LEARNING K-RESOLVED ELECTRONIC STRUCTURE VIA SOFT ENERGY OCCUPANCY PREDICTION

Dongik Park[1*] **Kunmin Jang**[1*] **Jaewon Bae**[1] **Dongin Kim**[2] **Chanyoung Park**[3†]

[1]NanoForge AI, Seoul, Republic of Korea
[2]Imperial College London, London, United Kingdom
[3]Korea Advanced Institute of Science and Technology, Daejeon, Republic of Korea

{dongik.park, alexjang, jaewon.bae}@nanoforgeai.com
dongin.kim20@imperial.ac.uk cy.park@kaist.ac.kr

[*]Equal contribution [†]Corresponding author

ABSTRACT

Predicting electronic structure from crystal geometry is essential for computational materials discovery, as it determines key physical quantities such as band gaps, DOS, and energy isosurfaces. While per-band prediction has been explored, it requires fixing the number of bands or indexing each band across k-points, limiting generality across materials. Predicting k-resolved electronic structure avoids these constraints; we propose kRESForge, which predicts energy bin occupancy at each k-point. Given a crystal structure and a query k-point, the model predicts a probability distribution over 256 energy bins spanning ± 10 eV from the Fermi level, naturally encoding prediction confidence. Band structure visualization follows directly from k-path queries, and downstream physical quantities such as band gaps, DOS, and energy isosurfaces can be derived through k-space aggregation without additional training. On 28,517 non-magnetic materials from the Materials Project, kRESForge achieves a band gap MAE of 0.39 eV, 90% metal/nonmetal classification accuracy, and DOS MAE of 2.64 states/eV.

1 INTRODUCTION

Electronic structure underlies a material's transport, optical, and thermal properties (Martin, 2020; Ashcroft & Mermin, 1976). Key physical quantities such as band gaps, DOS, and energy isosurfaces are derived from k-resolved electronic structure, i.e., eigenvalues as a function of k-point. Predicting k-resolved electronic structure from crystal geometry is thus a central challenge in computational materials discovery: once this challenge is addressed, downstream physical quantities follow naturally from k-space queries.

Density functional theory (DFT) (Hohenberg & Kohn, 1964; Kohn & Sham, 1965) and tight-binding (TB) (Slater & Koster, 1954) methods have enabled electronic structure prediction. However, DFT remains computationally prohibitive for high-throughput screening (Jones, 2015). Empirical TB is computationally efficient but requires system-specific parameterization: hopping integrals and on-site energies must be fitted for each material, and parameters rarely transfer across chemically distinct compounds (Goringe et al., 1997). Neither approach provides both computational efficiency and transferability to novel materials.

Machine learning surrogates address this gap: once trained, they offer rapid inference without material-specific fitting. However, to the best of our knowledge, no prior work has attempted to predict k-resolved electronic structure in this unified manner. An alternative approach is to predict electronic structure with respect to band. A recent study, Bandformer (Gong et al., 2024), predicts eigenvalues for a fixed set of bands along k-paths. However, per-band prediction is inherently limited: it requires fixing the number of bands under the assumption that band indices carry consistent meaning across materials, which is not generally valid. To truly predict k-resolved electronic structure, a principled approach must handle arbitrary k-points without relying on the band concept.

We propose **kRESForge**, which predicts energy bin occupancy at each k-point. Given a crystal structure and a query k-point, the model outputs a probability distribution over 256 energy bins spanning ± 10 eV from the Fermi level. This formulation provides native uncertainty through probabilistic predictions, all within a fixed-dimensional output and without assuming consistent band meaning across materials. Downstream physical quantities follow directly from the choice of k-sampling strategy without additional training: k-path queries yield band structures, k-grid aggregation yields DOS, and dense k-grid sampling enables energy isosurface extraction. On 28,517 non-magnetic materials from the Materials Project (Jain et al., 2013), kRESForge achieves a band gap MAE of 0.39 eV, 90% metal/non-metal classification accuracy, and DOS MAE of 2.64 states/eV.

