>\*2</sup>, Eliseo Ruiz<sup>\*2</sup>, and Javier Ruiz-Hidalgo<sup>\*1</sup>

<sup>1</sup>Image Processing Group – Signal Theory and Communications Department, Universitat Politècnica de Catalunya, Barcelona, Spain

<sup>2</sup>Inorganic and Organic Chemistry Department and Institute of Theoretical and Computational Chemistry, Universitat de Barcelona, Barcelona, Spain

<sup>3</sup>Materials Chemistry Department, Faculty of Chemistry and Biology, Universidad de Santiago de Chile, Santiago, Chile

**E-mail:** daniel.aravena.p@usach.cl; silvia.gomez@qi.ub.edu; eliseo.ruiz@qi.ub.edu; j.ruiz@upc.edu

## Abstract

Crystal structures are characterised by repeating atomic patterns within unit cells across three-dimensional space, posing unique challenges for graph-based representation learning. Current methods often overlook essential periodic boundary conditions and multiscale interactions inherent to crystalline structures. In this paper, we introduce PRISM, a graph neural network framework that explicitly integrates multiscale representations and periodic feature encoding by employing a set of expert modules, each specialised in encoding distinct structural and chemical aspects of periodic systems. Extensive experiments across crystal structure-based benchmarks demonstrate that PRISM improves state-of-the-art predictive accuracy, significantly enhancing crystal property prediction. The project page can be found at: <https://imatge-upc.github.io/PRISM/>.

## Introduction

Accurate prediction of crystal material properties is essential for accelerating the discovery of novel materials for applications such as energy storage, catalysis, and electronics [10, 17, 27, 1, 3, 21]. While density functional theory (DFT) provides reliable predictions [24], its high computational cost limits its applicability

for large-scale screening [11]. Consequently, machine learning approaches that leverage graph representations of crystal structures have emerged as a promising alternative [33, 4, 15, 31].

Graph Neural Networks (GNNs) have served as a powerful framework for modelling periodic crystal structures by representing atoms as nodes and inter-atomic relationships as edges. Early methods such as CGCNN [33] introduced multi-edge crystal graphs built upon Euclidean distances to capture short-range atomic interactions under periodic boundary conditions. Extensions like MEGNET [4] adopted enriched atom and bond embeddings, while GATGNN [15] incorporated attention mechanisms to aggregate local chemical and structural features more effectively. Angle-based line graphs in ALIGNN [31] and M3GNet [5] further enhance geometric expressiveness, although at an additional computational cost tied to the average number of neighbours per atom. Techniques like MatFormer [35] use self-loops for encoding periodic pattern awareness, whereas PotNet [13] approximates infinite interactions between every pair of atoms, incurring a quadratic overhead in the number of atoms. More recent work (eComformer [34], iComformer [34], and CartNet [30]) address orientation ambiguities, either by constructing orientation invariant or equivariant representations of the lattice vectors, or by combining rotational data augmentation with explicit direction vectors. Nonetheless, most rely primarily on distance-only or partially invariant embeddings, which can limit the capacity to model certain anisotropic or long-distance crystal properties.

Recent works [34, 30] use graph neural networks where the graph is created using atom geometric proximity policies such as k-Nearest Neighbours (k-NN) or radius. However, these techniques do not explicitly encode the unit cell; they focus solely on atom-level interactions with periodic boundary conditions applied in the geometrical space. Moreover, those techniques struggle to capture long-distance interactions, which are crucial for understanding the global structural context of crystalline materials.

Beyond geometry-based neighbourhoods such as fixed k-NN or radius graphs [19, 25], connecting points directly in the learned feature space often yields richer task-specific relations. DGCNN [32] recomputes the k-NN graph from the current feature space at every layer so edges link feature-nearest rather than Euclidean-nearest points and semantic information can move between regions that are far apart in space. Building on this principle, MUNEGC [20] constructs one graph in Euclidean space to preserve local geometric context and another in feature space to connect semantically similar yet spatially distant points. It then aggregates messages from both graphs to obtain a representation that blends geometric and semantic cues. In molecular modelling, So3krates [8] further extends these ideas by introducing feature-based neighbours that capture nonlocal quantum interactions, moving beyond traditional short-range or coordinate-based connections. Nevertheless, these methods were originally developed for finite or non-repetitive 3D tasks and often omit periodic boundary conditions, which complicates their application to crystalline systems.

Furthermore, a common strategy in molecules or 3D point-cloud tasks [20, 8, 22, 14] is to employ a double-neighbourhood approach that captures infor-

mation from both the geometrical domain and the feature space. However, when applied to crystals, it is crucial to encode periodicity directly within the feature space. Without this encoding, the feature space remains oblivious to the inherent repeating patterns, which can ultimately limit predictive performance. Additionally, as demonstrated by Keqiang et al. [34], the same material can have multiple equivalent cell representations. This variability must be addressed during the feature graph construction to ensure invariance with respect to the cell representation.

Despite their popularity, fixed-radius or k-NN graphs limit message passing to local neighbourhoods and cannot capture the long-distance lattice correlations that often govern crystal properties. Likewise, feature-space graph methods, developed primarily for finite point clouds, omit explicit periodic boundary conditions, thus failing to exploit translational symmetry. Our approach addresses both shortcomings by enforcing minimal-image periodicity in its feature-space graph construction and fusing this with multiscale atomistic and cell-level graphs, enabling unified encoding of short-range chemistry and long-range structural interactions.

To address these limitations, we propose an extended framework that explicitly incorporates unit cell information and enforces periodicity in the feature space. Our contributions are as follows:

- **Cell-Level Structural Encoding:** We introduce a dedicated module that captures and encodes cell-level structural details by adding a cell node, providing a comprehensive global perspective of the crystal’s periodic architecture.
- **Periodic Feature Encoding:** We develop a specialised module that explicitly incorporates periodic boundary conditions into the feature space, ensuring that the inherent repeating patterns and symmetry of crystalline materials are fully captured, ensuring invariance with respect to the cell representation.
- **Multiscale Fusion:** We design a comprehensive fusion module that seamlessly integrates the cell-level details with atom-level representations into a cohesive framework, thereby enhancing the overall predictive performance of the network.
- **Mixture-of-Experts Integration:** We introduce a principled mixture-of-experts mechanism that jointly leverages the cell-level structural encoding, periodic feature encoding, and multiscale, learning property-dependent gates and fusion weights to combine the different experts. This integration delivers strong predictive performance while remaining interpretable, by explicitly showing how much each expert contributes to each prediction.

We validate our approach by benchmarking it on Jarvis [6], Materials Project [4], and Matbench [7] datasets. Our extensive experiments demonstrate that by explicitly encoding the unit cell and enforcing periodicity in the

feature space, our framework significantly improves predictive accuracy, paving the way for more accurate models in crystal property prediction.

## Results

To rigorously evaluate the effectiveness of our proposed PRISM model, we conduct extensive experiments using three widely-adopted crystal property benchmarks: JARVIS [6], the Materials Project [4], and MatBench [7]. These benchmarks cover a broad range of crystalline materials and property prediction tasks, spanning diverse data regimes from small-scale tasks such as *jdft2d* (636 two-dimensional crystal structures) to large-scale tasks such as *e\_form* (132,752 crystals). Further dataset details are provided in the Dataset Description Section of the Supplementary Information.

