


ARTICLE

Data-Driven Screening of High-Performance Interconnect Materials: Integrating Graph Learning with Engineering Safety Constraints

Jiayi Tang¹, Liang Cao^{2,*} , Guanghui Xu¹, Manqi Dong² and Ming Li³

¹Information Materials and Intelligent Sensing Laboratory of Anhui Province, Anhui University, No. 111 Jiulong Road, Hefei, China

²School of Integrated Circuits, Anhui University, No. 111 Jiulong Road, Hefei, China

³State Key Laboratory of Metal Matrix Composites, School of Materials Science and Engineering, Shanghai Jiao Tong University, No. 800 Dongchuan Road, Shanghai, China

*Corresponding Author: Liang Cao. Email: caolianghit@foxmail.com

Received: 03 March 2026; Accepted: 21 April 2026; Published: 15 June 2026

ABSTRACT: The accelerated design of next-generation semiconductor interconnects faces a critical “applicability gap”. Purely data-driven models effectively navigate vast chemical spaces, but they often yield candidates that are theoretically performant yet violate practical manufacturing constraints. To bridge this disconnect, this study proposes a neuro-symbolic decision support framework that systematically integrates inductive graph learning with deductive engineering logic for Safe-by-Design material screening. The framework operates through a hierarchical dual-stream architecture. First, an inductive Graph Neural Network (GNN) engine transforms 3D crystal structures into topological graph representations to predict thermodynamic stability and metallicity with high discriminative power (AUC = 0.868). Second, a deductive safety layer enforces explicit domain ontology, including toxicity thresholds, raw material costs, and reactivity limits, to preemptively prune high-risk candidates. Operationally, this hybrid approach reduces the candidate search space of over 20,000 compounds by approximately 96% within minutes, demonstrating orders-of-magnitude computational efficiency gains over traditional *ab initio* high-throughput screening. The system’s reliability is further validated through structural perturbation analysis and high-fidelity physics simulations, identifying robust binary compounds such as HfB and NbAl₃ that exhibit cohesive energies up to 2.1 times that of copper. These results demonstrate the efficacy of integrating symbolic reasoning with deep learning to create transparent, reliability-aware computational tools for early-stage engineering decision-making.

KEYWORDS: Neuro-symbolic AI; GNN; decision support systems; material informatics; safe-by-design; engineering constraints

1 Introduction

The continuous scaling of semiconductor technology towards the sub-10 nm regime imposes escalating complexity on the co-design of material systems and process control protocols, as highlighted in the latest IRDS 2024 Metrology Roadmap [1]. As traditional metallization standards, primarily copper, encounter fundamental physical limitations due to electron surface scattering and electromigration [2–5], the engineering objective shifts from a single-variable optimization of conductivity to a high-dimensional Constraint Satisfaction Problem (CSP) [6]. An ideal interconnect candidate must simultaneously satisfy rigorous electrical, mechanical, and thermodynamic requirements to withstand fabrication stresses and ensure operational reliability [7,8]. Consequently, the discovery of such materials requires navigating a

vast design space where trade-offs between theoretical performance and manufacturing feasibility are non-trivial [1,9].

Traditional discovery paradigms, relying heavily on intuition or stochastic trial-and-error, are increasingly inefficient against the exponential demands of integrated circuit scaling [10,11]. While *ab initio* simulations based on Density Functional Theory (DFT) offer predictive fidelity [12,13], their computational cost restricts their utility to small-scale validation rather than large-scale search space exploration [1,14]. The emergence of data-driven Artificial Intelligence (AI) tools has begun to address this scalability bottleneck [15,16]. However, a critical “applicability gap” remains in the domain of metallic interconnects: purely data-driven models, which typically operate as opaque black boxes, frequently generate candidates that are theoretically performant yet violate stringent engineering constraints—such as toxicity, cost, and thermodynamic stability [17,18]. As highlighted in recent closed-loop discovery and hybrid learning paradigms, bridging the disconnect between algorithmic prediction and real-world engineering feasibility remains a fundamental barrier to the adoption of AI in safety-critical manufacturing workflows [19–21].

To address this challenge, establishing a trustworthy decision support system that enforces algorithmic safety is essential. This study introduces a novel neuro-symbolic computational framework designed to accelerate the rational discovery of binary interconnect materials. By integrating data-driven inductive learning with domain-specific deductive logic—specifically utilizing hybrid GNN/DFT approaches [22,23]—this approach shifts the search paradigm toward “Safe-by-Design” engineering informatics. Unlike static screening methods, this framework explicitly models the interaction between material properties and process constraints through a dual-stream architecture [22,24].

The proposed framework employs an inductive GNN engine for high-throughput screening [23,25], achieving an Area Under the Curve (AUC) of 0.868. Adopting the informed machine learning paradigm [26], a deductive knowledge layer integrates expert domain ontology as hard safety constraints to preemptively prune candidates based on cost, toxicity, and reactivity rules. Furthermore, to mitigate selection bias, we introduce a diversity-aware ranking mechanism that optimizes the candidate portfolio [25]. The framework’s reliability is further validated through structural perturbation analysis, confirming its robustness against data uncertainty. This integration ensures that identified candidates, such as NbAl_3 and HfB , are not only high-performing—demonstrating cohesive energies comparable to those of copper—but are also structurally reliable and manufacturable [2,27,28].

2 System Architecture and Methodology

2.1 Architectural Overview

As illustrated in Fig. 1, the proposed neuro-symbolic decision support framework is architected as an integrated discovery loop grounded in the latest paradigms of AI-driven materials design [29]. The framework harmonizes three core functional units: Data Ingestion & Partitioning, the Inductive Inference Engine, and the Deductive Safety & Optimization Layer. Distinct from conventional sequential screening, this architecture utilizes the deductive layer as a symbolic governor to ensure that the inductive engine’s probabilistic outputs remain within the high-dimensional engineering and physical constraint manifold, a strategy recently emphasized for ensuring reliability in functional materials discovery [26,29]. This hierarchical design allows for the seamless coupling of data-driven pattern recognition with rule-based engineering logic. The workflow initiates with the retrieval of crystallographic data from the Materials Project (MP) database [30], which is subsequently partitioned into training, validation, and testing sets to ensure rigorous model evaluation [31]. Unlike monolithic deep learning models, this architecture separates the probabilistic prediction of material properties from the deterministic enforcement of safety constraints,

ensuring that the final material recommendations are both high-performing and rigorously compliant with manufacturing standards [32].

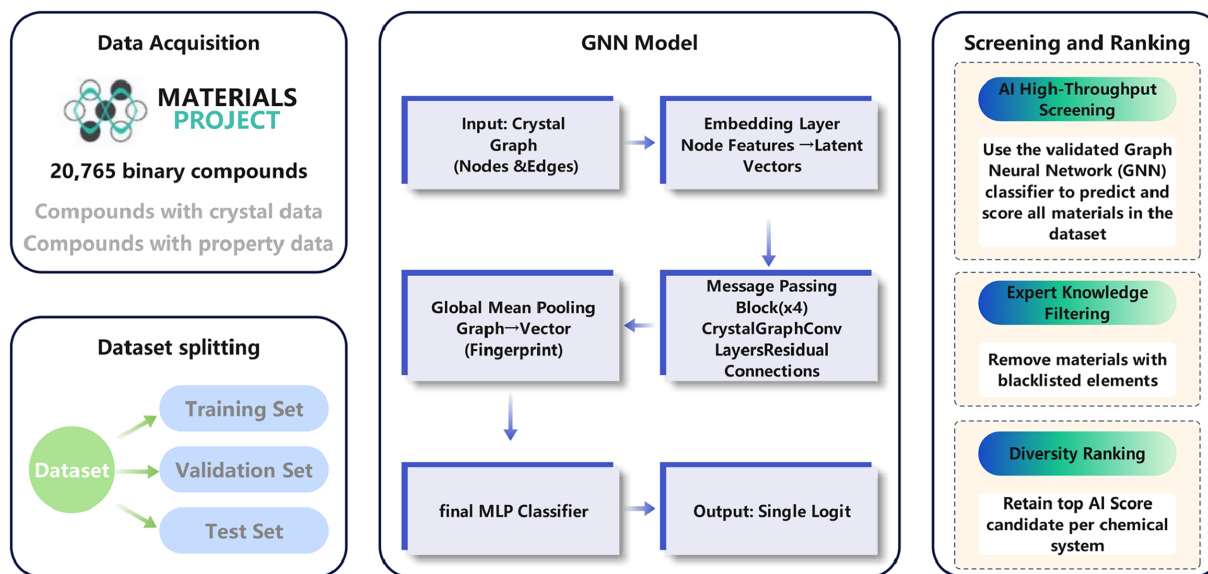


Figure 1: Architectural schematic of the proposed neuro-symbolic decision support framework. The system integrates three functional modules: data ingestion & partitioning, the inductive inference engine (GNN), and the deductive safety & optimization layer.