Our main contributions are:

- We formulate k-resolved electronic structure prediction as soft energy occupancy prediction in kRESForge, enabling a band-agnostic approach with native uncertainty estimates.
- We show that the model directly yields band structure visualization via k-path queries, and enables extraction of downstream physical quantities such as band gaps, DOS, and energy isosurfaces through k-sampling strategy, all without task-specific training.

2 METHOD

Problem formulation. We model electronic structure as a k-conditioned prediction problem over a configurable energy grid. Given a crystal structure and a query k-point $\mathbf{k}$, we predict which energy bins contain electronic states. We discretize the energy axis into $N_{\text{bins}} = 256$ bins over $E_F \pm 10$ eV ($\sim$78 meV resolution). The model outputs $\mathbf{p}(\mathbf{k}) \in [0,1]^{N_{\text{bins}}}$, where $p_l(\mathbf{k})$ is the probability that energy bin l contains an eigenvalue. This fixed binning yields a band-count-invariant output, and the probabilistic outputs provide a native uncertainty signal (values near 0 or 1 indicate confident predictions, while intermediate values suggest uncertainty).

Architecture. kRESForge consists of three modules (Figure 1):

(1) Crystal graph encoder. We adopt iComFormer (Yan et al., 2024), an attention-based GNN for periodic crystals. Atoms are embedded from atomic numbers via a learnable lookup, and interatomic distances are encoded with Gaussian radial basis functions. After L layers of edge-aware gated attention, we obtain per-atom representations $\{\mathbf{h}_i\}_{i=1}^{N} \subset \mathbb{R}^d$.

(2) k-point embedding. Fractional coordinates $\mathbf{k} = (k_1, k_2, k_3)$ defined with respect to the reciprocal lattice vectors, are encoded via sinusoidal functions (Cui et al., 2023):

$$\mathbf{h}_{\mathbf{k}} = \text{MLP}\Big([\sin(2\pi k_j),\ \cos(2\pi k_j)]_{j=1}^{3}\Big) \in \mathbb{R}^d, \tag{1}$$

which respects Brillouin zone periodicity by construction. This embedding conditions the decoder on $\mathbf{k}$, enabling k-resolved predictions with a shared fixed-dimensional output.

(3) Transformer decoder (Vaswani et al., 2017). We define $T = N_{\text{bins}}/P$ query tokens (patch size P), where each token $1 \le t \le T$ corresponds to a patch of P consecutive energy bins. Each token is initialized as $\mathbf{q}_t = \bar{\mathbf{h}} + \mathbf{e}_t + \mathbf{h}_{\mathbf{k}}$, where $\bar{\mathbf{h}} = \sum_i \mathbf{h}_i$ is the sum-pooled structure representation and $\{\mathbf{e}_t\}_{t=1}^{T}$ are learnable positional embeddings for each query token. Each transformer block applies self-attention among tokens, cross-attention to per-atom features $\{\mathbf{h}_i\}$, and a feedforward network (pre-LayerNorm). A linear head maps each token to P logits, yielding $\mathbf{p} \in [0,1]^{N_{\text{bins}}}$ after sigmoid.

Training. We train with binary cross-entropy loss on ground-truth labels obtained by discretizing DFT eigenvalues into energy bins. The model outputs soft occupancy probabilities, enabling flexible downstream extraction.

Downstream extraction. Because kRESForge predicts eigenvalue presence at arbitrary k-points without relying on band indices, the same trained model can extract diverse physical quantities at inference by simply varying the k-sampling strategy, all without task-specific retraining.

Band structure and band gap. Querying along high-symmetry k-paths produces band structure visualizations spanning the full $E_F \pm 10$ eV range. To extract the band gap, we aggregate per-bin predictions over the k-path to obtain a spectral sum $S(E_l) = \sum_{\mathbf{k}} p_l(\mathbf{k})$ and apply a hard threshold $\tau_{\text{gap}} = 9$ (optimized on the validation set). We compute the band gap as the energy distance between

