Inverse Design of Novel Superconductors via Guided Diffusion

Anonymous Author(s)

Affiliation Address email

Abstract

The inverse design of materials with specific desired properties, such as hightemperature superconductivity, represents a formidable challenge in materials science due to the vastness of chemical and structural space. This work introduces a guided diffusion framework to accelerate the discovery of novel superconductors. Our approach begins with a foundation model trained on crystal structures from the Alexandria Database, which is then fine-tuned using a labeled dataset of 7,217 conventional superconductors to generate new structures conditioned on critical temperature, T_c . Employing classifier-free guidance, we generated 200,000 potential crystal structures. These candidates were subsequently subjected to a rigorous multi-stage computational screening workflow, utilizing machine learning models and density functional theory calculations to assess stability and electronic properties. Notably, our generative model demonstrated effective property-driven design by shifting the distribution of generated materials toward targeted $T_{\rm c}$ values within the training regime. This process successfully identified 773 promising superconducting candidates with predicted $T_c > 5K$. This end-to-end workflow, from generation to new candidate superconductors, illuminates a powerful pathway for materials discovery, demonstrating the significant potential of the AI-driven framework to accelerate discovery.

1 Introduction

2

3

4

6

8

9

10

11

12

13

14 15

16

17

18

- The discovery of novel materials with desired properties remains a fundamental challenge in materials science, as traditional methods involving direct simulation or experimental synthesis and characterization are frequently impeded by high costs, lengthy development cycles, and a low yield of successful outcomes. While ab initio methods and machine learning (ML) models have advanced the ability to predict material properties from known crystal structures, the inverse problem—designing materials that exhibit specific target properties—remains difficult due to the vastness of chemical and structural space.
- Generative models, which have achieved remarkable success in domains such as images, text, and video, are now gaining traction in materials science to address this problem. Recent advances in structure-generating models, such as diffusion-based frameworks Ho et al. [2020] (e.g., DiffCSP Jiao et al. [2023], MatterGen Zeni et al. [2024]), flow-based Lipman et al. [2023] models (e.g., FlowMM Miller et al. [2024]), and frameworks based on stochastic interpolants Albergo et al. [2023] that unify diffusion and flow such as OMatG Höllmer et al. [2025], highlight the potential of generative approaches for material discovery.
- Superconductors, with their zero electrical resistance, hold transformative potential for technologies ranging from energy transmission and storage to high-field applications in medical imaging, particle accelerators, and materials processing Boeri et al. [2022], Larbalestier et al. [2001], Malozemoff et al.

[2005], Flükiger [2012]. These wide-ranging applications make the discovery of new superconductors, particularly those with higher critical temperatures $T_{\rm c}$, a high-priority target for inverse materials design. In the most widely used class of superconductors, i.e. electron-phonon superconductors, $T_{\rm c}$ is governed by the coupling between electrons and lattice vibrations Bardeen et al. [1957], Eliashberg [1960]. Traditionally DFT Giannozzi et al. [2020] coupled to Eliashberg theory Giustino et al. [2007] has been used to accurately predict the critical temperatures of superconductors, albeit at too high a computational cost to be suitable for high-throughput screening.

Attempts have been made previously to discover superconductors with generative artificial intelligence (AI), notably the work by Wines et al. Wines et al. [2023] using the crystal diffusion variational autoencoder (CDVAE Xie et al. [2022]) framework and ALIGNN Choudhary and DeCost [2021] model for superconductivity prediction. While a novel proof-of-concept, their study has two key limitations. First, their property predictor (ALIGNN) was trained on a relatively small dataset of 1,058 structures, which may limit generalization of the inference capability of the model. Second, all generated structures were reported in the low-symmetry space group P1, which may bias the structural diversity and affect physical realizability. Another work by Dordevic et al. Yuan and Dordevic [2024] introduced SuperDiff, a diffusion-based generative model for superconductors that implements iterative latent variable refinement (ILVR) Choi et al. [2021] to condition the structure generation on reference compounds. While SuperDiff demonstrated fast training and inference and was the first to generate entirely new superconductor families via conditioning, it operates purely on composition-level data without explicitly modeling atomic positions or crystal structures, limiting its ability to capture the underlying atomic structure.

The field lacks a true generative model for superconductors to solve the inverse problem, and one crucial challenge is to tackle the problem of lack of a large conditioned database of superconductors that is sufficient for training a big foundation model. In this work, we build on recent advances in generative AI and present an enhanced DiffCSP framework that leverages guided diffusion to generate crystal structures conditioned on target $T_{\rm c}$ values. To overcome the challenge posed by the limited size of labeled datasets for specific properties like superconductivity, our strategy separates the tasks of learning general crystal structure formation and specific property conditioning. We first develop a foundation model by training DiffCSP on a vast dataset of over two million crystal structures with up to 20 atoms in the unit cell, drawn from the Alexandria Database Schmidt et al. [2022, 2023]. This initial phase allows the model to learn to generate stable and novel crystal structures, independent of any particular material property. Subsequently, we fine-tune this pre-trained foundation model using a smaller, high-quality dataset of 7,217 conventional superconductors with calculated $T_{\rm c}$ Cerqueira et al. [2024] as illustrated in Fig. 1(a).