2.2 Graph Representation and Inductive Inference

The data foundation for the inductive module is sourced from the Materials Project database via its API [30]. We curated a dataset comprising 20,765 unique binary compounds, each associated with a verified crystal structure and calculated properties. To ensure methodological transparency, we explicitly note that the ground-truth property labels for this large-scale dataset were retrieved directly from the pre-calculated Density Functional Theory (DFT) records in the MP database, rather than from independent *ab initio* calculations by the authors. To operationalize the decision boundary, we defined a binary classification target, *is_compliant_candidate*. A positive label (1) is assigned if a material satisfies two physical criteria: thermodynamic stability ($energy_above_hull < 0.05$ eV/atom) and metallic character ($band_gap < 0.1$ eV) [13]. We explicitly acknowledge that these electronic descriptors, retrieved from the Materials Project database, are calculated using the PBE-GGA functional, which is documented to systematically underestimate electronic band gaps. Despite this inherent limitation, the use of GGA-level data remains the pragmatic standard for initial high-throughput screening across large chemical spaces comprising >20,000 compounds. To ensure the robustness of our discovery, candidates identified by the GNN are subsequently subjected to high-fidelity electronic structure verification—accounting for potential “false metallic” identifications—as detailed in Section 3.7. The resulting dataset contains 7723 positive samples (37.2%). Stratified sampling based on unique Materials Project IDs (mp-ids) was employed to partition the data into training (64%), validation (16%), and test (20%) sets, ensuring consistent class distribution across all subsets while rigorously preventing data leakage by ensuring structurally identical compounds are not shared between partitions [31,33].

The core of the inductive stream is a GNN designed to map crystal structures to property probabilities [34,35]. The transformation logic is explicitly detailed in Fig. 2a, which illustrates the encoding protocol from the physical 3D crystal lattice into a graph topology $G(V,E)$. In this process, atoms are mapped to nodes initialized with physicochemical descriptors—atomic number (Z), Pauling electronegativity (χ),

or economic landscapes without the prohibitive computational requirement of retraining the underlying neural architecture.

2.3 Model Optimization Protocol

The inductive model was trained using the Adam optimizer (learning rate = 0.001) for 100 epochs [38]. To ensure numerical stability during the binary classification task, *BCEWithLogitsLoss* was utilized as the objective function. The training process integrated a rigorous validation protocol, monitoring AUC, Precision, Recall, and F1-score after each epoch [28]. To maximize generalization capability and mitigate overfitting, the model checkpoint yielding the highest AUC on the validation set was preserved as the final inference engine.

2.4 Deductive Safety Layer and Decision Strategy

Following the inductive prediction, the framework employs a deterministic decision strategy to enforce algorithmic safety [29]. This strategy reflects the emerging transition from traditional black-box screening toward integrative, “Safe-by-Design” discovery frameworks. By architecting a hierarchical decision funnel, the system systematically reconciles probabilistic structural intuition with the rigorous requirements of semiconductor manufacturing. The quantitative impact of this multi-stage logic is visualized in the decision funnel in Fig. 2b, which demonstrates how the system systematically reduces the candidate pool through rigorous defined logic gates. Initially, the GNN acts as a high-throughput filter, applying a probability threshold of 0.5 to reduce the search space from 20,765 to 3754 candidates.

Subsequently, the process activates the Rule-Based Safety Layer. As visualized in the domain ontology map in Fig. 2c, this deductive module embeds expert knowledge as “hard constraints”, defining a forbidden set of elements associated with high toxicity (e.g., As, Pb), prohibitive cost (e.g., Au, Pt), or extreme reactivity. This safety guardrail effectively intercepts hazardous outputs, pruning approximately 78% of the algorithmically generated candidates and narrowing the viable pool to 826.

Finally, to counter “mode collapse” and ensure portfolio robustness, we applied a Bias Mitigation Strategy based on a Diversity-Aware Ranking algorithm [39]. This approach is further enhanced by active learning strategies [40,41] to refine the discovery of binary intermetallics. This strategy ensures the final recommendation list is robust, diverse, and compliant with engineering safety standards.

To ensure the deterministic enforcement of engineering constraints, the safety layer is formulated as a CSP. The decision function $\Psi(m)$ for a candidate material m within the global set M is defined as a conjunction of inductive and deductive predicates:

$$\Psi(m) = G(m) \wedge \neg Tox(m) \wedge Eco(m) \wedge Stab(m) \quad (1)$$

where $G(m)$ represents the inductive GNN prediction probability exceeding the threshold $\tau = 0.5$. The symbolic predicates are implemented via a modular, pluggable interface: $Tox(m)$ evaluates the presence of prohibited elements (e.g., As, Pb) based on toxicity ontologies; $Eco(m)$ enforces economic viability thresholds (e.g., cost metrics for Au, Pt); and $Stab(m)$ ensures compliance with deductive thermodynamic stability rules. This architectural separation between neural inference and symbolic verification—often referred to as a “decoupled neuro-symbolic system”—is a deliberate design choice. Unlike monolithic end-to-end models, this modular interface facilitates “dynamic rule-updating”, enabling the system to adapt to evolving regulatory landscapes or supply-chain fluctuations without the prohibitive computational cost of model retraining. Such an approach aligns with modern requirements for transparent and reliability-aware decision support in complex engineering domains.

To further optimize the decision space and prevent algorithmic mode collapse—a common failure in purely data-driven discovery—the candidate pool is refined through a Diversity-Aware Ranking algorithm. By grouping candidates into their respective chemical systems and selecting only the “System Champion” based on the highest $G(m)$ score, the framework ensures a heterogeneous representation of the material search space, providing engineers with a robust portfolio of structural alternatives. This methodology resonates with current trends in AI-assisted materials discovery, such as “self-driving laboratories” and “closed-loop discovery paradigms” [19,29]. By integrating expert domain ontology as hard logic gates, our framework ensures that the discovery process remains grounded in physical reality and economic feasibility, addressing the critical “applicability gap” prevalent in purely data-driven material informatics.

2.5 Uncertainty Quantification via Conformal Prediction

To provide formal safety guarantees for the inductive inference engine, we employed Split Conformal Prediction (SCP) to calibrate the model’s predictive uncertainty. Unlike heuristic methods, SCP offers distribution-free mathematical guarantees. We define a non-conformity score s_i for a calibration set of size n , representing the heuristic error of the GNN model. For a target error rate $\alpha = 0.1$ (corresponding to a 90% confidence interval), we compute the empirical quantile \hat{q} as follows:

$$\hat{q} = \text{Quantile} \left(s_1, \dots, s_n; \frac{[(n+1)(1-\alpha)]}{n} \right) \quad (2)$$

The prediction interval for a new material m_{test} is then constructed by bounding the GNN point-estimate \hat{y} with \hat{q} . This calibrated envelope guarantees that the true stability property falls within the predicted interval with a marginal probability of at least $1 - \alpha$, establishing a rigorously defined safety buffer for downstream engineering decisions.

3 Results and Discussion

3.1 Performance Validation of the Inductive Inference Engine

The Inductive Inference Engine constitutes the primary probabilistic filter within the decision framework. Its operational reliability was rigorously evaluated on an independent test set comprising 4153 compounds [42]. The validation analysis focuses on two critical engineering metrics: numerical stability during optimization and resource efficiency in candidate identification.

3.1.1 Numerical Stability and Convergence Analysis

The training dynamics, as visualized in the learning curves in Fig. 3a, indicate a stable optimization trajectory. The training loss exhibits a smooth, monotonic decay, converging without significant oscillation, while the validation AUC steadily climbs and reaches a plateau after approximately 80 epochs. This convergence behavior suggests that the GNN architecture effectively extracts generalized structural features from the chemical graph representations without succumbing to overfitting or gradient instability. Consequently, the training protocol was standardized at 100 epochs to maintain the optimal balance between feature abstraction fidelity and computational cost.

3.1.2 Discriminative Capability and Resource Efficiency

The optimization trajectory and discriminative power of the inductive engine are quantified in Fig. 3. As illustrated in Fig. 3a, the model demonstrates robust convergence stability over 100 epochs. The synchronized decay of both training and validation losses, without significant divergence or rebound, confirms that the

architecture effectively captures generalizable structural patterns and remains resistant to overfitting. This stable learning process is further reflected in the steady rise of the validation AUC, which plateaus at approximately 0.84, providing a reliable foundation for the engine's predictive performance.