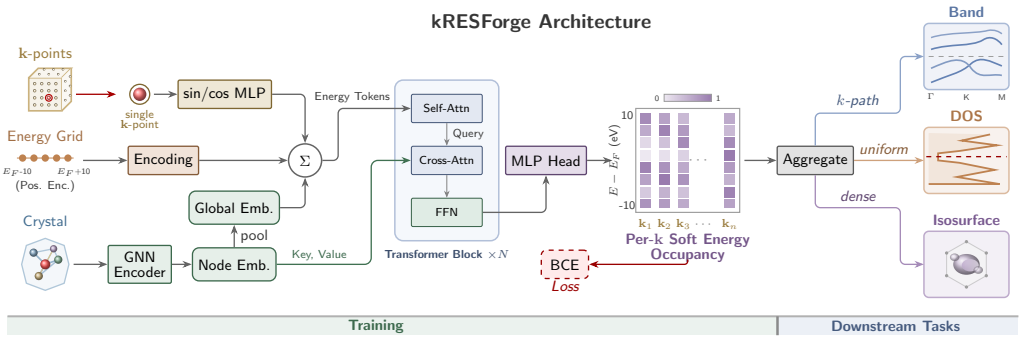

Figure 1: Overview of kRESForge. The GNN encoder (iComFormer) produces per-atom features $\{\mathbf{h}_i\}$ independently of the k-point. The decoder conditions on $\mathbf{k}$ via sinusoidal embedding and predicts per-k-point soft energy occupancy $\mathbf{p}(\mathbf{k})$. Different k-sampling strategies yield downstream physical quantities: k-path queries produce band structures, uniform k-grid aggregation yields DOS, and dense k-grid sampling enables energy isosurface extraction.

the highest-energy bin with $S(E_l) > \tau_{\text{gap}}$ below E_F and the lowest-energy bin with $S(E_l) > \tau_{\text{gap}}$ above E_F. Materials with $S(E_F) > \tau_{\text{gap}}$ are classified as metals (band gap $= 0$).

Density of states. We obtain the DOS by aggregating predictions over a $8 \times 8 \times 8$ Monkhorst–Pack grid (Monkhorst & Pack, 1976). We apply a confidence squashing operator to down-weight low-probability bins before aggregation using $\tau_{\text{dos}} = 0.2$ (optimized on the validation set):

$$\tilde{p}_i(\mathbf{k}) = \sigma\left(\frac{p_i(\mathbf{k}) - 0.5}{\tau_{\text{dos}}}\right), \tag{2}$$

$$D_{\text{raw}}(E_i) = \textstyle\sum_{\mathbf{k}} \tilde{p}_i(\mathbf{k}), \tag{3}$$

where σ is the sigmoid function. Because soft occupancy predicts state presence rather than multiplicity, D_{raw} underestimates the true spectral weight. We correct this with a multiplicative scale factor $\hat{s}$, predicted via simple linear regression on $\sum_l D_{\text{raw}}(E_l)$, N_e, and N_{atoms}:

$$D(E_i) = \hat{s} \cdot D_{\text{raw}}(E_i). \tag{4}$$

Energy isosurface. Per-k-point predictions over a dense k-grid yield 3D energy isosurfaces. For metals, we extract Fermi surfaces by thresholding the maximum predicted probability near E_F; for non-metals, energy isosurfaces at VBM and CBM energies reveal band-edge topology in k-space.

3 Experiments

We do not compare directly with per-band methods such as Bandformer (Gong et al., 2024), as their fixed band indices formulation differs fundamentally from our band-agnostic approach; we instead evaluate on physically meaningful downstream quantities.

Setup. We use DFT band structures from the Materials Project (Jain et al., 2013), computed along Latimer–Munro high-symmetry k-paths (Munro et al., 2020), selecting 28,517 non-magnetic materials (train: 22,815 / val: 2,899 / test: 2,803). We assess the quality of predicted electronic structures via three derived physical quantities: (1) band gap (MAE in eV), (2) metal/non-metal classification (accuracy), and (3) density of states (DOS) via k-grid aggregation (MAE in states/eV).

Results. Table 1 summarizes quantitative performance. kRESForge achieves a band gap MAE of 0.39 eV (0.97 eV for non-metals) with Pearson $r = 0.80$, and metal/non-metal classification accuracy of 90% (F1 = 0.89). For DOS, k-grid aggregation without calibration yields MAE 2.73 states/eV; calibration with structural features reduces this to 2.64 states/eV (Pearson $r = 0.66$). The $\sim$78 meV bin resolution sets a lower bound on band gap precision.