For reproducibility and fair comparison, we strictly follow the official benchmark splits when available. For JARVIS, we use the fixed 80/10/10 train/validation/ test split protocol adopted in prior work (ALIGNN/MatFormer) [35]. For the Materials Project benchmarks, we follow the same fixed splits used in MatFormer [35]. We additionally report the standard deviation across different random initialisation seeds to quantify robustness, as advocated in CartNet [30]. For MatBench, we follow the official MatBench evaluation protocol based on five-fold nested cross-validation: each outer fold uses 80% of the data for training/validation and 20% for testing; within the 80% training portion, we use 10% for validation, and we report mean and standard deviation over multiple test folds [7].

Additionally, we compare PRISM with several state-of-the-art baselines, including CGCNN [33], Schnet [29], MEGNET [4], GATGNN [15], ALIGNN [31], MatFormer [35], PotNet [13], M3GNet [5], MODNet [7], coGN [28], eComFormer [34], iComFormer [34], and CartNet [30]. Based on the results of the Table 5 from the Discussion Section, in all our experiments we used CartNet [30] as expert module. The Detailed Architecture Section from the Supplementary Information provides detailed implementation information for all the architecture used. Unless otherwise stated, we use a cutoff radius  $r_c = 5 \text{ \AA}$  for all experiments. All remaining hyperparameters are reported in the Supplementary Information. Additional details about the computational details, PRISM’s hyperparameters, and training configurations can be found in the Training Details Section from the Supplementary Information.

### Jarvis Dataset

Table 1 shows that among the tested architectures, PRISM delivers the highest performance in every task. In detail, our method improves formation energy error by about 4.36%, enhances band gap (OPT) accuracy by roughly 5.24%, and lowers total energy error by nearly 0.9% compared with the runner up. It also reduces band gap (MBJ) error by approximately 6.54% and cuts Ehull error by close to 45.2%.

Table 1: MAE results for the different tested architectures on the test split of the JARVIS dataset. The best result is in **bold** and the second-best is underlined.

Method	Form. Energy (meV/atom)	Band gap (OPT) (meV)	Total energy (meV/atom)	Band gap (MBJ) (meV)	Ehull (meV/atom)
CGCNN	63	200	78	410	170
SchNet	45	190	47	430	140
MEGNET	47	145	58	340	84
GATGNN	47	170	56	510	120
ALIGNN	33	142	37	310	76
Matformer	32.5	137	35	300	64
PotNet	29.4	127	32	270	55
eComFormer	28.4	124	32	280	<u>44</u>
iComFormer	27.2	<u>122</u>	<u>28.8</u>	<u>260</u>	47
CartNet	27.05 ± 0.07	115.31 ± 3.36	26.58 ± 0.28	253.03 ± 5.20	43.90 ± 0.36
PRISM	<b>25.87 ± 0.36</b>	<b>109.26 ± 2.54</b>	<b>26.34 ± 0.38</b>	<b>236.49 ± 5.56</b>	<b>24.07 ± 0.62</b>

## Materials Project Dataset

As shown in Table 2, among highly competitive baselines PRISM achieves the best performance for formation energy, band gap and bulk modulus. Specifically, our approach reduces the formation energy error by about 5.0% and the band gap error by nearly 5.8% compared with the second best method. In bulk modulus prediction PRISM attains the top result, while in shear modulus it shows an error that is around 2.8% higher than the best value.

Table 2: MAE results for the different tested architectures on the test split of the Materials Project dataset. Best result in **bold** and second-best underlined.

Method	Form. Energy (meV/atom)	Band Gap (meV)	Bulk Moduli (log(GPa))	Shear Moduli (log(GPa))
CGCNN	31	292	0.047	0.077
SchNet	33	345	0.066	0.099
MEGNET	30	307	0.060	0.099
GATGNN	33	280	0.045	0.075
ALIGNN	22	218	0.051	0.078
Matformer	21	211	0.043	0.073
PotNet	18.8	204	0.040	<u>0.065</u>
eComFormer	18.16	202	0.0417	0.0729
iComFormer	18.26	193	<u>0.038</u>	<b>0.0637</b>
CartNet	<u>17.47 ± 0.38</u>	<u>190.79 ± 3.14</u>	<b>0.033 ± 0.94 · 10<sup>-3</sup></b>	<b>0.0637 ± 0.0008</b>
PRISM	<b>16.59 ± 0.10</b>	<b>179.71 ± 1.58</b>	<b>0.033 ± 1.17 · 10<sup>-3</sup></b>	<u>0.0655 ± 0.0008</u>

## MatBench Dataset

Table 3 compares PRISM against state-of-the-art baselines on two MatBench tasks. On the e\_form benchmark (132 752 samples), PRISM achieves the lowest MAE and RMSE ( $15.20 \pm 0.31$ ,  $30.43 \pm 1.38$  meV/atom), improving upon the next-best models (eComFormer and iComFormer at  $16.5 \pm 0.3$  meV/atom). For jdf2d (636 samples), PRISM attains a MAE of  $38.41 \pm 12.44$  GPa and RMSE of  $97.90 \pm 38.25$  GPa, closely following iComFormer’s top MAE ( $34.8 \pm 9.9$  GPa)

and RMSE ( $96.1 \pm 46.3$  GPa) while outperforming other methods. We note that jdft2d comprises only 636 two-dimensional crystals; this low-data regime is likely not enough to achieve the mixture of experts full potential but still achieve a competitive result. The significant error deviation in the metrics suggests that additional labelled data or targeted augmentation could yield further gains. These results demonstrate PRISM’s superior accuracy in large-scale formation-energy predictions and its competitive performance under scarce-data scenarios.

Table 3: MAE and RMSE with mean and std from all the test splits from the Matbench dataset. Best result in **bold** and second-best underlined.

Method	e_form (meV)		jdft2d (GPa)	
	MAE	RMSE	MAE	RMSE
MODNet	$44.8 \pm 3.9$	$88.8 \pm 7.5$	<b><math>33.2 \pm 7.3</math></b>	<u><math>96.7 \pm 40.4</math></u>
ALIGNN	$21.5 \pm 0.5$	$55.4 \pm 5.5$	$43.4 \pm 8.9$	$117.4 \pm 42.9$
coGN	$17.0 \pm 0.3$	$48.3 \pm 5.9$	$37.2 \pm 13.7$	$101.2 \pm 55.0$
M3GNet	$19.5 \pm 0.2$	-	$50.1 \pm 11.9$	-
eComFormer	<u><math>16.5 \pm 0.3</math></u>	$45.4 \pm 4.7$	$37.8 \pm 9.0$	$102.2 \pm 46.4$
iComFormer	<u><math>16.5 \pm 0.3</math></u>	<u><math>43.8 \pm 3.7</math></u>	<u><math>34.8 \pm 9.9</math></u>	<b><math>96.1 \pm 46.3</math></b>
PRISM	<b><math>15.20 \pm 0.31</math></b>	<b><math>30.43 \pm 1.38</math></b>	$38.41 \pm 12.44$	$97.90 \pm 38.25$

## Discussion

We begin by quantifying how each specialist in PRISM contributes to accuracy and efficiency. The aim is to understand which signals, namely local atomistic geometry, feature-space similarity that respects periodicity, cross-scale aggregation, and explicit cell-scale periodicity, drive improvements on each target, and whether combining them yields complementary gains. The PRISM architecture and its expert modules are described in the Methods Section.