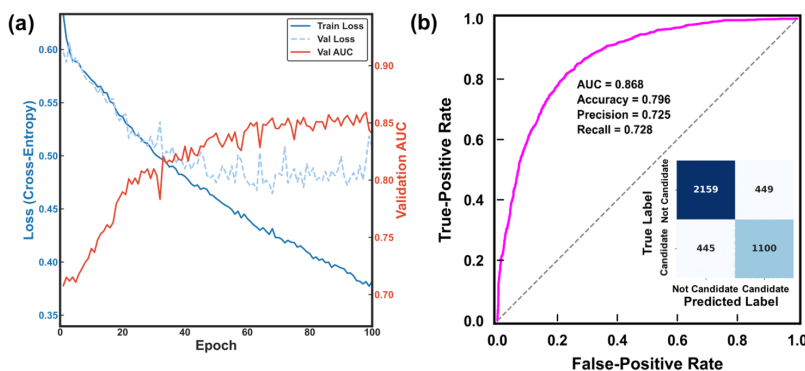


Figure 3: Quantitative performance validation of the inductive inference engine. (a) Optimization trajectory over 100 epochs, illustrating the synchronized convergence of training loss (solid blue line) and validation loss (dashed light-blue line) on the left ordinate, alongside the evolution of validation AUC (red line) on the right ordinate. (b) Discriminative capability analysis via the ROC curve and the embedded confusion matrix evaluated on the independent test set.

The final discriminative capability on the independent test set is detailed in Fig. 3b. The system achieves a test AUC of 0.868, indicating a high probability of correctly ranking compliant candidates above non-compliant ones. However, in a resource-constrained engineering discovery pipeline, the cost of validating false positives often outweighs the benefit of exhaustive recall. The embedded confusion matrix in Fig. 3b provides a granular view of this trade-off. At the standard decision threshold of 0.5, the engine attains a Precision of 72.5%. Specifically, out of 4153 test samples, the system successfully filtered out 2159 negative samples (True Negatives) while generating only 449 False Positives. In an engineering context, this high precision implies that approximately 3 out of every 4 candidates recommended by the inductive engine are valid, thereby significantly minimizing the computational resources wasted on downstream verification of infeasible materials. Simultaneously, the model maintains a Recall of 0.728 and an overall Accuracy of 0.796, confirming its role as a robust and efficient pre-screening module for the subsequent safety constraint layers.

3.2 Comparative Benchmarking of Graph Architectures

To objectively benchmark the efficacy of the proposed Inductive Inference Engine, a quantitative comparative study was conducted against two foundational graph neural network architectures in materials informatics: the Crystal Graph Convolutional Neural Network (CGCNN) [33] and the Message Passing Neural Network (MPNN) [35]. To ensure a controlled experimental environment, all three architectures were trained under identical protocols for 100 epochs, utilizing the same train-test splits.

As summarized in Table 1, the performance stratification reveals critical insights into the limitations of composition-only approaches. While the Random Forest and XGBoost baselines achieve respectable AUC scores (approximately 0.80), their predictive capability saturates at a distinct ceiling. Notably, their Precision metrics stagnate around 65%, implying that nearly 35% of the candidates they recommend are false positives. This empirical evidence underscores a fundamental limitation in material informatics: chemical composition alone serves as a necessary but insufficient descriptor for high-fidelity interconnect screening. The distinct performance leap observed in our GNN (AUC 0.868, Precision 72.5%) confirms that latent topological features—such as local coordination environments, bond directionality, and interstitial voids—are decisive

factors in determining electronic metallicity and thermodynamic stability. By capturing these structural-property relationships, the graph-based approach bridges the “applicability gap” that purely stoichiometric models fail to span.

Table 1: Comparative performance metrics on the independent test set.

Model Category	Model Architecture	AUC	Accuracy	Precision	Recall	F1-Score
Composition-based	Random Forest [43]	0.808	0.740	0.646	0.683	0.642
	XGBoost [44]	0.798	0.741	0.657	0.612	0.634
Structure-based	CGCNN [33]	0.844	0.773	0.685	0.722	0.703
	MPNN [35]	0.851	0.779	0.692	0.729	0.710
	Our GNN Classifier	0.868	0.796	0.725	0.728	0.727

From a systems engineering perspective, the choice of model architecture is governed not just by accuracy, but by the cost of error. In a resource-constrained discovery pipeline, the “cost of false positives”—the resources wasted on validating infeasible candidates via expensive DFT calculations or physical synthesis—constitutes the primary operational bottleneck. Consequently, Precision becomes the dominant metric for evaluating screening efficiency. The proposed architecture establishes a significant operational advantage, outperforming the standard MPNN (69.2%) and CGCNN (68.5%) in Precision by a margin of over 3%. Although MPNN exhibits a marginal advantage in Recall (0.729 vs. 0.728), this negligible 0.1% gain incurs a detrimental drop in Precision, which effectively introduces more noise into the candidate pool. Our architecture’s design priority—optimizing the signal-to-noise ratio—directly translates to reduced computational overhead for downstream verification, aligning with the core tenet of engineering informatics: delivering transparent, high-confidence decision support rather than mere statistical correlation.

3.3 Operational Feasibility and Engineering Validation

While predictive accuracy establishes the baseline competency of a data-driven model, its actual deployment in an industrial high-throughput screening pipeline demands a rigorous assessment of operational feasibility, specifically regarding computational scalability and epistemic reliability. As illustrated in Fig. 4a, the proposed graph inference engine demonstrates a linear time complexity ($O(N)$) with respect to system size, contrasting sharply with the cubic scaling ($O(N^3)$) typical of conventional *ab initio* simulations [45]. The inference latency increases linearly from 0.15 ms for binary unit cells to merely 8.7 ms for complex supercells containing 432 atoms, implying a throughput advantage of five orders of magnitude over traditional DFT methods. This efficiency is critical for the “Safe-by-Design” paradigm, allowing for the rapid exploration of vast chemical spaces that remain inaccessible to purely first-principles approaches.

Furthermore, the robustness of the decision boundary against data imperfections was evaluated to ensure the model captures intrinsic physical laws rather than superficial correlations. Real-world material databases often contain crystallographic measurement noise; thus, perturbation stability is critical. As depicted in Fig. 4b, the global AUC of the model exhibits a graceful, monotonic degradation even under severe feature perturbations, maintaining robust predictive power without catastrophic failure. This resilience suggests that the graph encoder has learned topological invariants that remain stable under distortion, which is essential for screening materials in early-stage discovery where precise lattice parameters may be uncertain.

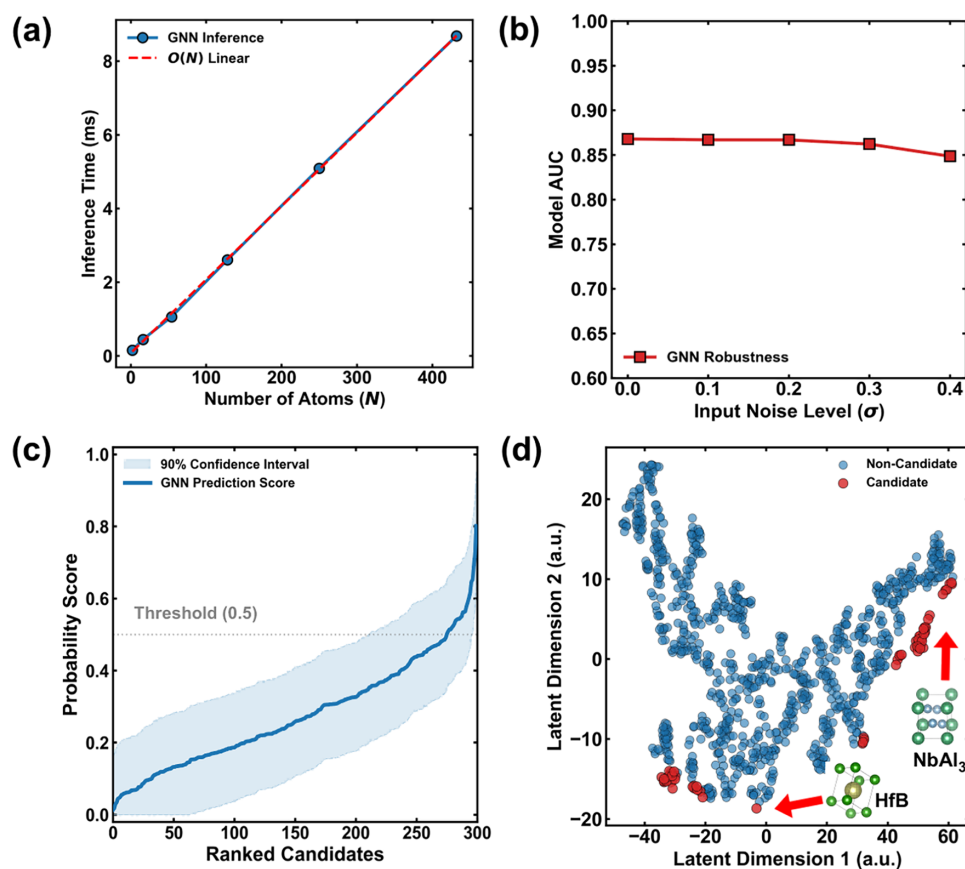


Figure 4: Operational validation of the decision framework. (a) $O(N)$ scalability showing linear inference latency suitable for large-scale screening. (b) Robustness analysis demonstrating performance stability under input feature perturbations. (c) Reliability manifold showing the 90% confidence interval calibrated via CP. (d) t-SNE projection of the latent space revealing clear topological clustering of viable candidates. Representative 3D atomic structures (NbAl₃ and HfB) are shown as insets to illustrate the structural similarities within the high-confidence candidate regions (red dots).