Figure 2 demonstrates representative predictions for a metal ($TiAlCu_2$) and a semiconductor ($SrZrO_3$), and the model captures key electronic features for both materials. For the metal (a),

Table 1: Quantitative evaluation on the test set (non-magnetic materials).

	MAE	**Pearson** r	**Acc.**	**F1**
Band gap (total)	0.39 eV	0.80	—	—
Band gap (non-metal)	0.97 eV	—	—	—
Metal / non-metal	—	—	0.90	0.89
DOS (uncalibrated)	2.73 states/eV	0.66	—	—
DOS (calibrated)	2.64 states/eV	0.66	—	—

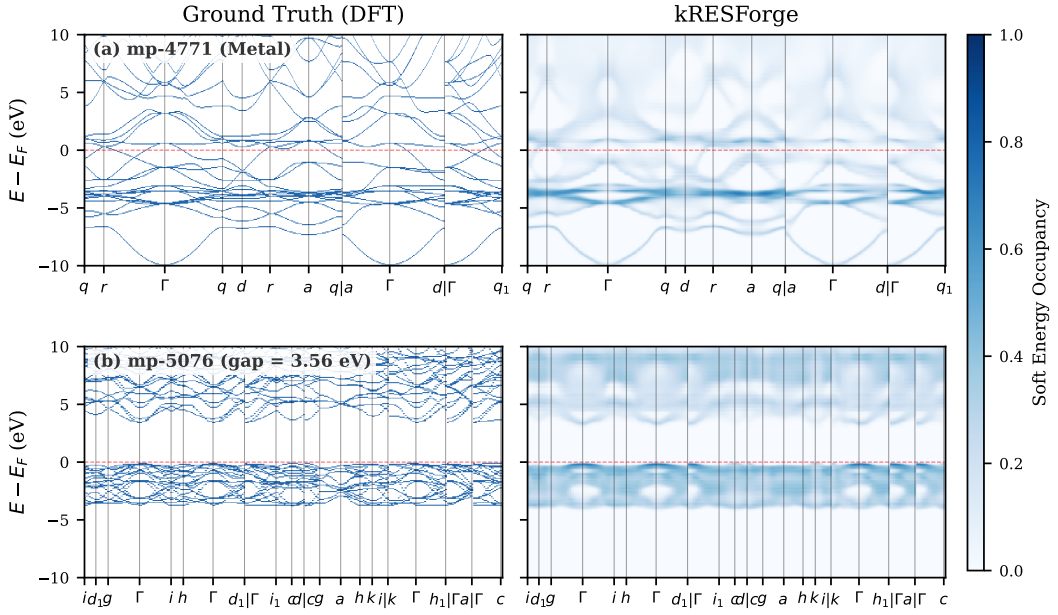

Figure 2: Soft energy occupancy predictions by kRESForge. *(a)* $TiAlCu_2$ (metal) and *(b)* $SrZrO_3$ (semiconductor), each showing DFT band structure (left) and predicted soft energy occupancy (right) across $E_F \pm 10$ eV.

the model correctly predicts continuous states across the Fermi level and captures the flat band region near -3 eV. For the semiconductor (b), the band gap region shows near-zero soft occupancy, and the model captures the overall band shape near VBM and CBM.

Figure 3 shows DOS derived from aggregating per-k-point soft energy occupancy over an $8 \times 8 \times 8$ k-grid, compared with DFT ground truth. For both materials, kRESForge captures key spectral features including overall shape, peak positions, and peak spacing. The semiconductor (b) correctly exhibits zero DOS in the band gap region, consistent with the k-resolved predictions in Figure 2.

Figure 4 illustrates energy isosurfaces extracted from $21 \times 21 \times 21$ k-grid predictions (interpolated to $105 \times 105 \times 105$) for the same materials. For the metal (top), kRESForge reproduces the Fermi surface topology, capturing the overall shape of the DFT reference. For the semiconductor (bottom), the predicted VBM and CBM isosurfaces correctly localize in k-space, with VBM concentrated near Γ and CBM distributed at different k-points, consistent with the DFT reference.