Table 4 presents a controlled progression from an atomistic-only baseline to the full mixture. Starting from CartNet with only an atomistic radius graph (formation energy  $27.05 \pm 0.07$  meV/atom and Ehull  $43.90 \pm 0.36$  meV/atom), adding the Similarity expert already lowers Ehull to  $34.91 \pm 1.28$  meV/atom with a moderate increase in time per epoch. Enforcing periodic invariance in feature space further improves both targets, reaching  $26.89 \pm 0.19$  meV/atom for formation energy and  $32.42 \pm 0.51$  meV/atom for Ehull at the same parameter count. The Multiscale expert on its own reduces formation-energy error to  $26.92 \pm 0.17$  meV/atom, consistent with its role in aggregating atomic information without introducing new geometric edges. Adding the Cell-Space expert on top of Multiscale strengthens formation energy to  $26.05 \pm 0.21$  meV/atom, showing the benefit of explicit lattice-scale periodic cues. The best overall performance appears when all experts are fused, with  $25.87 \pm 0.36$  meV/atom for formation energy and  $24.07 \pm 0.62$  meV/atom for Ehull. Together, these results indicate that feature-space similarity under periodic boundary conditions is especially effective

for Ehull, while Multiscale and Cell-Space contribute strongly to formation energy, and their combination is needed to reach the highest accuracy.

Table 4: Effect of expert modules on JARVIS. Checkmarks indicate which specialists are enabled (Atomistic, Similarity/Feat., Multiscale, and Cell).  $\checkmark_{PBC}$  denotes feature-space similarity made periodic-invariant via minimum-distance edges. We report mean  $\pm$  s.d. over seeds for formation energy and Ehull (meV/atom), time per epoch (s), and parameter count (M). Best and second-best per column are highlighted in **bold** and underline, respectively.

Atom.	Feat.	Multi.	Cell	Form. Energy (meV/atom)	Ehull (meV/atom)	Time/epoch (s)	#Params
$\checkmark$				$27.05 \pm 0.07$	$43.90 \pm 0.36$	<b><math>8 \pm 0.05</math></b>	<b>2.5M</b>
$\checkmark$	$\checkmark$			$27.10 \pm 0.27$	$34.91 \pm 1.28$	<u><math>16.46 \pm 0.13</math></u>	4.9M
$\checkmark$	$\checkmark_{PBC}$			$26.89 \pm 0.19$	<u><math>32.42 \pm 0.51</math></u>	$16.83 \pm 0.12$	4.9M
$\checkmark$		$\checkmark$		$26.92 \pm 0.17$	$39.96 \pm 0.37$	$19.28 \pm 0.20$	<u>4.5M</u>
$\checkmark$		$\checkmark$	$\checkmark$	<u><math>26.05 \pm 0.21</math></u>	$38.88 \pm 1.10$	$26.61 \pm 0.20$	6.6M
$\checkmark$	$\checkmark_{PBC}$	$\checkmark$	$\checkmark$	<b><math>25.87 \pm 0.36</math></b>	<b><math>24.07 \pm 0.62</math></b>	$37.71 \pm 0.29$	9M

In summary, the mixture consistently improves both targets relative to the atomistic-only baseline, with a controlled growth in parameters and time per epoch. This supports the decision to integrate short-range geometry, periodic-invariant similarity, cell-level periodicity, and cross-scale aggregation within one model.

We next test whether these gains are consistent across message-passing backbones and use this comparison to choose the backbone for the main experiments. Beyond consistency, we also test whether the improvements can be explained simply by increasing parameter count, depth, width, or radius, or whether the mixture itself contributes additional signal.

Table 5 compares vanilla eComformer, iComformer, and CartNet variants against the same backbones augmented with all PRISM experts. For eComformer, the formation-energy error decreases from 28.40 to 27.74 meV/atom and Ehull decreases from 44.00 to 31.73 meV/atom when the experts are added. For iComformer, formation energy decreases from 27.2 to 26.46 meV/atom and Ehull decreases from 47 to 36.06 meV/atom. With CartNet as the backbone, PRISM attains the strongest results overall, reaching  $25.87 \pm 0.36$  meV/atom for formation energy and  $24.07 \pm 0.62$  meV/atom for Ehull.

Importantly, the backbone comparisons show that the gains are not a by-product of simply making models larger or extending the cutoff. Wider or deeper variants (512-dimensional or 16 layers) carry more parameters but do not match the accuracy achieved by the mixture with experts. Likewise, CartNet with a larger cutoff of 10 Å has the same parameter count as the baseline, increases time per epoch substantially, and still underperforms PRISM. These controls indicate that the mixture of experts contributes complementary information that is not recovered by scaling depth, width, or radius alone.

From an efficiency perspective, PRISM with CartNet runs at  $37.71 \pm 0.29$  s/epoch with 9M parameters, which is competitive with other state-of-the-art

Table 5: Ablation study of different state-of-the-art architectures used as experts on formation energy and Ehull on the JARVIS dataset. Best is in **bold** and second-best is underlined.

Architecture	Form. energy (meV/atom)	Ehull (meV/atom)	Time/epoch (seconds)	#Params
eComformer (vanilla)	28.40	44.00	<u>23.87 ± 0.88</u>	12.4 M
eComformer (512 dim)	31.04 ± 0.49	44.99 ± 2.74	40.11 ± 0.66	16.8 M
eComformer (16 layers)	29.53 ± 1.78	42.35 ± 5.03	73.63 ± 2.20	16.6 M
+ all PRISM experts	27.74 ± 0.17	<u>31.73 ± 2.69</u>	42.65 ± 1.07	15 M
iComformer (vanilla)	27.2	47	29.90 ± 0.38	<u>5 M</u>
iComformer (512 dim)	29.25 ± 0.61	43.73 ± 1.55	62.18 ± 0.17	25 M
iComformer (16 layers)	31.96 ± 2.08	75.48 ± 8.39	72.80 ± 1.31	16.5 M
+ all PRISM experts	<u>26.46 ± 0.62</u>	36.06 ± 0.79	41.29 ± 1.32	15 M
CartNet (vanilla)	27.05 ± 0.07	43.90 ± 0.36	<b>8 ± 0.05</b>	<b>2.5 M</b>
CartNet (10 Å)	27.61 ± 2.47	32.86 ± 1.58	82.58 ± 0.16	<b>2.5 M</b>
CartNet (16 layers)	28.91 ± 0.88	34.44 ± 3.39	49.64 ± 0.95	8.8 M
CartNet (512 dim)	27.18 ± 0.22	41.39 ± 1.31	27.38 ± 0.19	9.8 M
+ all PRISM experts	<b>25.87 ± 0.36</b>	<b>24.07 ± 0.62</b>	37.71 ± 0.29	9 M

baselines such as vanilla eComformer ( $23.87 \pm 0.88$  s/epoch) and iComformer ( $29.90 \pm 0.38$  s/epoch), while delivering a considerable improvement in accuracy. When the experts are added to these backbones, the training time increases modestly. iComformer rises by about 10 s/epoch to  $41.29 \pm 1.32$  s/epoch, and eComformer by about 20 s/epoch to  $42.65 \pm 1.07$  s/epoch. For CartNet the overhead is larger and scales nearly linearly with the number of experts. CartNet uses very simple input encoders and concentrates most of the computation in the message-passing updates, which makes the vanilla model extremely fast but also means that duplicating message-passing modules across experts increases cost almost proportionally. This is why PRISM with CartNet is on the order of four times slower than vanilla CartNet, yet it remains competitive with the other baselines even without the experts. Given the consistent gains across all backbones and the favourable accuracy–efficiency balance, we adopt CartNet as the message-passing core in the remainder of the study. The pattern observed across eComformer, iComformer, and CartNet shows that the proposed mixture is largely agnostic to the choice of message-passing architecture.

Beyond aggregate accuracy, the mixture-of-experts formulation offers a direct handle on interpretability via the learned fusion weights. We therefore examine how PRISM distributes responsibility across experts and whether these distributions reflect the chemistry and physics of each target. To do so, we summarise the learned fusion weights by averaging across layers and random seeds, and then analyse the resulting allocations to the *Atomistic*, *Similarity*, *Multiscale* and *Cell* experts. Figure 1 reports the mean fusion weights for each property, separated into atom-level and cell-level contributions.