However, in engineering decision support, computational speed and robustness must be complemented by formal safety guarantees. As a baseline, the classification decision threshold was established at 0.5. This value serves as the natural mathematical equilibrium for a binary classifier optimized via BCE loss. Given the relatively balanced class distribution of our dataset (37.2% positive samples), maintaining this 0.5 threshold avoids introducing artificial inductive bias while securing a high Precision (72.5%), which is critical for minimizing the downstream computational waste associated with false positives. To this end, Fig. 4c visualizes the reliability manifold calibrated via Conformal Prediction (CP) [46]. Specifically, we employed Split Conformal Prediction (SCP) to rigorously quantify epistemic uncertainty. By calculating non-conformity scores (representing the heuristic error of the model) on an independent calibration set, we determined the empirical quantile corresponding to a target error rate of $\alpha = 0.1$. The deep blue trajectory represents the sorted GNN confidence scores, while the surrounding light blue envelope denotes the resulting 90% predictive interval. This framework transforms standard point-estimates into calibrated engineering signals. Unlike heuristic uncertainty quantification methods such as Monte Carlo Dropout, the CP-calibrated interval provides a mathematically grounded safety buffer. This enables a human-in-the-loop

(HITL) workflow where candidates with high score-volatility—typically those near the decision threshold—are automatically flagged for high-fidelity verification. This ensures the system remains “self-aware” of its knowledge limits, preventing over-confident misclassifications in ambiguous regions of the chemical space.

Finally, to elucidate the discriminative power of the engine, the high-dimensional latent feature space was projected onto a 2D manifold using t-SNE (Fig. 4d) [47]. The visualization reveals a distinct topological separation between high-confidence viable candidates and non-candidates. To assist in better appreciating this structural clustering, representative 3D atomic structures (NbAl₃ and HfB) corresponding to the two distinct candidate regions marked by red dots have been visualized as insets. This clear clustering confirms that the neuro-symbolic architecture has successfully extracted discriminative structural fingerprints, fundamentally distinguishing stable interconnect phases from unstable ones based on their intrinsic topological signatures. Even for materials from entirely different chemical families, their projection into similar high-confidence regions underscores the GNN’s reliance on shared geometric motifs, such as highly coordinated metallic bonding environments. Such representation learning is a prerequisite for reliable material discovery in complex engineering systems [48].

3.4 Decision Logic and Constraint Satisfaction Analysis

The systemic robustness of the proposed neuro-symbolic framework resides in its capacity to reconcile probabilistic “intuition” with deterministic engineering “rationality” [49]. While the inductive GNN engine efficiently explores the high-dimensional structural search space, the deductive safety layer serves as a critical symbolic safeguard to guarantee manufacturing feasibility. To illustrate this hierarchical decision-making procedure, Table 2 displays representative reasoning trajectories for candidates that triggered distinct system responses during the final screening stage.

Table 2: Representative decision reasoning trajectories: bridging inductive predictions with deductive engineering constraints.

Material	GNN Score	Safety Check	Cost Metric	Final Decision	Decision Logic/Reasoning
NbAl ₃	0.952	PASS	PASS	ACCEPTED	Consensus on high stability and economic viability.
ZrSi ₂	0.921	PASS	PASS	ACCEPTED	Optimal balance between performance and abundance.
Na ₃ Be	0.985	FAIL	PASS	REJECTED	Safety veto: Toxic Be and Na ⁺ contamination risk.
Li ₃ Mg	0.992	FAIL	PASS	REJECTED	Engineering veto: Low melting point incompatible with BEOL.
Cs ₃ Pt	0.981	PASS	FAIL	REJECTED	Economic veto: Prohibitive precious metal cost (Pt).
SF ₆	≈0	PASS	PASS	REJECTED	Inductive rejection of non-metallic molecular topology.

A pivotal observation in the decision-making logic is the system’s treatment of high-confidence false positives, exemplified by Na₃Be and Li₃Mg. As summarized in Table 2, both candidates achieved peak AI scores (≈1.000) from the inductive engine, owing to their ideal metallic topologies and high thermodynamic

stability. Nevertheless, these candidates were preemptively rejected by the deductive safety layer due to non-structural industrial constraints.

For Na_3Be , the veto was triggered by the intrinsic systemic toxicity of beryllium (Be) and the high ionic mobility of sodium (Na), the latter of which introduces a fatal risk of mobile-ion contamination to gate oxides in CMOS fabrication. Similarly, Li_3Mg was excluded due to its exceptionally low melting point ($\approx 180^\circ\text{C}$), which fails to satisfy the thermal budget requirements for back-end-of-line (BEOL) thermal cycling or high-power operating environments.

In contrast, the framework reached unified consensus for candidates such as NbAl_3 and ZrSi_2 , whose high structural stability scores were accompanied by full compliance with toxicity, reactivity, and cost metrics. Furthermore, the engine's capability for physical gatekeeping is demonstrated in the case of SF_6 . Without symbolic intervention, the GNN accurately recognized the non-metallic, molecular character of the crystal and assigned a near-zero probability score (7.2×10^{-36}), thereby decisively filtering out wide-bandgap insulators from the metallic search space.

This hierarchical filtering scheme, in which the GNN prunes physically inconsistent outliers and the safety layer enforces industrial safety standards, ensures that the final candidate pool is not merely a collection of statistical correlations but a set of technologically actionable engineering solutions [50].

3.5 GNN Model Interpretability and Physical Grounding

To evaluate the scientific integrity of the inductive engine beyond black-box metrics, we conducted a Permutation Feature Importance (PFI) analysis [51]. This process identifies the physical descriptors governing the decision boundary by quantifying performance degradation under feature-specific perturbations—a robust post-modeling explainability approach for deciphering complex neural representations [51]. As visualized in the feature attribution profile in Fig. 5a, the atomic number (Z) serves as the primary identifier, while Pauling electronegativity (χ) and atomic radius (r) exhibit substantial importance scores of 0.177 and 0.085, respectively.

The alignment between these attribution scores and classical metallurgical theories suggests that the GNN has successfully reconstructed the fundamental laws governing material stability and electronic transport [52]. The model's sensitivity to r directly corresponds to the Hume-Rothery size-factor rule, where a lattice strain resulting from atomic size mismatch governs the thermodynamic stability of intermetallic phases. More significantly, the reliance on χ reflects the model's grasp of the metallic bonding mechanism. In interconnect materials, a low electronegativity difference ($\Delta\chi$) between constituents favors the formation of delocalized electronic states, which are the physical prerequisite for the high metallic conductivity required in modern integrated circuits. Furthermore, the significance of the Z relates to the material's electron density and cohesive energy. From a mechanistic perspective, this serves as a critical proxy for the activation energy barrier (E_a) of atomic diffusion, which determines the material's intrinsic resistance against electromigration—a primary failure mode in high-density microelectronics. This correlation is further validated in the case study of NbAl_3 (Fig. 5b), where the near-identical electronegativity (1.60 vs. 1.61) and matched atomic radii (1.46 vs. 1.43 Å) provide the physicochemical basis for the stable, low-strain intermetallic structure identified by the system [52].

The high-confidence portfolio derived from this physics-informed logic, spearheaded by HfB, NbAl_3 , and ZrSi_2 , exhibits AI scores rigorous clustering between 0.972 and 0.998, reflecting the engine's robust certainty in their stability and metallicity (Fig. 5c). Fig. 5d contrasts the thermomechanical performance of these candidates against the elemental copper standard, revealing a decisive strategic gain in reliability. The transition metal boride HfB demonstrates a cohesive energy of 7.28 eV/atom—representing a 2.1-fold increase

over Cu. From a mechanistic perspective, this superior cohesive energy serves as a critical proxy for a higher E_a for atomic diffusion. In metallic interconnects, electromigration failure is fundamentally a mass-transport process driven by momentum transfer between charge carriers and lattice ions. By the empirical relation where the diffusion barrier Q_{EM} scales with the cohesive energy, the strengthened interatomic bonding in HfB and NbAl₃ effectively raises the threshold for atomic hopping. Furthermore, these high binding energies inherently imply a higher energy penalty for vacancy formation, which provides essential thermodynamic stability against atomic-scale fluctuations and structural degradation during the high-temperature cycles (e.g., >400°C) of BEOL processing. Simultaneously, the Young's Modulus for the top performers consistently exceeds 230 GPa, approximately doubling the mechanical stiffness of Cu (~130 GPa). This mechanical rigidity is pivotal for resisting thermal-stress-induced voiding and maintaining interfacial integrity under the extreme thermomechanical loads of advanced packaging.

