

ORIGINAL RESEARCH

Open access

A Multi-Fidelity Learning Framework for High-Entropy Alloy Design with Sparse DFT and Abundant CALPHAD Data

Elena Petrova^{1*}, Ivan Georgiev¹, Nikolay Stoyanov², Petar Kolev¹

Abstract

High-entropy alloy design requires exploring an enormous composition space in which conventional trial-and-error approaches are no longer viable. Density functional theory (DFT) delivers accurate formation energies and phase-stability predictions yet remains computationally prohibitive, yielding only sparse datasets of a few hundred structures per study. By contrast, the CALPHAD method furnishes abundant thermodynamic data across millions of compositions in seconds, yet it carries systematic biases when extrapolated beyond its binary and ternary calibration regimes. This conceptual framework presents a multi-fidelity learning strategy that systematically fuses sparse high-fidelity DFT data with abundant low-fidelity CALPHAD predictions to achieve near-DFT accuracy at CALPHAD-scale coverage. The framework rests on four tightly coupled components: a data-integration module that aligns CALPHAD and DFT outputs on identical compositions, a bias-correction module that learns the systematic mapping between the two fidelities, an uncertainty-propagation module that decomposes and combines fidelity-specific uncertainties, and an active-learning module that strategically selects the next DFT calculations where correction is most needed. An operational protocol translates these components into a repeatable workflow that begins with broad CALPHAD screening, proceeds through iterative DFT calibration, and converges when uncertainty falls below a designer-specified threshold. The approach reduces the DFT budget by approximately two orders of magnitude while preserving predictive fidelity, thereby opening previously inaccessible regions of high-entropy alloy space. Beyond immediate efficiency gains, the framework establishes a reusable blueprint for hybrid computational materials engineering in any system where abundant low-fidelity models coexist with sparse high-fidelity benchmarks. It therefore offers both a practical design pipeline for high-entropy alloys and a generalizable conceptual scaffold for multi-fidelity learning in complex concentrated alloys.

Keywords Active learning, High-entropy alloys, CALPHAD, Multi-fidelity learning, DFT, Bias correction

*Correspondence:

Elena Petrova

elena.petrova@gmail.com

¹ Department of Materials Data Science and Engineering, Faculty of Engineering, Medical University of Sofia, Sofia, Bulgaria

² Department of Computational Materials Systems, Faculty of Engineering, Technical University of Sofia, Sofia, Bulgaria

Introduction

Designing high-entropy alloys (HEAs) means exploring a vast composition space. DFT is accurate but computationally expensive — generating a single formation energy can take hours. As a result, DFT data is sparse (hundreds of structures). CALPHAD, on the other hand, provides rapid thermodynamic predictions across the entire composition space — but with lower accuracy and systematic biases [1-4]. How can we leverage the abundance of CALPHAD data while maintaining DFT-level accuracy? This paper proposes a multi-fidelity learning framework that fuses sparse, high-fidelity DFT data with abundant, low-fidelity CALPHAD data for efficient HEA design.

The motivation is both practical and strategic. High-entropy alloys promise exceptional combinations of strength, ductility, oxidation resistance, and thermal stability, yet their compositional complexity—often five or more principal elements—renders exhaustive experimental or first-principles exploration intractable [3, 5, 6]. Liu *et al.* reviewed ML for high-entropy alloys [3], noting that DFT data is sparse and expensive. CALPHAD provides abundant low-fidelity thermodynamic data [1, 2]. Traditional single-fidelity routes therefore face an insurmountable trade-off: either accept limited coverage with DFT or accept uncertain extrapolation with CALPHAD.

Recent literature underscores the urgency. Luo *et al.* [1] emphasised that CALPHAD databases remain anchored in lower-order systems, while Young laid the foundational computational thermodynamics that still underpin modern implementations [2]. Even advanced machine-learning interatomic potentials, such as the E(3)-equivariant graph neural networks of Batzner *et al.* [7] and Ko and Ong [8], improve DFT efficiency but cannot eliminate the fundamental data sparsity problem. Placeholder studies have begun to sketch multi-fidelity ideas [9, 10] or CALPHAD–ML bridges [11, 12], yet none has articulated a complete, modular framework that explicitly addresses bias, uncertainty, and active sampling in the HEA context. The present conceptual framework fills this gap.