4 Discussion & Conclusion

Discussion The soft energy occupancy formulation departs from the traditional approach of predicting individual physical properties separately. Rather than tracking individual band identities, kRESForge predicts per-energy bin eigenvalue presence at each k-point. By predicting k-resolved electronic structure in this unified manner, the model enables both physical interpretation and extrac-

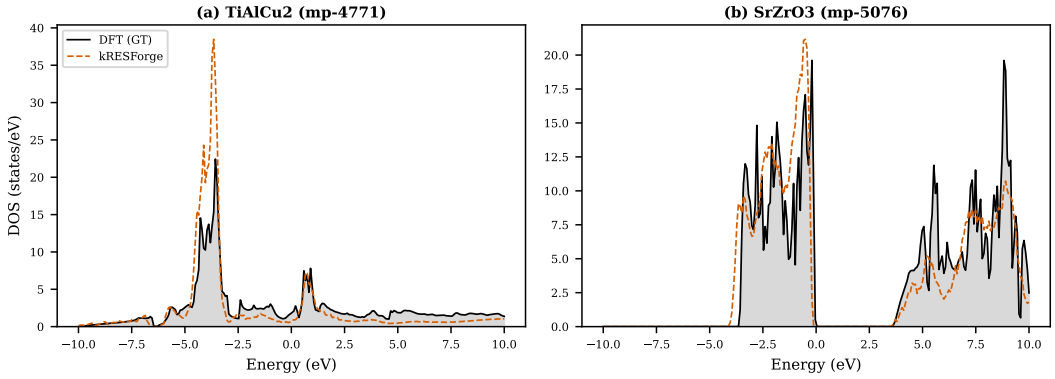

Figure 3: DOS extraction for the same materials as Figure 2. *(a)* $TiAlCu_2$ (metal) and *(b)* $SrZrO_3$ (semiconductor). Ground truth (solid) and kRESForge prediction (dashed) across the full $E_F \pm 10$ eV range.

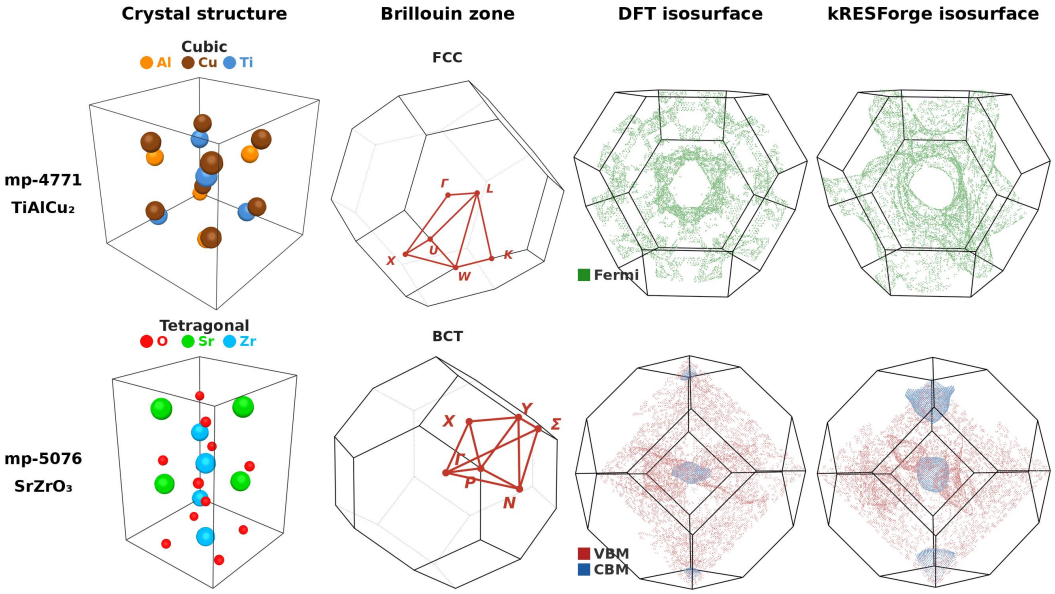

Figure 4: Predicted energy isosurfaces in the Brillouin zone. Top row: Fermi surface of $TiAlCu_2$ (metal). Bottom row: VBM and CBM energy isosurfaces of $SrZrO_3$ (semiconductor). Each row shows crystal structure, Brillouin zone, DFT reference, and kRESForge prediction from left to right.

tion of multiple downstream quantities (band gaps, DOS, and energy isosurfaces), pointing toward a foundation model approach for understanding electronic properties in materials.