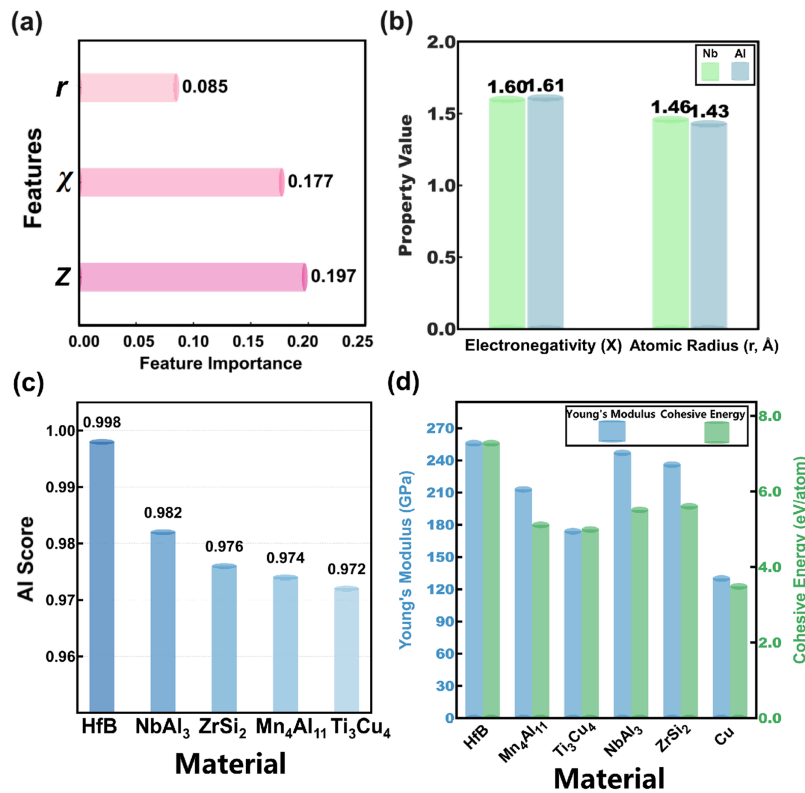


Figure 5: (a) Global feature importance analysis: this chart reveals that Z , χ , and r are the most crucial atomic descriptors for the model's predictions, and confirms that the model learns fundamental physicochemical principles governing chemical bonding and crystal structure. (b) The NbAl₃ case study: physicochemical validation of GNN feature importance. The close matching of electronegativity (~1.6) and atomic radius (~1.4 pm) between Nb and Al validates the model's reliance on these two key features for predicting metallic bonding and lattice stability. (c) The AI scores for the five top-tier candidate materials identified by our rigorous workflow: HfB, NbAl₃, ZrSi₂, Mn₄Al₁₁, and Ti₃Cu₄. (d) Comparison of Young's modulus (GPa) and cohesive energy (eV/atom) among top candidate materials and copper.

3.6 Execution of the Neuro-Symbolic Decision Funnel

The integration of the inductive GNN engine with deductive symbolic constraints was operationalized through a hierarchical decision funnel. This architecture was designed to systematically reduce the candidate

space from the initial 20,765 compounds to a prioritized set of materials that satisfy both performance and engineering compliance criteria [53].

As illustrated in the workflow (Fig. 2b), the first filtering stage utilized the GNN-based AI Score to assess the joint probability of thermodynamic stability and metallic character. By setting a classification threshold of 0.5, the search space was reduced by 81.9%, leaving 3754 candidates for further evaluation. This stage serves as a probabilistic coarse-grained filter, prioritizing compounds with structural features that align with the learned patterns of stability.

The second stage of the funnel involved the activation of the Rule-Based Safety Layer, governed by the domain ontology defined in Section 2.3. This layer executed a deterministic pruning of the 3754 candidates based on binary constraints including material cost, toxicity (e.g., Lead, Arsenic), and chemical reactivity. This deductive stage intercepted approximately 78% of the AI-selected candidates. The resulting pool of 826 materials represents a subset that meets the “Safe-by-Design” requirements necessary for integration into semiconductor fabrication lines.

To prevent algorithmic convergence on redundant chemical systems—a known limitation in purely data-driven material discovery—a Diversity-Aware Ranking algorithm was implemented. The 826 remaining candidates were clustered based on their chemical systems, and only the highest-scoring material from each system was retained as a “system champion”. This process identified 150 unique system champions representing 98 distinct space groups. This final ranking ensures that the workflow explores a heterogeneous range of crystal structures rather than clustering around a single chemical family.

The application of this hierarchical funnel demonstrates how symbolic logic can be utilized to constrain and refine the outputs of probabilistic neural networks, ensuring that the final material recommendations are both diverse and compliant with industrial safety standards.

3.7 Thermomechanical Properties of Candidate Materials

To validate the engineering reliability of the candidates identified by the neuro-symbolic framework, an *in-silico* audit was performed using high-fidelity physics simulations [54]. This stage serves as the final ground-truth verification to ensure that the probabilistic outputs of the inductive engine align with established mechanical and electronic standards. While our audit relies on static DFT calculations, it is worth noting that emerging neural time-series architectures now enable the acceleration of more complex dynamic simulations, including free-energy perturbations [55], offering a complementary path for evaluating material behavior under rare-event conditions.

The primary objective of this study is to identify alternatives that resolve the thermomechanical bottlenecks of conventional Cu interconnects. As visualized in Fig. 5d, the candidates selected by the framework demonstrate a strategic advantage in mechanical integrity over the Cu standard. Beyond cohesive energy, we evaluated the thermodynamic phase stability of these candidates via their energy-above-hull (E_{hull}). All identified candidates exhibit $E_{hull} < 0.02$ eV/atom, confirming their resistance to phase decomposition or unwanted transformations during the repetitive thermal annealing cycles (up to 400°C) characteristic of BEOL processing.

The intrinsic bonding strength was quantitatively evaluated via cohesive energy, which serves as a critical proxy for electromigration resistance. From a mechanistic perspective, this superior cohesive energy reflects a higher E_a for atomic diffusion. Since electromigration failure is fundamentally a mass-transport process driven by momentum transfer between charge carriers and lattice ions, the strengthened interatomic bonding in candidates like HfB and ZrSi₂ effectively raises the threshold for vacancy formation and atomic hopping. Specifically, HfB exhibits a cohesive energy of 7.281 eV/atom, a 2.1-fold increase over the Cu standard

(~ 3.49 eV/atom). Other candidates, such as NbAl_3 (5.515 eV/atom) and Ti_3Cu_4 (4.994 eV/atom), also demonstrate enhanced bonding strengths. Coupled with Young's modulus values that significantly exceed the Cu baseline (~ 130 GPa), these materials provide a robust “safety margin” against thermal-stress-induced voiding—a critical failure mode arising from the thermal expansion mismatch between interconnects and surrounding dielectrics.

To ensure charge carrier mobility, we analyzed the electronic band structures and Density of States (DOS). We acknowledge that the PBE-GGA functional used in standard databases systematically underestimates electronic band gaps and may neglect spin-orbit coupling (SOC) effects, which can be significant for heavy elements such as Hf and Nb. As Sun et al. [56] demonstrated in their development of strongly constrained and appropriately normed semilocal density functionals, such limitations of conventional GGA functionals highlight the need for rigorous verification of electronic descriptors.

However, as shown in the high-fidelity band structures (Fig. 6), the identified candidates exhibit such pronounced and multiple band-crossings at the Fermi level (E_F) that their metallic character is robust against the typical energy shifts introduced by more advanced functionals or SOC corrections, a consideration recently highlighted in materials informatics by Liu et al. [57]. The band structures in Fig. 6 confirm this intrinsic metallic behavior, with no observable gap at E_F . Complementing this, the DOS analysis in Fig. 7 quantifies the concentration of electronic states available at E_F . The measured $N(E_F)$ values range from 0.16 to 6.77 states/eV/atom. Notably, Ti_3Cu_4 (Fig. 6d) exhibits a high $N(E_F)$ of 6.77, suggesting a strong potential for intrinsic conductivity despite the complex intermetallic lattice. This multi-dimensional verification ensures that the framework's recommendations are physically sound and functionally viable for advanced technology nodes.