To ground the discussion, we briefly recall what each target represents. The Formation Energy is the energy per atom of a compound relative to the




Figure 1: Fusion-weight analysis. Average fusion weights over layers and seeds on the JARVIS test set. Each subpanel shows atom-level (*Multiscale*, *Atomistic*, *Similarity*) and cell-level (*Multiscale*, *Cell*) contributions for a single target property.

isolated neutral atoms; it is a crucial parameter for quantifying thermodynamic stability. Total Energy is the ground-state energy of the relaxed cell. Band Gap (OPT) is the Kohn–Sham gap computed with the OPT (optB88-vdW) functional [12]. Band Gap (MBJ) applies a modified exchange functional that includes an additional term proposed by Becke–Johnson [26], which mitigates the underestimation of gaps by Generalised-Gradient Approximation (GGA) [23] functionals. Finally, Ehull (energy above the convex hull) measures the distance of a material from the stability hull defined by all competing phases of the same composition; stable entries have Ehull=0.

Two robust patterns are evident in Figure 1. First, for properties with a strong global or band-structure character the cell-level pathway dominates. Band Gap (OPT) and Ehull place most weight on the *Cell* expert, and Band Gap (MBJ) still prefers *Cell* over cell-level *Multiscale*. Second, within atom-level fusion, energy-like targets favour *Atomistic* and *Multiscale* approaches, whereas the *Similarity* expert becomes more prominent for electronic and stability descriptors.

Although both targets are reported per atom in JARVIS, the model allocates weights differently to formation energy and total energy. For formation energy, the atom-level fusion is dominated by *Atomistic* and *Multiscale* share (see Figure 1), because formation energy is built as a difference between the compound and its elemental references; this reference normalisation cancels much of the atomic contributions, leaving local bonding, coordination, hybridisation and short interatomic distances as the primary determinants. By contrast, total energy per atom is an absolute quantity that still aggregates electrostatic and dispersion effects into the energy density experienced by each site; dividing by atom count removes trivial size scaling but not the underlying long-range couplings. Consequently, PRISM assigns a larger share to the cell-level *Multiscale* pathway for total energy, where the superatom aggregator pools information across multiple neighbourhood radii to capture these extended interactions.

The behaviour of the band gap in Figure 1 further illustrates the role of scale. With OPT, the gap is primarily governed by global band-structure features, such as periodicity, crystal symmetry, and lattice parameters; therefore, the model assigns overwhelming weight to the *Cell* expert. Introducing MBJ

changes the picture. MBJ raises conduction bands by including an additional kinetic-energy-density-dependent potential in the exchange term of the functional. This correction improves both periodic gap properties and the local chemical description (orbital energies, p–d hybridisation, and crystal-field splittings). This increases the relative contribution of atom-level paths, notably *Similarity* and *Multiscale*, compared with OPT, even though *Cell* remains the single most significant contributor. Importantly, this shift is not primarily due to van der Waals dispersion, since OPT already includes a dispersion correction (optB88-vdW). The additional multiscale weight with MBJ reflects the need to propagate the proper local electronic structure to a long-range level and to capture subtle short-to-mid-range electronic effects beyond a pure cell-level summary.

The weights for Ehull shown in Figure 1 are also physically reasonable. Because Ehull compares a structure’s formation energy with all competing phases of the same composition, it is strongly influenced by stoichiometry, density and lattice metrics. This explains the pronounced weight on the *Cell* expert. At the same time, near-degenerate polymorphs and ordering motifs can shift a compound on or off the hull; capturing these subtle differences benefits from a *Multiscale* expert, together with a smaller *Atomistic* and *Similarity* contribution.

Taken together, the property-aware allocation of weights accords with chemical and physical intuition. The fusion layer therefore improves interpretability in two complementary ways. It reveals which structural scale (atomic, similarity, cell or multiscale) is most informative for each property, and it enables the prediction to be decomposed into expert-wise contributions. As an additional quantitative check, correlations between mean weights and interpretable descriptors, such as local coordination for *Atomistic*, dispersion and sublattice identities for *Similarity*, lattice translation for *Cell*, and long-range effects for *Multiscale*, support a mechanistic interpretation rather than using a non-physically informed model.

Overall, these results support the central design choice of PRISM. The model explicitly encodes unit-cell periodicity in both geometry and feature space, whilst fusing local and global information through a multiscale mixture-of-experts design (*Atomistic*, *Similarity*, *Multiscale*, and *Cell*). Across JARVIS, Materials Project, and MatBench, this combination yields consistent improvements over competitive baselines and is not replicated by straightforward scaling of depth, width, or cutoff radius. At the same time, the analysis highlights several current limitations that delineate clear avenues for extension. First, the present study focuses exclusively on periodic crystal structures, which limits direct transfer to non-periodic settings such as molecular property prediction. Second, the current Similarity expert encodes a single unit cell; as a result, infinite lattice repetitions are not represented in the feature graph, and crystals in which feature-level similarity appears primarily across periodic replicas may not fully benefit from this pathway. Third, whilst the Multiscale expert captures repetition frequencies at the cell level via the superatom mechanism, it does not yet incorporate broader coarse-grained experts that may better represent hierarchical materials or multiscale phenomena beyond the unit-cell abstraction. Finally, our approach is centred on a chosen unit-cell representation; although alternative representations

such as supercells are compatible in principle, exploiting them would require retraining.

These considerations suggest several practical directions for future work. Extending PRISM to molecular contexts would require adapting the graph construction and fusion mechanisms to account for the absence of translational symmetry and for the different balance of local versus global signals. Within periodic systems, incorporating larger supercells or explicitly modelling lattice repetitions within the feature-space graph could strengthen the Similarity pathway and reduce failure cases where intra-cell similarity is sparse. More broadly, introducing additional coarse-grained experts, operating at intermediate structural resolutions between atomistic sites and the superatom, may further enhance PRISM’s ability to model complex hierarchical materials and long-range couplings. In this sense, PRISM provides both an effective predictive model and a flexible, interpretable scaffold for systematically integrating new periodic and multiscale inductive biases.

## Methods

PRISM (Periodic Representation with multIscale and Similarity graph Modelling) is a multigraph-based neural network architecture tailored for crystalline materials. It employs a collection of expert modules, each specialised in encoding complementary structural and chemical features of periodic systems. Drawing inspiration from ensemble learning [9, 18], every expert operates on a distinct graph topology to capture interactions at different spatial scales, thereby enabling PRISM to jointly model both local atomic environments and global lattice periodicity.

Our PRISM framework constructs a dual-scale, multi-graph representation that is iteratively refined through an ensemble of experts. As illustrated in Figure 2a, we maintain both atomistic embeddings  $\mathbf{h}_i$  (highlighted in blue) and a global superatom embedding  $\mathbf{h}_s$  (highlighted in yellow), each connected via their respective edge sets. Each PRISM layer aggregates and fuses information through four specialised experts: the *Cell Expert* encodes lattice periodicity by constructing a superatom graph that captures global cell-level repetitions; the *Multiscale Expert* mediates bidirectional information flow between atomistic and superatom embeddings to integrate local and global context; the *Atomistic Expert* models short-range interatomic interactions through a radial cutoff graph under periodic boundary conditions; and the *Similarity Expert* connects atoms with similar electronic environments in feature space under periodic boundary conditions, thereby enabling long-range relational modeling while preserving periodic invariance. All embeddings are updated in parallel within each layer, and the final atomistic embeddings are used to predict the target property (see Figure 2b).