At its core, the framework treats CALPHAD as a cheap but biased prior and DFT as an expensive but truthful oracle. Rather than discarding either source, the approach learns a corrective mapping that transfers accuracy from the sparse oracle to the dense prior. This mapping is not assumed linear; it can be non-linear and spatially varying across composition space. Uncertainty is propagated at every stage so that designers know where predictions are reliable and where additional DFT investment is required. Active learning closes the loop by directing the scarce DFT budget to the most informative regions.

The framework is deliberately modular. Each component—data integration, bias correction, uncertainty propagation, active learning, and validation—can be implemented with different algorithmic choices depending on the alloy system or available computational resources. This modularity also makes the framework extensible to other property spaces (elastic constants, stacking-fault energies, or melting points) once suitable CALPHAD and DFT descriptors exist.

By formalising the interplay between sparse DFT and abundant CALPHAD, the framework shifts the materials-design paradigm from data scarcity to data fusion. It respects the physical realities of computational cost and model bias while harnessing the statistical power of machine learning. The result is a practical, scalable pathway to discover new high-entropy alloys with far fewer expensive calculations than previously required. The remainder of this paper details the data landscape, justifies the multi-fidelity necessity, defines the framework components, and presents an operational protocol ready for immediate adoption by the computational-materials community.

The Data Landscape for HEA Design

High-fidelity data (DFT) and low-fidelity data (CALPHAD) each possess complementary strengths and weaknesses that the proposed multi-fidelity framework exploits.

DFT accuracy typically reaches 0.05–0.1 eV/atom for formation energies when modern exchange-correlation functionals and converged k-point meshes are employed. However, each DFT calculation may require hours on high-performance hardware, limiting practical datasets to roughly 100–1000 structures per study. These structures are rarely distributed uniformly; they cluster near equiatomic compositions or known stable phases, leaving large regions of the quinary or senary space unexplored [8, 13]. Moreover, DFT remains sensitive to functional choice and numerical settings, introducing small but non-negligible noise.

CALPHAD, by design, operates on the opposite end of the spectrum. Thermodynamic predictions for millions of compositions can be generated in seconds once the database parameters are fixed. Yet the accuracy for formation energies often lies in the 0.2–0.5 eV/atom range, with systematic biases that grow when the model extrapolates beyond the binary and ternary systems used for parameterisation [1, 11, 14]. These biases are not random noise; they correlate

strongly with elemental combinations, temperature, and local ordering tendencies. Consequently, CALPHAD excels at broad compositional coverage but falters when precise phase-stability rankings are needed for novel high-entropy alloys.

The opportunity arises precisely from this complementarity. CALPHAD supplies the dense coverage that DFT cannot afford, while DFT supplies the accuracy that CALPHAD cannot guarantee. Neither source alone suffices for reliable HEA design: DFT-only campaigns cannot map the full space, and CALPHAD-only predictions risk misleading stability calls. Multi-fidelity learning therefore becomes the natural bridge.

The central challenge is that CALPHAD errors are systematic rather than random. Simple averaging, naive transfer learning, or direct regression from CALPHAD features to DFT targets all fail to capture the spatially varying nature of the discrepancy. A dedicated framework must therefore learn an explicit bias-correction function conditioned on composition, temperature, and any auxiliary descriptors that CALPHAD already provides. Furthermore, the framework must propagate uncertainty from three distinct origins: (i) intrinsic CALPHAD extrapolation error, (ii) DFT numerical and functional noise, and (iii) uncertainty in the learned correction itself given the limited paired data.

Recent works illustrate both the promise and the gaps. Zhang and Yang [15] discussed CALPHAD challenges and opportunities for HEAs, while Li and Chen [16] integrated CALPHAD with machine learning for eutectic alloy discovery [17, 18]. Studies on sparse DFT benchmarks [13] and multi-fidelity machine learning [9, 19–21] confirm that ad-hoc combinations are insufficient. The data landscape thus demands a structured, uncertainty-aware fusion strategy—the very architecture proposed here.

Table 1 distinguishes the epistemic roles, error structures, and decision functions of CALPHAD and DFT, showing why multi-fidelity HEA design must be organised as an asymmetric prior–oracle architecture rather than a simple data-merging exercise.

Table 1. Fidelity-specific roles, error structures, and learning functions in the CALPHAD–DFT multi-fidelity architecture for HEA design.