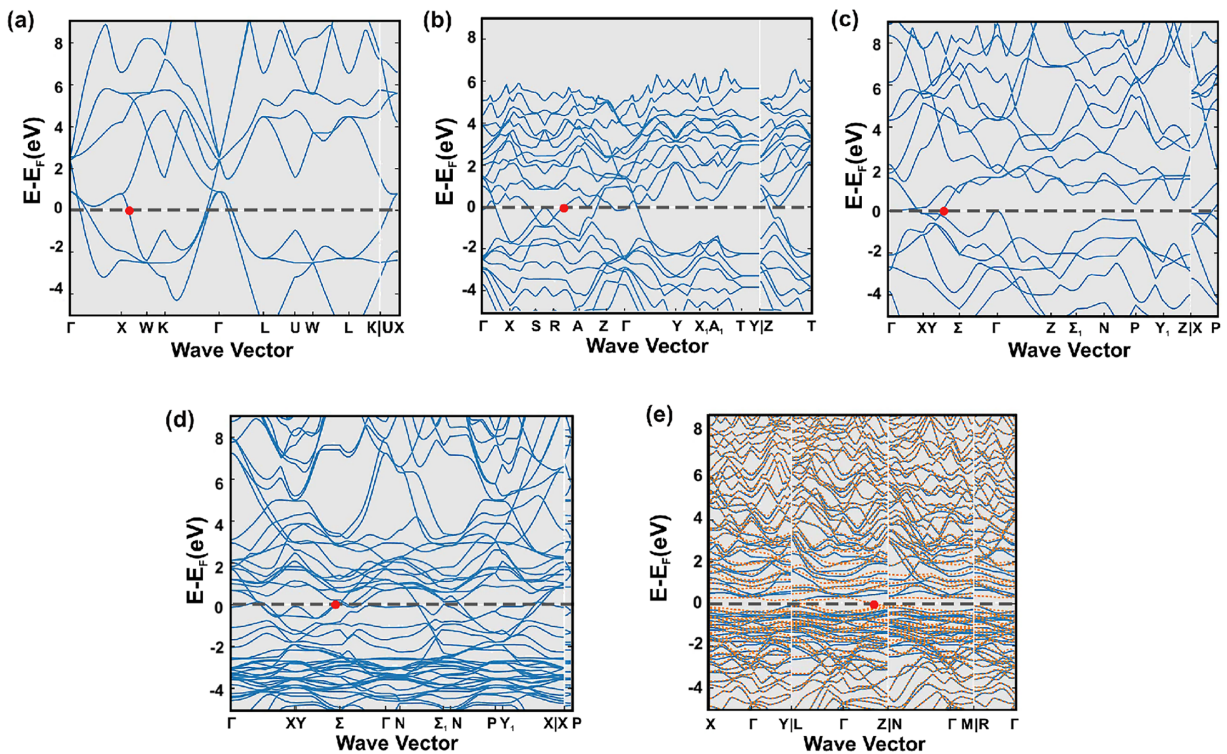


Figure 6: Electronic band structure of GNN-identified metallic candidates. (a–e) Band structures for five materials: (a) HfB, (b) ZrSi₂, (c) NbAl₃, (d) Ti₃Cu₄, and (e) Mn₄Al₁₁. The crossing of electronic bands through the Fermi level (E_F , marked by the red dots) confirms the metallic behavior.

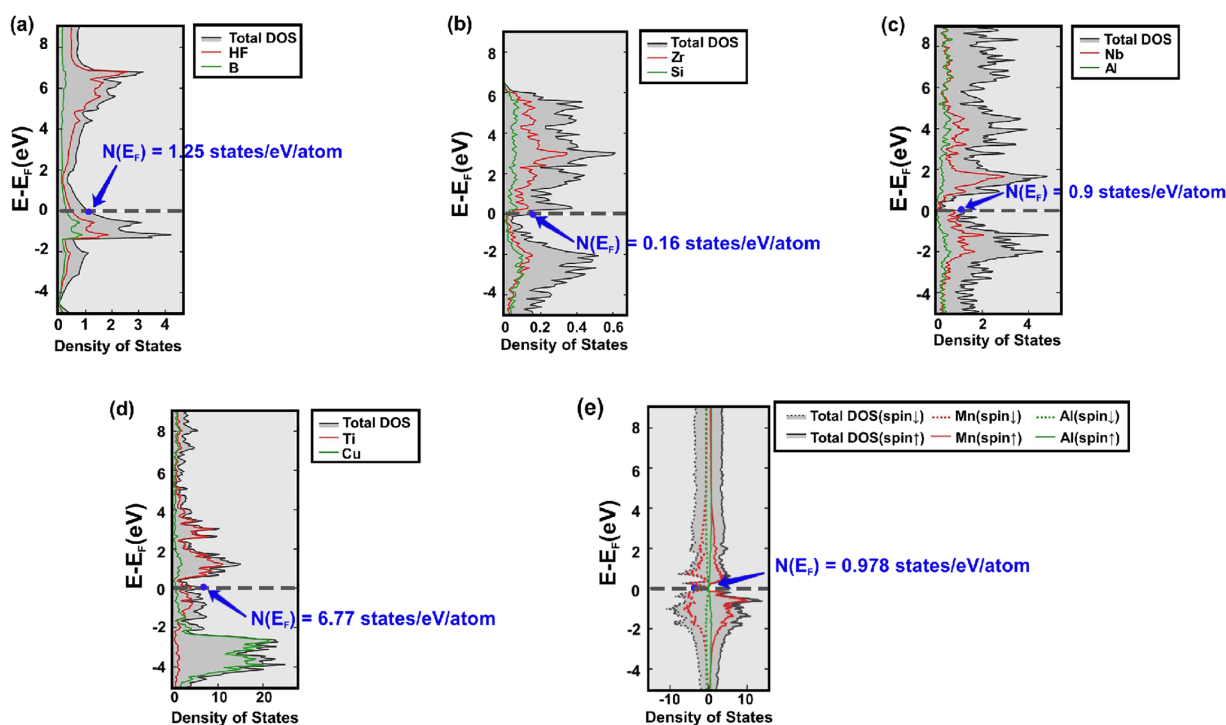


Figure 7: DOS analysis for GNN-identified metallic candidates. (a–e) DOS plots for five materials: (a) HfB, (b) ZrSi₂, (c) NbAl₃, (d) Ti₃Cu₄, and (e) Mn₄Al₁₁. The measured $N(E_F)$ value at the Fermi level ($E_F = 0$ eV) is indicated, correlating with the material's intrinsic conductive potential.

3.8 Methodological Superiority and Comparative Analysis

The decision framework presented in this study represents a shift from traditional single-stage screening paradigms toward a multi-stage, informatics-driven approach [58]. To evaluate its practical utility, the proposed workflow was benchmarked against the conventional *ab initio* screening methodology reported in recent literature. This comparison highlights the advantages of the neuro-symbolic architecture in terms of computational scalability and engineering relevance.

The primary bottleneck of traditional discovery is the prohibitive computational cost of high-fidelity simulations. Literature-standard methods often require hundreds of CPU/GPU hours to evaluate a pre-selected set of fewer than 150 candidates. In contrast, our framework leverages an inductive GNN engine as a high-throughput pre-filter. As demonstrated in the decision funnel analysis, the GNN can evaluate the initial pool of 20,765 compounds in less than one hour on standard hardware, effectively reducing the search space by over 80% with minimal resource consumption. This orders-of-magnitude increase in screening velocity facilitates the exploration of vast chemical spaces that remain inaccessible to purely *ab initio* methods [59].

Furthermore, to address the critical challenge of out-of-distribution (OOD) robustness, we evaluated the framework's extrapolation capability across distinct chemical families. While many AI systems in materials science struggle with extrapolation, our GNN engine demonstrated significant stability across varying chemical spaces. As evidenced in the latent space projection (Fig. 4d), the model successfully identified high-potential candidates from fundamentally different families—specifically transition metal borides (e.g., HfB) and intermetallic aluminides (e.g., NbAl₃). The accurate classification of these materials, which possess disparate bonding characters and electronic environments, suggests that the GNN has captured intrinsic

structural-property invariants (such as local coordination stability and electronegativity-driven metallicity) rather than merely memorizing stoichiometric patterns from the training distribution.

Finally, the inclusion of a diversity-aware ranking algorithm ensures a broad chemical representation that traditional, intuition-driven selection often lacks. By partitioning candidates into their respective chemical systems and selecting only the “System Champion” based on the highest $G(m)$ score, the framework ensures a heterogeneous exploration of the material landscape spanning 98 distinct space groups. This systematic exploration confirms that the neuro-symbolic approach offers a more robust and generalizable paradigm for reliable material discovery in complex engineering systems.

4 Conclusion

This study has established a hierarchical neuro-symbolic framework designed to optimize the discovery of functional materials for semiconductor interconnect applications. By integrating an inductive GNN engine with deductive symbolic constraints, the developed workflow addresses the dual challenges of computational scalability and decision reliability in materials informatics. The primary methodological contribution lies in the operationalization of a multi-stage decision funnel that effectively navigates a chemical space of over 20,000 compounds. This “Safe-by-Design” paradigm ensures that candidates comply with industrial standards regarding toxicity, cost, and reactivity [49], positioning the framework as a generalizable template for engineering CSP. The systematic audit via physics-based simulations confirms the framework’s capacity to identify high-potential materials, such as HfB and NbAl₃, which exhibit a 2.1-fold increase in cohesive energy compared to conventional copper interconnects.