To formalise the input to our PRISM architecture, we represent each crystal as a graph over its atomic sites and periodic lattice vectors, defined as follows. Given a crystal structure, we define the set of atoms, represented as nodes  $\mathcal{V}$ , which

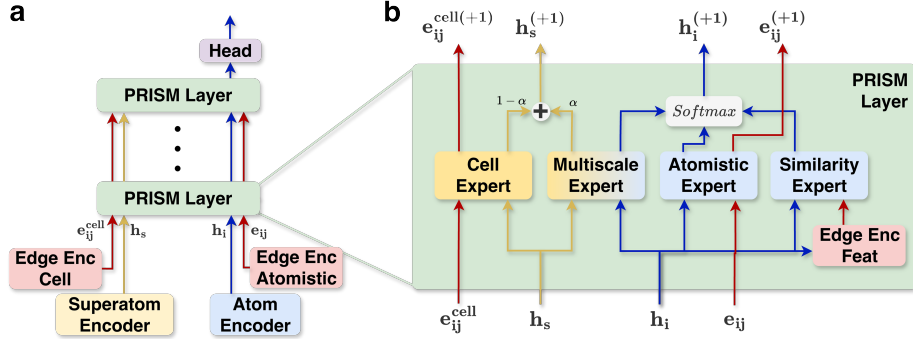


Figure 2: **a** Overview of the PRISM architecture. Atomic and superatom embeddings are initialised via dedicated encoders, and each PRISM layer updates them until the final representation that is used to predict the property. **b** The PRISM layer architecture. Each of the representations are aggregated and fused using four expert modules: Cell (global lattice periodicity), Multiscale (atom-superatom bidirectional interactions), Atomistic (radius graph with periodic boundaries), and Similarity (feature-space graph with periodic boundaries). In all graph-based experts,  $e_{ij}$  is the learned edge embedding for the connection  $(i, j)$ , computed by the edge encoder for that specific graph and used in the expert.

represents the  $N$  atoms within a unit cell. Each node  $i \in \mathcal{V}$  is characterised by its Cartesian coordinates  $\mathbf{r}_i \in \mathbb{R}^3$  and an initial latent feature embedding  $\mathbf{h}_i^{(0)} \in \mathbb{R}^{dim}$  derived from its atomic number, where  $dim$  denotes the dimensionality of the latent feature vector. The periodic boundaries of the crystal are defined by the lattice matrix  $\mathbf{L} = [\mathbf{l}_1, \mathbf{l}_2, \mathbf{l}_3] \in \mathbb{R}^{3 \times 3}$ , with lattice vectors  $\mathbf{l}_j \in \mathbb{R}^3$  specifying the geometry of the unit cell. Based on this formalism, we propose four distinct expert modules, each designed to encode the crystal’s complementary structural and chemical characteristics by operating over specialised graph topologies defined on the atomic set  $\mathcal{V}$ .

### Atomistic Expert

To accurately model the fundamental chemical bonds and local physical forces that govern crystalline stability and emergent properties, it is essential to capture interactions at the atomistic scale. This expert captures short-range atomic interactions by constructing a graph  $\mathcal{G}_{atomistic}$  connecting atoms  $i$  and  $j$  when their minimum-image distance under Periodic Boundary Conditions satisfies:

$$\|\mathbf{r}_i - \mathbf{r}_j + \mathbf{L}\mathbf{n}_{lmn}\| < r_c, \quad (1)$$

where  $\mathbf{n}_{lmn} \in \mathbb{Z}^3$  is the lattice shift vector of the image.




Figure 3: Atomistic expert graph  $\mathcal{G}_{\text{atomistic}}$  under periodic boundary conditions. A radius-based neighbourhood with cutoff  $r_c$  (green spheres) is built around each reference atom (blue), and edges (red) connect atoms within  $r_c$  while accounting for periodic images. **a** Three reference atoms are highlighted, and a radius graph is created around each of them. The resulting lateral links connect adjacent layers of the material, allowing message passing to propagate across layers and correctly encode the geometry. **b** In this material, the cutoff does not reach lateral neighbours, so edges form mainly along the vertical direction. Message passing is then propagated only vertically, and the model cannot capture the lateral geometry; stacking additional message-passing layers does not restore the missing lateral pathways.

Although a radius-based construction is one of the most common choices when building graphs for GNNs for materials, it has an important limitation about propagating the information, since it assumes that all the regions are connected. If two regions are not connected through the graph, messages cannot propagate across these gaps, and even stacking multiple message-passing layers that follow this connectivity does not restore the missing pathways. Increasing  $r_c$  can mitigate this by adding longer-range edges. This, however, weakens locality, often degrades overall performance, and has an expensive computational cost, as discussed in the Discussion Section. As illustrated in Figure 3, when the layers of the material are close enough to have a connection inside the  $r_c$ , the information is propagated through all the system. When the material layers are more separated, the information is only propagated vertically, causing that the message passing is not encoding the structure correctly.

## Similarity Expert

To address the propagation bottlenecks that the atomistic radius graph  $\mathcal{G}_{\text{atomistic}}$  can induce, we introduce the *Similarity* expert. This expert propagates semantically relevant chemical and structural cues across distant yet similar atoms, capturing global correlations that spatial proximity alone cannot reveal. It explicitly models similarity by constructing a feature-space graph,  $\mathcal{G}_{\text{feat}}$ , where atoms are connected if the Euclidean distance between their learned feature embeddings is below a feature cutoff  $r_f$ . Because our message-passing modules encode geometry on edges, we then instantiate geometric edges between these

feature-connected nodes using their minimum-periodic displacement. Linking atoms with similar chemistry or local environments enables global information flow and supports long-range dependencies beyond radius-only neighbourhoods.

For each edge connecting two atoms  $i$  and  $j$  based on feature similarity, we compute an associated geometric edge attribute that represents the minimum periodic distance between these atoms. This step ensures that the graph representation is invariant to equivalent unit-cell configurations, an essential requirement given that crystalline structures can have multiple equivalent representations due to symmetry.

Specifically, the calculation involves converting atomic coordinates into fractional coordinates to account explicitly for the periodic boundary conditions imposed by the lattice vectors:

$$\mathbf{d}_{\text{frac}} = \mathbf{L}^{-1}(\mathbf{r}_i - \mathbf{r}_j), \quad (2)$$

where  $\mathbf{d}_{\text{frac}}$  denotes the fractional distance vector between atoms  $i$  and  $j$ . To identify the minimal periodic distance, we shift this fractional distance into the fundamental periodic cell by applying:

$$\mathbf{d}_{\text{frac}}^{\text{PBC}} = \mathbf{d}_{\text{frac}} - \lfloor \mathbf{d}_{\text{frac}} + 0.5 \rfloor, \quad (3)$$

where  $\lfloor \cdot \rfloor$  denotes the floor operation, effectively mapping the fractional coordinates into the range  $[-0.5, 0.5]$ . Finally, we convert the fractional minimal vector back into Cartesian coordinates:

$$\mathbf{d}_{\text{min}} = \mathbf{L} \mathbf{d}_{\text{frac}}^{\text{PBC}}. \quad (4)$$

This approach guarantees that the resulting geometric edge attributes consistently represent minimal distances under periodic boundary conditions and are invariant to different but equivalent unit-cell representations. The detailed construction of  $\mathcal{G}_{\text{feat}}$  is visually illustrated in Figure 4.

Due to the dynamic nature of the feature-space graph, where different edges may emerge or disappear between message-passing layers, edge features cannot be precomputed or reused. Thus, at each layer, we independently reconstruct the edges and re-encode their geometric attributes through a dedicated encoder named *Feature-Edge Encoder*. This contrasts with the static edge encoders employed in other experts, ensuring accurate and consistent edge feature representations across the evolving graph topology.