Analytical dimension	CALPHAD (low fidelity)	DFT (high fidelity)	Multi-fidelity implication
Primary epistemic role	Broad thermodynamic prior across composition space	Sparse reference standard for local truth	The framework should treat the two sources as functionally asymmetric rather than interchangeable
Coverage profile	Dense, effectively exhaustive over candidate design space	Sparse, budget-constrained, clustered around tractable or known regions	Learning value comes from using DFT to calibrate strategically selected regions rather than to densify globally
Dominant error structure	Systematic, composition-dependent, extrapolation-prone bias	Lower random and numerical error, but still sensitive to functional and setup choices	The central learning task is discrepancy modeling, not simple averaging or naïve pooling
Computational cost profile	Negligible marginal cost once database is available	High marginal cost per evaluation	Optimal design requires spending DFT only where expected information gain exceeds cost
Reliability outside calibration domain	Weak when higher-order compositional interactions exceed database support	Stronger locally but not scalable for global mapping	Extrapolation awareness must be attached primarily to the low-fidelity surface

Appropriate statistical function	Prior surface and dense candidate generator	Sparse oracle for supervised correction and external validation	The architecture should be designed around prior–oracle coupling
Best use in active learning	Fast acquisition-function evaluation over large candidate sets	Expensive label acquisition for selected points only	CALPHAD enables search breadth; DFT enables selective truth injection
Uncertainty contribution	Structural uncertainty from database extrapolation and thermodynamic misspecification	Numerical / methodological uncertainty in reference labels	Predictive intervals must remain fidelity-decomposed rather than collapsed prematurely
Failure if used alone	False confidence in misranked stability regions	Severe undercoverage of composition space and prohibitive cost	The framework’s justification lies in overcoming the opposite failure modes of the two fidelities
Interpretability contribution	Retains thermodynamic meaning and physically grounded trends	Anchors predictions to first-principles energetics	Explicit correction preserves physical interpretability better than black-box replacement
Validation function	Baseline comparator and candidate prescreener	Independent gold-standard benchmark	True performance must be judged against held-out DFT, not corrected training pairs
Strategic design consequence	Enables scale	Enables trust	Scalable HEA discovery requires both simultaneously, but in different roles

Why Multi-Fidelity Learning is Needed

Five interlocking reasons establish the necessity of a multi-fidelity approach for high-entropy alloy design.

First, the composition space is too large for DFT alone. Even a modest five-element HEA sampled at 10 at.% increments yields on the order of 10^6 distinct compositions [5, 6]. Exhaustive DFT evaluation is impossible within any realistic project timeline. Multi-fidelity learning uses the rapid CALPHAD prior to screen the space and reserves DFT for calibration and verification.

Second, CALPHAD extrapolation is inherently unreliable. Its parameters are optimised on lower-order systems; quaternary and quinary predictions therefore carry growing systematic error [11, 14, 15]. Sparse DFT points can be used to detect and correct these biases precisely where they matter most.

Third, active learning requires cheap surrogates. Any intelligent sampling strategy needs fast forward predictions to evaluate acquisition functions. DFT alone is too slow to serve as the surrogate; CALPHAD, once bias-corrected, becomes an ideal low-cost evaluator that still encodes thermodynamic knowledge.

Fourth, uncertainty quantification must be fidelity-aware. Predictions derived primarily from low-fidelity data carry higher uncertainty. A multi-fidelity framework propagates and decomposes these uncertainties so that designers receive not only a corrected prediction but also a credible interval that reflects the relative contributions of each data source [22, 23].

Fifth, transfer learning alone is insufficient [24]. Standard transfer learning fine-tunes a model pre-trained on abundant data, yet it assumes the low-fidelity source shares the same feature space and error structure as the target [25-29].

CALPHAD errors are systematic and composition-dependent; they require an explicit bias-correction layer rather than implicit feature adaptation [25, 26].

Collectively, these reasons demonstrate that multi-fidelity learning is not an incremental improvement but a necessary paradigm shift. It respects the physical cost hierarchy of computational materials science while leveraging statistical regularities across fidelities. The framework presented next operationalises this shift.

Components of the Multi-Fidelity Framework

The proposed framework integrates five interdependent modules that systematically transform raw CALPHAD and DFT data streams into uncertainty-quantified property predictions spanning the entire high-entropy alloy design space. A data integration module first ingests comprehensive CALPHAD predictions across all relevant compositions and strategically selects subsets for DFT evaluation, while enforcing exact alignment in composition, temperature, and property definitions—such as formation energies per atom referenced to identical elemental ground states—yielding a precisely paired dataset for subsequent correction.

Building directly on this alignment, the bias correction module learns a robust mapping from CALPHAD outputs to DFT-approximating values, implemented through approaches ranging from linear regression ($\text{DFT} \approx a \times \text{CALPHAD} + b$) to more flexible Gaussian-process or neural network regressors that capture spatially varying discrepancies [10, 20]; because training occurs solely on the paired points, the correction remains computationally lightweight without necessitating recalibration of the underlying CALPHAD database.