Despite the framework’s capacity for rapid screening, several inherent limitations warrant discussion. The current inductive engine utilizes isotropic graphs, which may overlook long-range structural correlations governing anisotropic mechanical properties. Furthermore, while the safety layer enforces rigorous engineering standards, it currently relies on static constraints. To bridge the gap between theoretical screening and real-world industrial volatility, future iterations will leverage the system’s modular interface to facilitate dynamic rule-updating, enabling autonomous adaptation to supply-chain fluctuations or evolving regulatory landscapes without model retraining [60].

To address these limitations and transition the framework toward a more resilient decision support system, future research must focus on several key evolutionary paths. Future iterations of this framework will incorporate Dynamic Attention Mechanisms or Equivariant Graph Neural Networks to bridge this representational gap, enabling the system to predict directional material responses with greater fidelity for high-performance interconnect applications. Concurrently, the implementation of a HITL arbitration protocol will serve as the core of the active learning cycle. By leveraging the quantified epistemic uncertainty, the system will be designed to automatically trigger requests for expert intervention or targeted first-principles verification when encountering ambiguous candidates (i.e., $\sigma \approx 0.12$). This synergy ensures that human domain expertise and algorithmic efficiency are mutually reinforced, allowing the deductive layer to dynamically update its rule-base through iterative feedback. Furthermore, transitioning from static rules toward a Knowledge-Graph based reasoning system would enable more context-aware decisions that integrate real-world manufacturing constraints. The ultimate validation will require experimental assessment within an industrial process environment. In summary, the methodology established here provides a scalable and trustworthy template for materials engineering, evolving from a probabilistic discovery tool into a robust, human-centric, and reliability-aware decision system for accelerating innovation in critical functional domains [50].

Acknowledgement: The authors express their gratitude to the Materials Project (MP) for providing the open-access crystallographic and electronic structure datasets.

Funding Statement: This work is supported by the National Natural Science Foundation of China (No. 62504001) and the Start-Up Research Funding for the “Young Talents” Program at Anhui University.

Author Contributions: Jiayi Tang conceived and designed the study, developed the GNN architecture and diversity-aware ranking algorithms, performed data collection and analysis, and wrote the original manuscript; Liang Cao provided supervision, managed project administration, refined the methodology, and contributed to manuscript revisions; Guanghui Xu provided overall supervision, secured resources and funding, and contributed to manuscript revisions; Ming Li provided expertise in data interpretation and investigation, and conducted technical review of the manuscript; Manqi Dong contributed to visualization and formatting, and provided technical assistance during manuscript preparation. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The raw material data used in this study were retrieved from the Materials Project (<https://materialsproject.org/>). The custom Python scripts developed for the GNN classification and diversity analysis are available from the corresponding author upon reasonable request.

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Obeng Y, Mansfield E, Davydov A, Barnes B, Vladar A. International roadmap for devices and systems™ 2024 metrology. [cited 2026 Feb 25]. Available from: <https://irds.ieee.org/editions/2024>.
2. Gall D. The search for the most conductive metal for narrow interconnect lines. *J Appl Phys.* 2020;127(5):050901. doi:10.1063/1.5133671.
3. Gall D, Cha JJ, Chen Z, Han HJ, Hinkle C, Robinson JA, et al. Materials for interconnects. *MRS Bull.* 2021;46(10):959–66. doi:10.1557/s43577-021-00192-3.
4. Tu KN. Recent advances on electromigration in very-large-scale-integration of interconnects. *J Appl Phys.* 2003;94(9):5451–73. doi:10.1063/1.1611263.
5. Kumar S, Multunas C, Defay B, Gall D, Sundararaman R. Ultralow electron-surface scattering in nanoscale metals leveraging Fermi-surface anisotropy. *Phys Rev Materials.* 2022;6(8):085002. doi:10.1103/physrevmaterials.6.085002.
6. Das Sharma D, Pasdast G, Tiagaraj S, Aygün K. High-performance, power-efficient three-dimensional system-in-package designs with universal chiplet interconnect express. *Nat Electron.* 2024;7(3):244–54. doi:10.1038/s41928-024-01126-y.
7. Sun LG, Wu G, Wang Q, Lu J. Nanostructural metallic materials: structures and mechanical properties. *Mater Today.* 2020;38(31):114–35. doi:10.1016/j.mattod.2020.04.005.
8. Zhang YW, Wu BY, Chen KC, Wu CH, Lin SY. Highly conductive nanometer-thick gold films grown on molybdenum disulfide surfaces for interconnect applications. *Sci Rep.* 2020;10(1):14463. doi:10.1038/s41598-020-71520-x.
9. Van Troeye B, Sankaran K, Tokei Z, Adelman C, Pourtois G. First-principles investigation of thickness-dependent electrical resistivity for low-dimensional interconnects. *Phys Rev B.* 2023;108(12):125117. doi:10.1103/physrevb.108.125117.
10. Curtarolo S, Hart GLW, Nardelli MB, Mingo N, Sanvito S, Levy O. The high-throughput highway to computational materials design. *Nat Mater.* 2013;12(3):191–201. doi:10.1038/nmat3568.
11. Ren F, Ward L, Williams T, Laws KJ, Wolverton C, Hattrick-Simpers J, et al. Accelerated discovery of metallic glasses through iteration of machine learning and high-throughput experiments. *Sci Adv.* 2018;4(4):eaq1566. doi:10.1126/sciadv.aq1566.

12. Dudarev SL, Botton GA, Savrasov SY, Humphreys CJ, Sutton AP. Electron-energy-loss spectra and the structural stability of nickel oxide: an LSDA+U study. *Phys Rev B*. 1998;57(3):1505–9. doi:10.1103/physrevb.57.1505.
13. Lee J, Ikeda Y, Tanaka I. First-principles screening of structural properties of intermetallic compounds on martensitic transformation. *npj Comput Mater*. 2017;3(1):52. doi:10.1038/s41524-017-0053-8.
14. Mounet N, Gibertini M, Schwaller P, Campi D, Merkys A, Marrazzo A, et al. Two-dimensional materials from high-throughput computational exfoliation of experimentally known compounds. *Nat Nanotechnol*. 2018;13(3):246–52. doi:10.1038/s41565-017-0035-5.
15. Madika B, Saha A, Kang C, Buyantogtokh B, Agar J, Wolverton CM, et al. Artificial intelligence for materials discovery, development, and optimization. *ACS Nano*. 2025;19(30):27116–58. doi:10.1021/acsnano.5c04200.
16. Eslamlou AD, Ghasemlou A, Barros B, Riveiro B. A hybrid data-physics framework with conformal GNN for enhanced damage identification. *Adv Eng Inform*. 2025;68(13):103718. doi:10.1016/j.aei.2025.103718.
17. Nematov D, Hojamberdiev M. Machine learning-driven materials discovery: unlocking next-generation functional materials—a review. *Comput Condens Matter*. 2025;45(1):e01139. doi:10.1016/j.cocom.2025.e01139.
18. Chávez-Angel E, Eriksen MB, Castro-Alvarez A, Garcia JH, Botifoll M, Avalos-Ovando O, et al. Applied artificial intelligence in materials science and material design. *Adv Intell Syst*. 2025;7(8):2400986. doi:10.1002/aisy.202400986.
19. Zuo Y, Qin M, Chen C, Ye W, Li X, Luo J, et al. Accelerating materials discovery with Bayesian optimization and graph deep learning. *Mater Today*. 2021;51(4A):126–35. doi:10.1016/j.mattod.2021.08.012.
20. Njirjak M, Žužić L, Babić M, Janković P, Oтовиć E, Kalafatovic D, et al. Reshaping the discovery of self-assembling peptides with generative AI guided by hybrid deep learning. *Nat Mach Intell*. 2024;6(12):1487–500. doi:10.1038/s42256-024-00928-1.
21. Mo M, Yu H, Zhou X, Yang W, Leng C, Zhou W, et al. CMD-FEP: machine-learned free-energy prediction for efficient screening of material interfacial binder. *Adv Funct Mater*. 2026, e29571. doi:10.1002/adfm.202529571.
22. Matsokin NA, Eremin RA, Kuznetsova AA, Humonen IS, Krautsou AV, Lazarev VD, et al. Discovery of chemically modified higher tungsten boride by means of hybrid GNN/DFT approach. *npj Comput Mater*. 2025;11(1):163. doi:10.1038/s41524-025-01628-z.
23. Du H, Wang J, Hui J, Zhang L, Wang H. DenseGNN: universal and scalable deeper graph neural networks for high-performance property prediction in crystals and molecules. *npj Comput Mater*. 2024;10(1):292. doi:10.1038/s41524-024-01444-x.
24. Sun A, Xing S, Deng X, Shen R, Yan A, Hu F, et al. Edge-guided inverse design of digital metamaterial-based mode multiplexers for high-capacity multi-dimensional optical interconnect. *Nat Commun*. 2025;16(1):2372. doi:10.1038/s41467-025-57689-7.
25. Sun J, Li D, Zou J, Zhu S, Xu C, Zou Y, et al. Accelerating the discovery of acceptor materials for organic solar cells by deep learning. *npj Comput Mater*. 2024;10(1):181. doi:10.1038/s41524-024-01367-7.
26. von Rueden L, Mayer S, Beckh K, Georgiev B, Giesselbach S, Heese R, et al. Informed machine learning—a taxonomy and survey of integrating prior knowledge into learning systems. *IEEE Trans Knowl Data Eng*. 2023;35(1):614–33. doi:10.1109/tkde.2021.3079836.
27. Costa PMFJ, Gautam UK, Bando Y, Golberg D. Direct imaging of Joule heating dynamics and temperature profiling inside a carbon nanotube interconnect. *Nat Commun*. 2011;2(1):421. doi:10.1038/ncomms1429.
28. Debroy S, Sivasubramani S, Vaidya G, Acharyya SG, Acharyya A. Temperature and size effect on the electrical properties of monolayer graphene based interconnects for next generation MQCA based nanoelectronics. *Sci Rep*. 2020;10(1):6240. doi:10.1038/s41598-020-63360-6.
29. Cheng M, Fu CL, Okabe R, Chotrattanapituk A, Boonkird A, Hung NT, et al. Artificial intelligence-driven approaches for materials design and discovery. *Nat Mater*. 2026;25(2):174–90. doi:10.1038/s41563-025-02403-7.
30. Jain A, Montoya J, Dwaraknath S, Zimmermann NER, Dagdelen J, Horton M, et al. The materials project: accelerating materials design through theory-driven data and tools. In: Andreoni W, Yip S, editors. *Handbook of materials modeling. Methods: theory and modeling*. Cham, Switzerland: Springer International Publishing; 2018. p. 1–34. doi:10.1007/978-3-319-42913-7_60-1.