The probabilistic outputs reveal interpretable patterns that may reflect underlying physics. As shown in Figure 2, prediction confidence varies across energy regions, with the model capturing key features such as flat bands and band edges. However, regions with increased prediction uncertainty suggest greater modeling challenges, potentially due to intrinsic physical complexity, encoder expressiveness, or training data limitations. Intermediate probability values do not represent physical electron occupations, but rather indicate model uncertainty.

Furthermore, by predicting soft occupancy for each energy bin at each k-point, our formulation produces a k-resolved, energy-resolved output that bears resemblance to the spectral function $A(\mathbf{k}, \omega)$ measured in angle-resolved photoemission spectroscopy (ARPES) (Damascelli et al., 2003), suggesting potential alignment with experimental spectral data. This resemblance suggests potential extension beyond DFT-derived training data to models trained directly on experimental spectra.

Conclusion. We propose kRESForge, which reformulates electronic structure prediction as soft energy occupancy prediction. By predicting per-k-point occupancy over 256 energy bins, the model has no dependence on band count, while probabilistic outputs provide inherent uncertainty estimates. On 28,517 non-magnetic materials, kRESForge achieves a band gap MAE of 0.39 eV and 90% metal/non-metal accuracy, while enabling DOS and Fermi surface extraction from the same predictions. This unified formulation offers a scalable alternative to eigenvalue regression for high-throughput materials screening. Future work includes extending to spin-polarized systems and exploring alignment with ARPES measurements.

ACKNOWLEDGEMENTS

This work was supported by the Technology development Program (RS-2025-25462152) funded by the Ministry of SMEs and Startups (MSS, Korea).

REFERENCES

Neil W Ashcroft and N David Mermin. *Solid State Physics*. Holt, Rinehart and Winston, 1976.

Yong Cui, Kai Chen, Liwei Zhang, Han Wang, Lei Bai, Daniel Elliston, and Wei Ren. Atomic positional embedding-based transformer model for predicting the density of states of crystalline materials. *The Journal of Physical Chemistry Letters*, 14:7924–7932, 2023.

Andrea Damascelli, Zahid Hussain, and Zhi-Xun Shen. Angle-resolved photoemission studies of the cuprate superconductors. *Reviews of Modern Physics*, 75(2):473–541, 2003.

Weiyi Gong, Tao Sun, Hexin Bai, Jeng-Yuan Tsai, Haibin Ling, and Qimin Yan. Graph transformer networks for accurate band structure prediction: An end-to-end approach, 2024.

C M Goringe, D R Bowler, and E Hernández. Tight-binding modelling of materials. *Reports on Progress in Physics*, 60(12):1447–1512, 1997.

Pierre Hohenberg and Walter Kohn. Inhomogeneous electron gas. *Physical Review*, 136(3B):B864–B871, 1964.

Anubhav Jain, Geoffroy Hautier, Charles J Moore, Shyue Ping Ong, Christopher C Fischer, Tim Mueller, Kristin A Persson, and Gerbrand Ceder. A high-throughput infrastructure for density functional theory calculations. *Computational Materials Science*, 50(8):2295–2310, 2011.

Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, and Kristin A Persson. Commentary: The Materials Project: A materials genome approach to accelerating materials innovation. *APL Materials*, 1(1):011002, 2013.

R O Jones. Density functional theory: Its origins, rise to prominence, and future. *Reviews of Modern Physics*, 87(3):897–923, 2015.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.

Walter Kohn and Lu Jeu Sham. Self-consistent equations including exchange and correlation effects. *Physical Review*, 140(4A):A1133–A1138, 1965.

Georg Kresse and Jürgen Furthmüller. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Physical Review B*, 54(16):11169–11186, 1996.

Georg Kresse and Daniel Joubert. From ultrasoft pseudopotentials to the projector augmented-wave method. *Physical Review B*, 59(3):1758–1775, 1999.