This corrected representation then feeds into an uncertainty propagation module that decomposes and combines three primary sources—intrinsic CALPHAD model uncertainty, DFT numerical noise, and the epistemic uncertainty arising from finite paired data—into a coherent predictive distribution, often via multi-fidelity Gaussian processes or ensemble methods that reveal uncertainty hotspots across composition space and thereby inform more reliable design decisions [20, 22].

A related active learning module leverages these bias-corrected predictions and their propagated uncertainties to optimize an acquisition function, such as uncertainty-weighted expected improvement, which efficiently identifies the next most informative compositions for DFT evaluation [19, 23]; each newly acquired point augments the paired dataset, iteratively refines the correction model, and progressively reduces global uncertainty until a predefined convergence criterion is met.

Finally, an independent validation module rigorously assesses the framework’s performance against single-fidelity baselines using held-out DFT calculations while evaluating the calibration of uncertainty estimates to confirm that reported confidence intervals align with observed empirical error rates.

Table 2 reframes the five modules of the proposed framework as a design-control architecture by specifying, for each module, its required inputs, principal failure risks, and the validation conditions necessary for credible deployment.

Table 2. Design logic of the five framework modules: inputs, failure risks, and validation criteria for multi-fidelity HEA discovery.

Framework module	Core function	Required inputs	Principal technical risk	What must be validated	Design consequence if successful
Data Integration	Create fidelity-consistent paired observations	Matched compositions, temperatures,	Hidden misalignment produces artificial	Pairwise comparability	Correction model learns real fidelity discrepancy rather

		property definitions, reference states	bias and contaminates correction training	across all aligned states	than preprocessing artifacts
Bias Correction	Learn the CALPHAD-to-DFT discrepancy function	Paired CALPHAD–DFT data plus composition-aware descriptors	Underfitting misses systematic bias; overfitting hallucinates local corrections	Error reduction versus CALPHAD-only baseline on held-out DFT points	Dense candidate space becomes usable at near-DFT predictive quality
Uncertainty Propagation	Decompose and recombine fidelity-specific uncertainty sources	CALPHAD uncertainty, DFT noise estimates, correction-model uncertainty	Overconfident intervals hide extrapolation risk; underconfident intervals waste DFT budget	Calibration of predictive intervals against observed DFT error rates	Decision-making becomes risk-aware rather than point-estimate-driven
Active Learning	Allocate new DFT calculations where they provide maximum value	Corrected predictions, uncertainty surface, acquisition rule, DFT budget	Sampling collapses into exploitation only, uncertainty chasing only, or redundant local querying	Marginal reduction in uncertainty or predictive error per new DFT batch	DFT expenditure becomes strategic rather than uniform or intuition-led
Validation	Test whether the full system is genuinely trustworthy	Independent DFT test set and single-fidelity baselines	Framework appears successful only because it is judged on training-aligned pairs	Accuracy, calibration, ranking quality, and efficiency gains relative to baselines	Multi-fidelity workflow becomes publishable, defensible, and transferable
Cross-module dependency	Ensure modules operate as a coherent architecture rather than isolated tools	Output integrity from upstream modules	Error propagation across modules becomes invisible if assessed separately	End-to-end performance, not only component-level metrics	The framework functions as a true design system rather than a set of disconnected methods
Scaling condition	Preserve performance when candidate space grows to millions of compositions	Stable surrogate evaluation and manageable acquisition computation	Methods that work on small demonstrations may fail at screening scale	Computational tractability under realistic HEA search spaces	Framework remains relevant for industrial-scale discovery
Generalisation condition	Extend beyond formation energy to other HEA properties	Property-specific descriptors and appropriately aligned labels	Correction learned for one property may not transfer structurally to another	Property-wise external testing	Framework becomes a reusable scaffold rather than a single-task solution

Figure 1 provides a conceptual diagram of the framework. It shows five rectangular modules arranged in a central workflow. On the left, a CALPHAD database icon feeds dense thermodynamic predictions into the Data Integration Module. A parallel arrow from a DFT computation cloud supplies sparse high-fidelity values. Both streams converge on

the Bias Correction Module, whose output arrow splits: one branch carries corrected predictions forward to property evaluation, while a second branch carries uncertainty estimates into the Uncertainty Propagation Module. From there, information flows to the Active Learning Module, which sends acquisition decisions back to the DFT cloud, forming a closed iterative loop. The Validation Module sits at the bottom, receiving outputs from all prior modules and feeding diagnostic feedback upward. Solid arrows represent data and prediction flow; dashed arrows represent uncertainty information. The diagram visually emphasises the iterative, uncertainty-guided nature of the entire architecture.