The main limitation of this expert is that  $\mathcal{G}_{\text{feat}}$  is built with atoms restricted to a single unit cell. Some crystals exhibit no feature-level similarity within the cell but do exhibit similarity across periodic replicas. In such cases,  $\mathcal{G}_{\text{feat}}$  can become a disconnected graph and fail to improve message-passing propagation, as can be seen in Figure 4c.

## Cell-Space Expert

To address the limitations introduced by the Atomistic and Similarity experts in the message passing, such as the examples in Figure 3b and Figure 4c, we propose

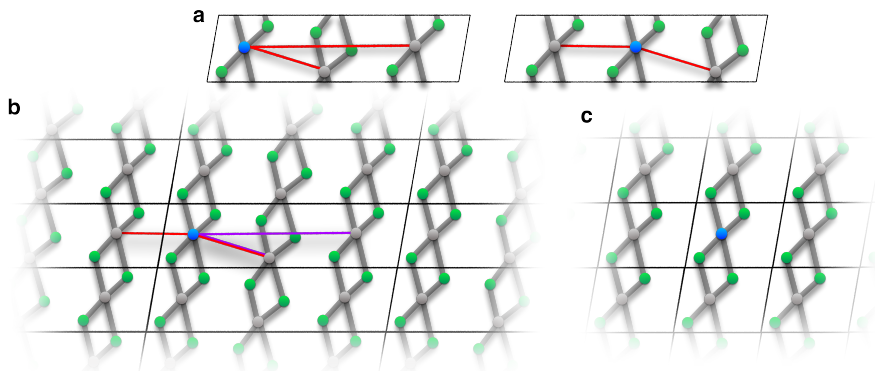


Figure 4: Illustration of the feature-space neighbourhood  $\mathcal{G}_{\text{feat}}$  under periodic boundary conditions. The neighbourhood is built around the blue reference atom and edges (red) link feature-similar atoms within the shown cell. **a** Two equivalent unit-cell choices lead to different similarity edges if periodicity is ignored, producing inconsistent graphs across equivalent cells. **b** Proposed periodic-invariant construction. Candidate neighbours (magenta) are first identified in feature space; for each candidate, we then select the minimal periodic image and add the corresponding edge (red), yielding consistent and invariant graphs. **c** Corner case in which no feature-similar atoms lie inside the reference unit cell. Similar atoms exist only as periodic replicas, resulting in an unconnected  $\mathcal{G}_{\text{feat}}$  graph and a vertically connected  $\mathcal{G}_{\text{atomistic}}$ , which leads to poor information flow for this structure. Nevertheless, the Cell-Space Expert is able to propagate the information in such structures since the  $\mathcal{G}_{\text{cell}}$  graph is connected to the image replicas.

the Cell-Space expert. It captures long-range cell-to-cell interactions, such as delocalised electronic coupling, collective lattice correlations, and boundary-driven surface effects, which atom-level graphs cannot represent. The design encodes global periodic repetitions by introducing a single superatom node  $s$ . We initialise the embedding of  $h_s$  with an attention-weighted sum of the atomic embeddings within the unit cell. See the Detailed Architecture section in the Supplementary Information for implementation details.

We construct a radius-based graph  $\mathcal{G}_{\text{cell}}$  around the superatom node with a cutoff radius  $R_c$ , significantly larger than the atomistic radius  $r_c$  ( $R_c \gg r_c$ ). Edges in this graph are initialised based on the geometric distances within this extended radius  $R_c$  and are processed using an independent edge encoder. This edge encoder is analogous to the one used in the atomistic expert but tailored specifically to handle the larger cutoff radius. The significantly increased radius ensures that the graph explicitly captures lattice repetitions at the superatom scale. Figure 5 illustrates the detailed construction of the Cell-Space Expert graph  $\mathcal{G}_{\text{cell}}$ . As discussed in Matformer [35] and iComformer [34], the distance

between the repetitions of the same atom is invariant to equivalent unit-cell transformations, and the rotational symmetries can be overcome thanks to defining an invariant cell [34], equivariant message passing updates [34] or via rotational data augmentation [30].




Figure 5: Cell-Space expert graph construction. The superatom (central orange node) connects to periodic replicas within a large cutoff  $R_c$  (green circle), encoding lattice-scale periodicity through edges (red). **a–b** Two equivalent unit-cell choices produce the same set of replica connections because cell transformations preserve the relative distances and directions between the superatom and its periodic images. The only change comes from global rotations, which can be handled by defining an invariant cell [34], using equivariant message-passing updates [34], or applying rotational data augmentation [30].

The Cell-Space expert arises from the way electronic coupling between neighbouring cells varies with cell size. For small unit cells, atoms in adjacent images lie within the immediate chemical environment and therefore produce dense connectivity around the superatom. In contrast, for larger unit cells, significant cross-cell interactions occur only for atoms near the boundaries, and the Atomistic Expert and the Similarity Expert already capture these boundary effects. In other words, for small unit cells, this expert is highly influential and forms many connections, whereas for larger cells, its reduced connectivity is compensated by more connections formed by the Atomistic and Similarity experts.

Other methods, such as Matformer [35], also aim to overcome the limitations of relying only on the Atomistic experts. Matformer adds self-loops to each atom to encode distance features derived from the unique distances to each of the image repetitions. The main limitation is that, even though each atom is connected to its periodic image, the method does not capture intra-cell symmetries as effectively as the Similarity expert. In addition, every atom stores the same self-loop distances information, which introduces redundancy that scales with the number of atoms. Collapsing this shared information into a single superatom removes the redundancy and makes the approach more efficient while preserving the relevant periodic information. Our Cell-Space expert also scales negatively with the number of atoms in the unit cell. As the number of atoms increases, the typical cell size grows, periodic images move farther apart for a fixed cutoff  $R_c$ , and the superatom graph  $\mathcal{G}_{\text{cell}}$  contains fewer connections. Further comparisons on scalability and computational cost are provided in the

Supplementary Information. The remaining design question for this expert is how to couple the superatom to the underlying atomic representations so that equivalent unit-cell transformations produce identical outputs.

Finally, to capture interactions across distinct structural scales, we introduce the Multiscale expert, which explicitly connects the global superatom representation with atomic-level embeddings without breaking the invariance between equivalent unit-cell representations. The primary objective of this expert is to enable an efficient bidirectional information flow, allowing local atomic features to aggregate into a global representation and simultaneously distributing global contextual information back to the atomic nodes. Such coupling across scales is inspired by multiscale modeling frameworks commonly employed in computational physics [16, 2], where macroscale properties emerge from microscale interactions, and global constraints influence local states. Formally, we define the bipartite multiscale graph  $\mathcal{G}_{\text{multiscale}}$  as follows:

$$\mathcal{G}_{\text{multiscale}} = \{(s, i) \mid s \leftrightarrow i, \forall i \in \mathcal{V}\}, \quad (5)$$

where the superatom node  $s$  is bidirectionally connected to every atom  $i$  within the set of atomic nodes  $\mathcal{V}$  representing the unit cell. The connectivity ensures that all atomic nodes share a direct path to exchange information with the superatom node, enforcing global coherence across scales. Figure 6 visually illustrates the construction and connectivity of the multiscale graph  $\mathcal{G}_{\text{multiscale}}$ . The message-passing operation between scales can be mathematically represented as follows:

$$\mathbf{h}_s^{(l+1)} = \mathbf{h}_s^{(l)} + \phi_{\text{multiscale}} \left( \{\mathbf{h}_i^{(l)}\}_{i \in \mathcal{V}} \right). \quad (6)$$

$$\mathbf{h}_i^{(l+1)} = \mathbf{h}_i^{(l)} + \phi_{\text{multiscale}} \left( \mathbf{h}_s^{(l+1)} \right). \quad (7)$$

Employing the same graph neural network transformation  $\phi_{\text{multiscale}}$  in both aggregation and distribution steps ensures consistency in multiscale information exchange. The Multiscale expert is deliberately geometry-agnostic and operates only on the  $N$  atomic embeddings from the unit cell. We modify the cross-scale message passing to exclude the edge features, which carry the geometric information.  $\phi_{\text{multiscale}}$  consumes only feature embeddings, bipartite edges carry no distances, angles, or relative-position encodings, and updates depend solely on  $\{\mathbf{h}_i\}$  and  $\mathbf{h}_s$ . As illustrated in Figure 6, each equivalent unit-cell transformation contains the same  $N$  atoms. Their coordinates may vary under alternative unit-cell choices, yet the set of  $N$  atomic embeddings remains the same.