31. Yuan R, Liu Z, Balachandran PV, Xue D, Zhou Y, Ding X, et al. Accelerated discovery of large electrostrains in BaTiO₃-based piezoelectrics using active learning. *Adv Mater.* 2018;30(7):1702884. doi:10.1002/adma.201702884.
32. Moon JH, Jeong E, Kim S, Kim T, Oh E, Lee K, et al. Materials quest for advanced interconnect metallization in integrated circuits. *Adv Sci.* 2023;10(23):2207321. doi:10.1002/advs.202207321.
33. Bang K, Hong D, Park Y, Kim D, Han SS, Lee HM. Machine learning-enabled exploration of the electrochemical stability of real-scale metallic nanoparticles. *Nat Commun.* 2023;14(1):3004. doi:10.1038/s41467-023-38758-1.
34. Scarselli F, Gori M, Tsoi AC, Hagenbuchner M, Monfardini G. The graph neural network model. *IEEE Trans Neural Netw.* 2009;20(1):61–80. doi:10.1109/TNN.2008.2005605.
35. Zhou Y, Zheng H, Huang X, Hao S, Li D, Zhao J. Graph neural networks: taxonomy, advances, and trends. *ACM Trans Intell Syst Technol.* 2022;13(1):15. doi:10.1145/3495161.
36. Isayev O, Oses C, Toher C, Gossett E, Curtarolo S, Tropsha A. Universal fragment descriptors for predicting properties of inorganic crystals. *Nat Commun.* 2017;8(1):15679. doi:10.1038/ncomms15679.
37. Gilmer J, Schoenholz SS, Riley PF, Vinyals O, Dahl GE. Neural message passing for quantum chemistry. In: *Proceedings of the 34th International Conference on Machine Learning—Vol. 70; 2017 Aug 6–11; Sydney, Australia.* p. 1263–72.
38. Reyad M, Sarhan AM, Arafa M. A modified Adam algorithm for deep neural network optimization. *Neural Comput Appl.* 2023;35(23):17095–112. doi:10.1007/s00521-023-08568-z.
39. Noel J, Monterola C, Tan DS. Improving recommendation diversity without retraining from scratch. *Int J Data Sci Anal.* 2025;20(2):1151–60. doi:10.1007/s41060-024-00518-9.
40. Wang Y, Tang Q, Xu X, Weng P, Ying T, Yang Y, et al. Accelerated discovery of magnesium intermetallic compounds with sluggish corrosion cathodic reactions through active learning and DFT calculations. *Acta Mater.* 2023;255:119063. doi:10.1016/j.actamat.2023.119063.
41. Cui G, Guo Z, Ren X, Jiang Y, Jin X, Wu Y, et al. Active learning for the discovery of binary intermetallic compounds as advanced interconnects. *J Phys Chem Lett.* 2025;16(14):3579–88. doi:10.1021/acs.jpcclett.5c00386.
42. Wu Z, Ramsundar B, Feinberg EN, Gomes J, Geniesse C, Pappu AS, et al. MoleculeNet: a benchmark for molecular machine learning. *Chem Sci.* 2018;9(2):513–30. doi:10.1039/c7sc02664a.
43. Breiman L. Random forests. *Mach Learn.* 2001;45(1):5–32. doi:10.1023/A:1010933404324.
44. Chen T, Guestrin C. XGBoost: a scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; 2016 Aug 13–17; San Francisco, CA, USA,* p. 785–94. doi:10.1145/2939672.2939785.
45. Talaei Khoei T, Ould Slimane H, Kaabouch N. Deep learning: systematic review, models, challenges, and research directions. *Neural Comput Appl.* 2023;35(31):23103–24. doi:10.1007/s00521-023-08957-4.
46. Angelopoulos AN, Bates S. A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv:2107.07511.* 2021. doi:10.48550/arxiv.2107.07511.
47. Van der Maaten L, Hinton G. Visualizing data using t-SNE. *J Mach Learn Res.* 2008;9(11):2579–605.
48. Wang X, Liu A, Kara S. Machine learning for engineering design toward smart customization: a systematic review. *J Manuf Syst.* 2022;65(1):391–405. doi:10.1016/j.jmsy.2022.10.001.
49. Hitzler P, Sarker MK, editors. *Neuro-symbolic artificial intelligence: the state of the art.* Amsterdam, The Netherlands: IOS Press; 2021.
50. Kügler P, Dworschak F, Schleich B, Wartzack S. The evolution of knowledge-based engineering from a design research perspective: literature review 2012–2021. *Adv Eng Inform.* 2023;55(1):101892. doi:10.1016/j.aei.2023.101892.
51. Minh D, Wang HX, Li YF, Nguyen TN. Explainable artificial intelligence: a comprehensive review. *Artif Intell Rev.* 2022;55(5):3503–68. doi:10.1007/s10462-021-10088-y.
52. Zhong X, Gallagher B, Liu S, Kailkhura B, Hiszpanski A, Han TY. Explainable machine learning in materials science. *npj Comput Mater.* 2022;8(1):204. doi:10.1038/s41524-022-00884-7.
53. Himanen L, Geurts A, Foster AS, Rinke P. Data-driven materials science: status, challenges, and perspectives. *Adv Sci.* 2020;7(2):1903667. doi:10.1002/advs.201903667.
54. Babuska I, Oden JT. Verification and validation in computational engineering and science: basic concepts. *Comput Meth Appl Mech Eng.* 2004;193(36–38):4057–66. doi:10.1016/j.cma.2004.03.002.

55. Mo M, Yu H, Wu C, Xu Z, Liu J, Huang WQ, et al. A neural time-series learning method for accelerating free-energy perturbation and rare-event molecular dynamics simulations. *J Chem Inf Model.* 2026;66(5):2651–62. doi:10.1021/acs.jcim.5c03127.
56. Sun J, Ruzsinszky A, Perdew JP. Strongly constrained and appropriately normed semilocal density functional. *Phys Rev Lett.* 2015;115(3):036402. doi:10.1103/PhysRevLett.115.036402.
57. Liu M, Gopakumar A, Hegde VI, He J, Wolverton C. High-throughput hybrid-functional DFT calculations of bandgaps and formation energies and multifidelity learning with uncertainty quantification. *Phys Rev Mater.* 2024;8(4):043803. doi:10.1103/PhysRevMaterials.8.043803.
58. Ward L, Agrawal A, Choudhary A, Wolverton C. A general-purpose machine learning framework for predicting properties of inorganic materials. *npj Comput Mater.* 2016;2(1):16028. doi:10.1038/npjcompumats.2016.28.
59. Agrawal A, Choudhary A. Perspective: materials informatics and big data: realization of the “fourth paradigm” of science in materials science. *APL Mater.* 2016;4(5):053208. doi:10.1063/1.4946894.
60. Satorras VG, Hoogeboom E, Welling M. E(n) equivariant graph neural networks. In: *Proceedings of the 38th International Conference on Machine Learning*; 2021 Jul 18–24; Virtual.