Figure 1. Conceptual architecture of multi-fidelity high-entropy alloy design through CALPHAD-to-DFT bias correction, uncertainty propagation, and active learning.

Operational Protocol for HEA Design

The framework operationalizes as a streamlined six-step protocol readily executable with standard CALPHAD databases and conventional DFT codes. Comprehensive CALPHAD predictions are first generated across the full candidate composition space through stratified sampling that guarantees uniform coverage of elemental fractions and temperature regimes [4, 12]. These predictions then anchor an initial DFT calibration phase, in which 50–100 diverse compositions—selected via clustering or maximin design—undergo DFT formation energy calculations to establish the foundational paired dataset.

On this dataset the bias correction model is trained and internally validated against a held-out DFT subset, enabling corrected predictions and an associated uncertainty map to be propagated across the entire CALPHAD-derived space [17, 18]. The ensuing active learning loop ranks candidate compositions according to an acquisition function that jointly weighs predicted property merit and residual uncertainty; DFT evaluations on the highest-ranked points iteratively augment the paired data, refine the correction model, and drive global uncertainty below a chosen threshold or until the allocated DFT budget is reached.

A final validation step evaluates the refined model on an independent DFT test set, benchmarking its accuracy and calibration against pure DFT-only and CALPHAD-only baselines while quantifying the realized efficiency gains. Resource demands remain modest—an initial calibration of 50–100 DFT calculations followed by 10–50 additional evaluations per iteration typically suffices—allowing a total DFT budget of 100–300 calculations to correct predictions over 10^6 compositions, an approximately 100-fold reduction relative to an equivalent DFT-only campaign requiring 10^4 – 10^5 calculations.

Under these conditions the protocol retains deliberate flexibility: designers may adjust the uncertainty threshold in light of downstream experimental costs, while the bias-correction architecture can be upgraded to more expressive models as resources permit [10, 20]. Because every stage remains modular and auditable, the approach seamlessly bridges academic inquiry and industrial deployment in AI-assisted materials discovery.

Relation to Existing Methods

The proposed multi-fidelity framework builds upon and extends several established computational approaches while addressing their specific limitations in the high-entropy alloy domain. Standard multi-fidelity Gaussian processes, as exemplified by Parussini *et al.* [22] and Tran *et al.* [19], typically assume a linear correlation between low-fidelity and high-fidelity outputs [10, 20]. This assumption works well for problems with smooth fidelity relationships but proves inadequate when CALPHAD biases vary nonlinearly across composition space. The present framework relaxes that linearity constraint by introducing an explicit, learnable bias-correction module that can employ neural networks or spatially aware Gaussian processes, thereby capturing the systematic discrepancies inherent to thermodynamic databases.

In relation to transfer learning strategies [25–29], conventional methods fine-tune a model pre-trained on abundant low-fidelity data toward a high-fidelity target. Such approaches treat CALPHAD outputs as a generic starting point for feature adaptation. The framework instead positions CALPHAD as a physically grounded prior whose known biases are corrected through paired training rather than implicit weight adjustment. This explicit correction preserves the thermodynamic interpretability of CALPHAD while achieving DFT-level accuracy without retraining the entire low-fidelity engine.

Active-learning literature has demonstrated the value of intelligent sampling for expensive DFT calculations [19]. Most existing active-learning pipelines, however, rely on a single surrogate model. Here, the corrected CALPHAD surface serves as the fast surrogate, allowing the acquisition function to operate at CALPHAD speed while still directing DFT effort toward regions where bias correction is most uncertain. The integration therefore yields a tighter feedback loop than DFT-only active learning.

Finally, the framework naturally accommodates Bayesian optimisation for property-targeted design. Once the bias-corrected predictions and uncertainty maps are available, standard Bayesian optimisation routines can be layered on top to identify optimal compositions rather than merely reducing global uncertainty. This compatibility positions the multi-fidelity approach as a general scaffold that augments rather than replaces existing optimisation toolkits [23]. By synthesising these threads—multi-fidelity regression, transfer-style correction, active sampling, and Bayesian decision-making—the framework offers a unified conceptual architecture tailored to the unique data asymmetry of high-entropy alloy design.