Because the superatom has no defined spatial position and messages do not use geometric features, the construction preserves invariance under equivalent unit-cell transformations. Further details on these modifications are provided in the Detailed Architecture section of the Supplementary Information.




Figure 6: Multiscale expert graph  $\mathcal{G}_{\text{multiscale}}$ . The superatom node (yellow) forms a bipartite graph to all atoms in the unit cell, with edges shown in red. In this expert, message passing is feature-only: no distances, directions, or geometric information are used. **a** and **b** show two equivalent unit-cell representations of the same material. Because the connectivity depends only on the set of atomic embeddings and not on their coordinates, both representations yield the same graph and updates, making this expert invariant to unit-cell transformations.

## Fusion of Experts

To integrate complementary information captured by the different expert modules, we introduce a learned fusion mechanism at each aggregation step. The fusion strategy separately handles the superatom and atomic node embeddings, reflecting their distinct roles within the model.

For the superatom representation, we employ a gating mechanism controlled by a trainable parameter  $\alpha$ . This parameter, modulated through the sigmoid function  $\sigma(\cdot)$ , dynamically balances the contributions from the Cell-Space and Multiscale experts:

$$\mathbf{h}_s^{(l+1)} = \sigma(\alpha)\mathbf{h}_s^{(\text{cell})} + (1 - \sigma(\alpha))\mathbf{h}_s^{(\text{multiscale})}, \quad (8)$$

where  $\mathbf{h}_s^{(\text{cell})}$  is the output superatom embedding from the Cell expert,  $\mathbf{h}_s^{(\text{multiscale})}$  is the output superatom embedding from the Multiscale expert, and  $\sigma$  denotes the sigmoid activation function, allowing the model to adaptively favor either the global periodic information from the Cell-Space expert or the hierarchical interactions from the Multiscale expert, depending on the context.

For atomic node representations, we employ a separate fusion approach. Here, we define learnable parameters  $\beta'$ ,  $\gamma'$ , and  $\delta'$ , corresponding to each expert module (Atomistic, Similarity, and Multiscale). These parameters are normalised via a softmax operation, producing weighting factors that sum to one, ensuring a meaningful convex combination of expert outputs:

$$[\beta, \gamma, \delta] = \text{Softmax}([\beta', \gamma', \delta']), \quad (9)$$

$$\mathbf{h}_i^{(l+1)} = \beta\mathbf{h}_i^{(\text{atomistic})} + \gamma\mathbf{h}_i^{(\text{feat})} + \delta\mathbf{h}_i^{(\text{multiscale})}. \quad (10)$$

where  $\mathbf{h}_i^{(\text{atomistic})}$  is the output of the  $i$ -th atom embedding from the Atomistic expert,  $\mathbf{h}_i^{(\text{feat})}$  is the output of the  $i$ -th atom embedding from the Similarity

expert, and  $\mathbf{h}_i^{(\text{multiscale})}$  is the output of the  $i$ -th atom embedding from the Multiscale expert.

This weighted fusion allows each atomic embedding to optimally integrate geometric, chemical, and scale-dependent information provided by the distinct expert modules, ensuring the representation remains flexible and expressive across diverse crystal structures.

## Data availability

The results reported in this article are based exclusively on publicly available benchmark datasets for crystal property prediction: JARVIS [6], the Materials Project [4], and MatBench [7]. For reproducibility, we use the official benchmark splits whenever available. In particular, for JARVIS we follow the fixed 80/10/10 train/validation/test split protocol adopted in prior work (ALIGNN/MatFormer) [35], and for Materials Project benchmarks we follow the same fixed splits used in MatFormer [35]. For MatBench, we follow the official evaluation protocol based on five-fold nested cross-validation [7].

All scripts required to download the raw data from the original sources, perform preprocessing, and reproduce the experiments in this article are available at <https://github.com/imatge-upc/PRISM>.

## Code availability

The code to recreate the results of the paper can be found at: <https://github.com/imatge-upc/PRISM>

## Acknowledgements

This work was supported by the Spanish Research Agency (AEI) under project PID2024-161868OB-I00, PID2021-122464NB-I00, TED2021-129593B-I00, CNS2023-144561, PID2024-155562NB-I00 and Maria de Maeztu CEX2021-001202-M. E. R. also acknowledges the Generalitat de Catalunya for an ICREA Academia grant.

## Author contributions statement

**À. S.:** Conceptualisation, Methodology, Software, Investigation, Validation, Writing - review & editing. **A. M-M.:** Conceptualisation, Supervision, Resources, Writing - review & editing. **J. C.:** Conceptualisation, Writing - review & editing. **D. A.:** Conceptualisation, Supervision, Resources, Writing - review & editing. **S. G-C.:** Conceptualisation, Supervision, Resources, Writing - review & editing, Funding acquisition. **E. R.:** Conceptualisation, Supervision, Resources, Writing - review & editing, Funding acquisition. **J. R-H.:** Conceptualisation, Supervision, Resources, Writing - review & editing, Funding acquisition.

## Conflicts of interest

The authors declare no competing financial or non-financial interests.

## Supplementary Information

The supplementary file (PDF) documents the complete PRISM architecture and reproducibility materials: a detailed description of the model components (Superatom encoder; baseline modifications on CartNet; atom and edge encoders; message-passing gates, envelope, and equations), the prediction head, and a scalability/complexity analysis with illustrative schematics, plus a comparison table of per-atom message-passing complexity and parameter counts. It further includes concise dataset summaries for JARVIS 3D DFT (v2021.8.18), Materials Project (2018.6.1), and MatBench tasks (e\_form, jdft2d); training setup (hardware and library versions), optimisation choices, and complete hyperparameter tables for each benchmark. All additional figures, tables, and references cited in the SI are provided within the PDF.