Richard M Martin. *Electronic Structure: Basic Theory and Practical Methods*. Cambridge University Press, 2nd edition, 2020.

Hendrik J Monkhorst and James D Pack. Special points for Brillouin-zone integrations. *Physical Review B*, 13(12):5188–5192, 1976.

Jason M Munro, Keith Latimer, Matthew K Horton, Shyam Dwaraknath, and Kristin A Persson. An improved symmetry-based approach to reciprocal space path selection in band structure calculations. *npj Computational Materials*, 6(1):112, 2020.

John P Perdew, Kieron Burke, and Matthias Ernzerhof. Generalized gradient approximation made simple. *Physical Review Letters*, 77(18):3865–3868, 1996.

John C Slater and George F Koster. Simplified LCAO method for the periodic potential problem. *Physical Review*, 94(6):1498–1524, 1954.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, 2017.

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 (ICLR)*, 2024.

A APPENDIX

A.1 IMPLEMENTATION DETAILS

Model configuration. The iComFormer encoder uses 4 attention-based message-passing layers (plus 1 edge update layer) with hidden dimension $d = 256$ and 1 attention head. Interatomic distances are encoded via Gaussian RBF with 256 basis functions. The k-point embedding maps the 6-dimensional sinusoidal encoding $[\sin(2\pi k_j), \cos(2\pi k_j)]_{j=1}^{3}$ through a two-layer MLP ($6 \to 64 \to 256$). The transformer decoder has 4 blocks (8 attention heads, FFN multiplier 4) with patch size $P = 4$, yielding $T = 64$ query tokens. We train for 600 epochs with Adam (Kingma & Ba, 2015) (lr $= 10^{-3}$), batch size 64, on a single NVIDIA A100 GPU. Checkpoints are saved every 20 epochs; the best is selected by band gap MAE on the validation set with threshold optimization.

Dataset details. We use band structures from the Materials Project (Jain et al., 2013). Non-magnetic materials are identified by filtering on the `is_spin_polarized` attribute. The final dataset contains 28,517 materials (train: 22,815 / val: 2,899 / test: 2,803).

Band gap extraction. From the predicted 256-bin occupancy vector, we extract band gaps by summing predictions over all k-points and applying a hard threshold $\tau = 9.0$ (optimized on the validation set) to identify occupied energy bins. The gap is measured as the energy distance between the last occupied bin below the Fermi level and the first occupied bin above.

DOS prediction. A soft calibration mask $\sigma((p - 0.5)/0.2)$ is applied before summation to suppress low-confidence predictions. Because the model predicts binary occupancy probabilities rather than state counts, the raw predicted DOS underestimates the true spectral weight. We predict a multiplicative scale factor via linear regression on three input-accessible features—raw DOS integral, total electron count, and number of atoms—fitted on the training and validation sets ($R^2 = 0.91$, test $R^2 = 0.93$).

A.2 DFT CALCULATIONS FOR ENERGY ISOSURFACES

DFT calculations for energy isosurfaces. Band structure data for training and evaluation are obtained from the Materials Project (Jain et al., 2013). For energy isosurface visualization (Figure 4), we performed additional DFT calculations on a dense k-grid using VASP (Kresse & Furthmüller, 1996) with the Perdew-Burke-Ernzerhof (PBE) exchange-correlation functional (Perdew et al., 1996) and projector augmented-wave (PAW) pseudopotentials (Kresse & Joubert, 1999). A plane-wave energy cutoff of 520 eV was used with an electronic convergence criterion of 2×10^{-6} eV. Electronic occupations were treated using Methfessel-Paxton smearing with a width of $\sigma = 0.2$ eV.

Spin-polarized calculations were performed with initial magnetic moments assigned as follows: for $TiAlCu_2$, 0.6 μ_B for Ti and 5.0 μ_B for Al and Cu; for $SrZrO_3$, 0.6 μ_B for all atoms. Hubbard U corrections were applied to transition metal d orbitals following the Materials Project standard (Jain et al., 2011): $U_{\text{Ti}} = 0.0$ eV, $U_{\text{Cu}} = 4.0$ eV (with $U = 5.0$ eV, $J = 1.0$ eV in Dudarev's formulation), and $U_{\text{Zr}} = 0.0$ eV.