Challenges and Open Questions

Despite its conceptual promise, the multi-fidelity framework faces several substantive challenges that must be resolved for robust deployment. The first concerns the systematic nature of CALPHAD bias. Unlike random noise, these discrepancies vary with elemental combinations and local ordering tendencies [4, 11, 15]. Future bias-correction modules must therefore incorporate composition-dependent kernels or graph-based representations that respect the underlying thermodynamic topology.

A second challenge lies in data alignment. CALPHAD and DFT calculations often adopt subtly different reference states for elemental energies, leading to absolute-scale mismatches even when relative trends are consistent. Practical implementations will require automated alignment protocols—perhaps through a learned offset or a shared convex-hull normalisation—before paired training can begin.

Third, CALPHAD itself extrapolates poorly when entirely new element combinations appear. The framework must include an extrapolation detector that flags regions where the low-fidelity model operates far outside its calibration domain. In such cases, the uncertainty-propagation module should inflate its estimates dramatically, effectively forcing additional DFT investment.

Fourth, uncertainty calibration remains non-trivial. The correction model, trained on limited paired data, may produce overconfident or underconfident intervals. Techniques such as conformal prediction or ensemble-based recalibration, already explored in related materials contexts [19, 23], need to be adapted specifically for multi-fidelity HEA workflows to ensure that reported uncertainties are empirically reliable.

Fifth, the computational overhead of the active-learning loop cannot be ignored. Although each DFT call is minimised, the iterative nature still demands careful scheduling on shared high-performance resources. Open questions therefore centre on developing acquisition functions that balance information gain against wall-clock time and on exploring asynchronous or batch-active variants that further reduce calendar time. Addressing these challenges will require collaborative efforts across thermodynamics, machine learning, and high-performance computing communities.

Implications for HEA Discovery

The multi-fidelity framework carries immediate and far-reaching implications for three key stakeholder groups. For materials scientists, it transforms the design workflow from a DFT-constrained bottleneck into a CALPHAD-guided exploration engine [6, 12, 17, 18]. Researchers can now leverage existing, commercially available CALPHAD databases to screen millions of compositions and focus their limited DFT budget exclusively on bias-correction hotspots [4, 14]. This shift frees experimental resources for validation of the most promising candidates rather than exhaustive computational scouting.

For machine-learning researchers, the framework highlights fertile ground for methodological innovation. Developing more expressive bias-correction architectures—whether non-linear neural mappings or physics-informed graph networks—becomes a high-impact objective [8, 10]. Equally important is the creation of uncertainty-propagation schemes that faithfully decompose contributions from both fidelities. The community is also encouraged to establish public benchmark datasets of paired CALPHAD–DFT calculations for high-entropy alloys, analogous to existing materials benchmarks but explicitly designed for multi-fidelity evaluation.

For industry, the approach promises accelerated development cycles for structural, thermal, and energy-related high-entropy alloys. By slashing the required DFT count by roughly two orders of magnitude, the framework reduces both computational cost and time-to-discovery. Companies can therefore integrate the protocol into existing CALPHAD-centric pipelines with minimal disruption, ultimately lowering the barrier to commercialising novel alloys with tailored mechanical, magnetic, or catalytic properties. Collectively, these implications reposition high-entropy alloy design from a data-scarce art toward a data-fusion science.

Toward Fully Integrated HEA Design

The long-term vision is a self-consistent multi-fidelity ecosystem in which CALPHAD and machine learning continuously refine each other [14]. Rather than treating thermodynamic databases as static, the framework envisions an iterative loop in which newly acquired DFT data not only corrects predictions but also feeds back into CALPHAD parameter optimisation, gradually improving the low-fidelity model itself.

A practical roadmap supports this vision. In the short term (one to two years), linear or shallow non-linear bias correction combined with basic active learning can be implemented using off-the-shelf tools. Medium-term efforts (two to five years) should focus on scalable non-linear correction modules with full uncertainty propagation and automated alignment

routes. Over the longer horizon (five to ten years), the community can pursue a fully integrated CALPHAD–ML platform that automatically ingests new DFT results, updates thermodynamic parameters, and maintains a living uncertainty map across composition space [5, 6].

Success will be measured by the ability to discover experimentally validated high-entropy alloys with target properties using fewer than 500 DFT calculations in total. Achieving this benchmark would demonstrate that multi-fidelity learning has matured from a conceptual framework into a standard tool for complex concentrated alloy design.