## References

- [1] J. G. Bednorz and K. A. Müller. Possible high  $t_c$  superconductivity in the  $\text{Ba-La-Cu-O}$  system. *Zeitschrift für Physik B Condensed Matter*, 64(2):189–193, 1986.
- [2] Matteo Capone, Marco Romanelli, Davide Castaldo, Giovanni Parolin, Alessandro Bello, Gabriel Gil, and Mirko Vanzan. A vision for the future of multiscale modeling. *ACS Physical Chemistry Au*, 4(3):202–225, 2024.
- [3] G. Ceder, Y.-M. Chiang, D. R. Sadoway, M. K. Aydinol, Y.-I. Jang, and B. Huang. Identification of cathode materials for lithium batteries guided by first-principles calculations. *Nature*, 392(6677):694–696, 1998.
- [4] Chi Chen, Weike Ye, Yunxing Zuo, Chen Zheng, and Shyue Ping Ong. Graph networks as a universal machine learning framework for molecules and crystals. *Chemistry of Materials*, 31(9):3564–3572, 2019.
- [5] Chi Chen, Weike Ye, Yunxing Zuo, Chen Zheng, and Shyue Ping Ong. Graph networks as a universal machine learning framework for molecules and crystals. *Nature Machine Intelligence*, 4(9):761–770, 2022.
- [6] Kamal Choudhary, Kevin F. Garrity, Andrew C. E. Reid, Brian DeCost, Adam J. Bicchi, Angela R. Hight Walker, Zachary Trautt, Jason Hatrick-Simpers, A. Gilad Kusne, Andrea Centrone, Albert Davydov, Jie Jiang, Ruth Pachter, Gowoon Cheon, Evan Reed, Ankit Agrawal, Xiaofeng Qian, Vinit Sharma, Houlong Zhuang, et al. The joint automated repository for various integrated simulations (jarvis) for data-driven materials design. *npj Computational Materials*, 6(1):173, 2020.

- [7] Alexander Dunn, Qi Wang, Alex Ganose, Daniel Dopp, and Anubhav Jain. Benchmarking materials property prediction methods: the matbench test set and automatminer reference algorithm. *npj Computational Materials*, 6(1):138, 2020.
- [8] Thorben Frank, Oliver Unke, and Klaus-Robert Müller. So3krates: Equivariant attention for interactions on arbitrary length-scales in molecular systems. *Advances in Neural Information Processing Systems*, 35:29400–29413, 2022.
- [9] M.A. Ganaie, Minghui Hu, A.K. Malik, M. Tanveer, and P.N. Suganthan. Ensemble deep learning: A review. *Engineering Applications of Artificial Intelligence*, 115:105151, 2022.
- [10] Martin A. Green, Anita Ho-Baillie, and Henry J. Snaith. The emergence of perovskite solar cells. *Nature Photonics*, 8(7):506–514, 2014.
- [11] Geoffroy Hautier, Anubhav Jain, and Shyue Ong. From the computer to the laboratory: Materials discovery and design using first-principles calculations. *Journal of Materials Science*, 47, 05 2012.
- [12] Ji ři Klimeř, David R. Bowler, and Angelos Michaelides. Van der waals density functionals applied to solids. *Phys. Rev. B*, 83:195131, May 2011.
- [13] Yuchao Lin, Keqiang Yan, Youzhi Luo, Yi Liu, Xiaoning Qian, and Shuiwang Ji. Efficient approximations of complete interatomic potentials for crystal property prediction. In *International Conference on Machine Learning*, pages 21260–21287. PMLR, 2023.
- [14] Yahui Liu, Bin Tian, Yisheng Lv, Lingxi Li, and Fei-Yue Wang. Point cloud classification using content-based transformer via clustering in feature space. *IEEE/CAA Journal of Automatica Sinica*, 11(JAS-2022-1385):231, 2024.
- [15] Steph-Yves Louis, Yong Zhao, Alireza Nasiri, Xiran Wang, Yuqi Song, Fei Liu, and Jianjun Hu. Graph convolutional neural networks with global attention for improved materials property prediction. *Phys. Chem. Chem. Phys.*, 22:18141–18148, 2020.
- [16] Mitchell Luskin and Christoph Ortner. Atomistic-to-continuum coupling. *Acta Numerica*, 22:397–508, 2013.
- [17] K. Mizushima, P.C. Jones, P.J. Wiseman, and J.B. Goodenough. Lixcoo2 (0<x<-1): A new cathode material for batteries of high energy density. *Materials Research Bulletin*, 15(6):783–789, 1980.
- [18] Ammar Mohammed and Rania Kora. A comprehensive review on ensemble deep learning: Opportunities and challenges. *Journal of King Saud University - Computer and Information Sciences*, 35(2):757–774, 2023.
- [19] Albert Mosella-Montoro and Javier Ruiz-Hidalgo. Residual attention graph convolutional network for geometric 3d scene classification. In *IEEE Conference on Computer Vision Workshop (ICCVW)*, 2019.

- [20] Albert Mosella-Montoro and Javier Ruiz-Hidalgo. 2d–3d geometric fusion network using multi-neighbourhood graph convolution for rgb-d indoor scene classification. *Information Fusion*, 76:46–54, 2021.
- [21] J. K. Nørskov, T. Bligaard, J. Rossmeisl, and C. H. Christensen. Towards the computational design of solid catalysts. *Nature Chemistry*, 1(1):37–46, 2009.
- [22] Jinyoung Park, Sanghyeok Lee, Sihyeon Kim, Yunyang Xiong, and Hyunwoo J. Kim. Self-positioning point-based transformer for point cloud understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21814–21823, June 2023.
- [23] John P. Perdew, Kieron Burke, and Yue Wang. Generalized gradient approximation for the exchange-correlation hole of a many-electron system. *Phys. Rev. B*, 54:16533–16539, Dec 1996.
- [24] John P. Perdew and Karla Schmidt. Jacob’s ladder of density functional approximations for the exchange-correlation energy. *AIP Conference Proceedings*, 577(1):1–20, 07 2001.
- [25] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5099–5108, 2017.
- [26] DP Rai, MP Ghimire, and RK Thapa. A dft study of bex (x= s, se, te) semiconductor: modified becke johnson (mbj) potential. *Semiconductors*, 48:1411–1422, 2014.
- [27] Rampi Ramprasad, Rohit Batra, Ghanshyam Pilania, Arun Mannodi-Kanakkithodi, and Chiho Kim. Machine learning in materials informatics: recent applications and prospects. *npj Computational Materials*, 3(1):54, 2017.
- [28] Robin Ruff, Patrick Reiser, Jan Stühmer, and Pascal Friederich. Connectivity optimized nested line graph networks for crystal structures. *Digital Discovery*, 3:594–601, 2024.
- [29] Kristof Schütt, Pieter-Jan Kindermans, Huziel Enoc Saucedo Felix, Stefan Chmiela, Alexandre Tkatchenko, and Klaus-Robert Müller. Schnet: A continuous-filter convolutional neural network for modeling quantum interactions. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

- [30] Àlex Solé, Albert Mosella-Montoro, Joan Cardona, Silvia Gómez-Coca, Daniel Aravena, Eliseo Ruiz, and Javier Ruiz-Hidalgo. A cartesian encoding graph neural network for crystal structure property prediction: application to thermal ellipsoid estimation. *Digital Discovery*, 4(3):694–710, 2025.
- [31] Kevin Tran, Zhenwei Yao, Simon Batzner, Jiahua Sun, Mordechai Kornbluth, Agha Kusne, Alexander Shapeev, Robert A DiStasio Jr, Nicola Marzari, et al. Benchmarking graph neural networks for materials chemistry. *npj Computational Materials*, 8(1):96, 2022.
- [32] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic graph cnn for learning on point clouds. *ACM Trans. Graph.*, 38(5), October 2019.
- [33] Tian Xie and Jeffrey C. Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.*, 120:145301, Apr 2018.
- [34] Keqiang Yan, Cong Fu, Xiaofeng Qian, Xiaoning Qian, and Shuiwang Ji. Complete and efficient graph transformers for crystal material property prediction. In *International Conference on Learning Representations*, 2024.
- [35] Keqiang Yan, Yi Liu, Yuchao Lin, and Shuiwang Ji. Periodic graph transformers for crystal material property prediction. *Advances in Neural Information Processing Systems*, 35:15066–15080, 2022.