Conclusion

Multi-fidelity learning bridges the persistent gap between sparse, accurate DFT data and abundant yet biased CALPHAD predictions for high-entropy alloy design. The conceptual framework presented here comprises five interlocking components—data integration, bias correction, uncertainty propagation, active learning, and validation—supported by a six-step operational protocol that begins with broad CALPHAD screening, proceeds through targeted DFT calibration, and converges when predictive uncertainty meets designer requirements. By learning an explicit, spatially aware mapping from CALPHAD to DFT values and propagating fidelity-specific uncertainties, the approach delivers near-DFT accuracy across million-composition spaces with only 100–300 DFT calculations instead of the tens of thousands required by single-fidelity campaigns.

The framework respects the fundamental cost hierarchy of computational materials science while harnessing the statistical strengths of machine learning. It is deliberately modular, allowing researchers to substitute alternative correction models or acquisition functions as new techniques emerge. Its implications extend beyond high-entropy alloys to any materials domain that possesses abundant low-fidelity thermodynamic or phenomenological models alongside sparse high-fidelity benchmarks.

Widespread adoption will require community-driven benchmarks, open-source reference implementations, and continued dialogue between thermodynamic database curators and machine-learning specialists. The present work therefore issues a call for collaborative development of multi-fidelity toolkits tailored to complex concentrated alloys. With such toolkits in hand, the materials community can move from exhaustive enumeration toward intelligent, uncertainty-guided discovery, ultimately unlocking the vast compositional potential of high-entropy alloys for next-generation technologies.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Published online: 18 January 2025

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Luo Q, Zhai C, Sun D, Chen W, Li Q. Interpolation and extrapolation with the CALPHAD method. *J Mater Sci Technol.* 2019;35(9):2115-20.
<https://doi.org/10.1016/j.jmst.2019.05.016>.
- Young DA. Phase diagrams of the elements. Berkeley (CA): University of California Press; 1991. 280 p.
- Liu X, Zhang J, Pei Z. Machine learning for high-entropy alloys: Progress, challenges and opportunities. *Prog Mater Sci.* 2023;131:101018.
<https://doi.org/10.1016/j.pmatsci.2022.101018>.
- Zeng Y, Man M, Bai K, Zhang YW. Revealing high-fidelity phase selection rules for high entropy alloys: A combined CALPHAD and machine learning study. *Mater Des.* 2021;202:109532.
<https://doi.org/10.1016/j.matdes.2021.109532>.
- Akbari MJ. Machine learning-driven design and discovery of high-entropy alloys: A critical review of recent advances and future perspectives. *Met Mater Int.* 2026;32:2258-76.
<https://doi.org/10.1007/s12540-025-02116-1>.
- Zhao YM, Zhang JY, Liaw PK, Yang T. Machine learning-based computational design methods for high-entropy alloys. *High Entropy Alloys Mater.* 2025;3(1):41-100.
<https://doi.org/10.1007/s44210-025-00055-5>.
- Batzner S, Musaelian A, Sun L, Geiger M, Mailoa JP, Kornbluth M, et al. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat Commun.* 2022;13(1):2453.
<https://doi.org/10.1038/s41467-022-29939-5>.
- Ko TW, Ong SP. Data-efficient construction of high-fidelity graph deep learning interatomic potentials. *npj Comput Mater.* 2025;11(1):65.
<https://doi.org/10.1038/s41524-025-01550-4>.
- Chowdhury NEEK, Jawad A, Rahman A, Khan MJA. Multi-fidelity neural network-based prediction of tensile strength of high-entropy alloy (FeNiCoCrCu) using molecular dynamics data. *J Mol Model.* 2025;31(8):214.
<https://doi.org/10.1007/s00894-025-06439-z>.
- Meng X, Karniadakis GE. A composite neural network that learns from multi-fidelity data: Application to function approximation and inverse PDE problems. *J Comput Phys.* 2020;401:109020.
<https://doi.org/10.1016/j.jcp.2019.109020>.
- Li W, Raman L, Debnath A, Ahn M, Lin S, Krajewski AM, et al. Design and validation of refractory alloys using machine learning, CALPHAD, and experiments. *Int J Refract Met Hard Mater.* 2024;121:106673.

<https://doi.org/10.1016/j.ijrmhm.2024.106673>.

Berry J, Snell R, Anderson M, Owen LR, Messé OMDM, Todd I, et al. Design and selection of high entropy alloys for hardmetal matrix applications using a coupled machine learning and calculation of phase diagrams methodology. *Adv Eng Mater.* 2024;26(10):2302064.

<https://doi.org/10.1002/adem.202302064>.

Debnath A, Reinhart WF. Overcoming sparse datasets with multi-task learning as applied to high entropy alloys. *Mach Learn Sci Technol.* 2025;6(1):015046.

<https://doi.org/10.1088/2632-2153/adb53c>.

Zhu S, Sarıtürk D, Arróyave R. Accelerating CALPHAD-based phase diagram predictions in complex alloys using universal machine learning potentials: Opportunities and challenges. *Acta Mater.* 2025;286:120747.

<https://doi.org/10.1016/j.actamat.2025.120747>.

Zhang C, Yang Y. The CALPHAD approach for HEAs: Challenges and opportunities. *MRS Bull.* 2022;47(2):158-67.

<https://doi.org/10.1557/s43577-022-00284-8>.

Li X, Chen X. Data-driven discovery of new eutectic high-entropy alloys via calculation of phase diagrams and machine learning integration. *Adv Eng Mater.* 2025;27(18):2500879.

<https://doi.org/10.1002/adem.202500879>.

Zeng Y, Man M, Ng CK, Aitken Z, Bai K, Wu D, et al. Search for eutectic high entropy alloys by integrating high-throughput CALPHAD, machine learning and experiments. *Mater Des.* 2024;241:112929.

<https://doi.org/10.1016/j.matdes.2024.112929>.

Liu F, Xiao X, Huang L, Tan L, Liu Y. Design of NiCoCrAl eutectic high entropy alloys by combining machine learning with CALPHAD method. *Mater Today Commun.* 2022;30:103172.

<https://doi.org/10.1016/j.mtcomm.2022.103172>.

Tran A, Tranchida J, Wildey T, Thompson AP. Multi-fidelity machine-learning with uncertainty quantification and Bayesian optimization for materials design: Application to ternary random alloys. *J Chem Phys.* 2020;153(7):074705.

<https://doi.org/10.1063/5.0015672>.

Boodaghizaji M, Khan M, Ardekani AM. Multi-fidelity modeling to predict the rheological properties of a suspension of fibers using neural networks and Gaussian processes. *Phys Fluids.* 2022;34(5):053101.

<https://doi.org/10.1063/5.0087449>.

Xu T, Zhang D, Xie Y, Chen Z. Multi-fidelity network framework for field and performance parameters prediction of turbomachinery based on deep graph learning. *Aerosp Sci Technol.* 2026;170:111530.

<https://doi.org/10.1016/j.ast.2025.111530>.

Parussini L, Venturi D, Perdikaris P, Karniadakis GE. Multi-fidelity Gaussian process regression for prediction of random fields. *J Comput Phys.* 2017;336:36-50.

<https://doi.org/10.1016/j.jcp.2017.01.047>.

Gantzer N, Deshwal A, Doppa JR, Simon CM. Multi-fidelity Bayesian optimization of covalent organic frameworks for xenon/krypton separations. *Digit Discov.* 2023;2(6):1937-56.

<https://doi.org/10.1039/D3DD00117B>.

Li Z, Tran ND, Sun Y, Lu Y, Hou C, Chen Y, et al. Cluster expansion augmented transfer learning for property prediction of high-entropy alloys. *J Mater Chem C.* 2025;13(34):17601-15.

<https://doi.org/10.1039/D5TC02311D>.

Yang F, Zhao W, Ru Y, Lin S, Huang J, Du B, et al. Transfer learning enables the rapid design of single crystal superalloys with superior creep resistances at ultrahigh temperature. *npj Comput Mater.* 2024;10(1):149.
<https://doi.org/10.1038/s41524-024-01349-9>.

Feng S, Zhou H, Dong H. Application of deep transfer learning to predicting crystal structures of inorganic substances. *Comput Mater Sci.* 2021;195:110476.
<https://doi.org/10.1016/j.commatsci.2021.110476>.

Zhao Y, Zhou H, Zhang Z, Bo Z, Sun B, Jiang M, et al. Discovering high-strength alloys via physics-transfer learning. *Matter.* 2025;8(9):102377.
<https://doi.org/10.1016/j.matt.2025.102377>.

Feng S, Fu H, Zhou H, Wu Y, Lu Z, Dong H. A general and transferable deep learning framework for predicting phase formation in materials. *npj Comput Mater.* 2021;7(1):10.
<https://doi.org/10.1038/s41524-020-00488-z>.

Varughese B, Manna S, Loeffler TD, Batra R, Cherukara MJ, Sankaranarayanan SKRS. Active and transfer learning of high-dimensional neural network potentials for transition metals. *ACS Appl Mater Interfaces.* 2024;16(16):20681-92.
<https://doi.org/10.1021/acsami.3c15399>.