This two-stage approach has the significant advantage of leveraging a massive dataset to learn structural plausibility, while requiring only a specialized dataset to learn the more subtle correlations that govern superconductivity. This method effectively guides the generation of novel crystal structures towards candidates that are both structurally stable and likely to be superconducting. We couple this approach to a subsequent screening framework that identifies the most-promising candidate superconductors. We apply this complete end-to-end discovery workflow by first using our guided diffusion framework to generate a large-scale library of 200,000 candidate crystal structures, as illustrated in Figure 1(a). These are then processed through a multi-stage computational screening pipeline that leverages machine learning models Chen et al. [2019], Chen and Ong [2022], Gibson et al. [2025] and density functional theory Giannozzi et al. [2020] to identify a subset of novel, stable candidates predicted to have higher $T_{\rm c}$, shown in Figure 1(b). This leads to prediction of 773 new superconducting candidates with $T_{\rm c} > 5$ K with $E_{\rm hull} < 200$ meV.

2 Methods

Our methodology for discovering novel superconductors integrates two primary components: (1)
a generative AI framework to propose novel crystal structures conditioned on superconducting
properties; (2) a high-throughput computational screening workflow driven by machine learning
models and ab initio calculations to assess the stability and promise of the generated candidates. The
components of our discovery pipeline, illustrated in Figure 1, is designed to first generate a vast
number of candidate structures and then systematically filter them to identify the most promising ones.
The generative AI stage itself involves three core steps: training a foundation model, fine-tuning it for
superconductivity, and guiding the generation process with classifier-free diffusion to create an initial

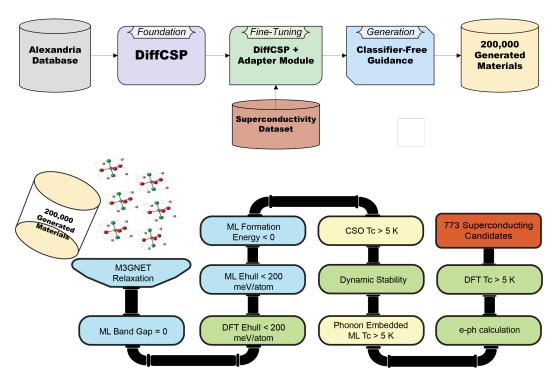


Figure 1: Workflow combining (a) guided diffusion and (b) multi-stage filtering to predict candidate superconductor. (a) Overview of the guided diffusion pipeline for superconductor discovery. Starting from the Alexandria database of 2,086,767 crystal structures, we train a foundation model based on the DiffCSP architecture, capable of generating stable crystals. An adapter module is added to DiffCSP architecture to fine-tune the foundation model using a dataset of 7183 superconductors Cerqueira et al. [2024], Gibson et al. [2025]. Using classifier free guidance, we use this AI to generate 200,000 crystal structures. (b) Overview of the multi-stage filtering process of the generated structures. We first pass the generated structures through initial ML relaxation, followed by search for metallic, thermodynamically stable, ML predicted $T_{\rm c} > 5$ K structures and $E_{hull} < 200$ meV/atom. We perform dynamic stability search on the remaining candidates, and calculated the stable structures' $T_{\rm c}$ with BEE-NET Gibson et al. [2025] using the PhDOS as embedding, selecting the $T_{\rm c} > 5$ K. The e-ph spectral function of the final candidates is evaluated using DFT, which is then used to get the final (DFT) $T_{\rm c}$ of the candidate structures.

ensemble of 200,000 structures (Fig. 1(a)). These candidates are then subjected to a rigorous, multistage filtering workflow (Fig. 1(b)) that uses a combination of machine learning interatomic potentials (MLIPs), the BEE-NET property predictor, and density functional theory (DFT) calculations to assess stability and predict T_c . The following subsections detail each step of this computational process.

96 2.1 Data

- We used two main datasets in this study: the Alexandria Materials Database, and a labeled dataset of superconductors Cerqueira et al. [2024] included the superconducting materials predicted in BEE-NET Gibson et al. [2025].
- The Alexandria Materials Database Schmidt et al. [2022, 2023] contains over 5 million crystal structures. To train our foundation model, we restricted the data to structures with up to 20 atoms per unit cell, giving us 2,086,767 structures, divided into subset of 1,857,222 crystal structures for training and another 229,545 for validation. This large, unlabeled dataset enables the model to learn the manifold of material space.
- For fine-tuning, we used the DS-A dataset from Cerqueira et al. Cerqueira et al. [2024], which comprises 7,217 dynamically stable metallic compounds with first-principles electron—phonon calculations. Each entry includes the relaxed crystal structure, density of states at the Fermi level, logarith-

mic average phonon frequency, and electron–phonon coupling constant. We used the Allen–Dynes formulaAllen and Dynes [1975] to estimate the superconducting transition temperature $T_{\rm c}$ of there materials. We end up with 6326 superconducting materials, a reduction due to the numerial instability from the Allen-Dynes formula. We use only the $T_{\rm c}$ values for fine-tuning, the dataset offers high-quality labels derived from the DFT-based workflows, making it well suited for guiding models toward superconductivity-relevant features. We extended this dataset by adding additional 857 superconducting candidates found in Gibson et al. [2025]– giving us dataset of size 7183 for fine-funing. The fine-tuning data was split into training, validation and test sets in 8:1:1 ratio.

2.2 Foundation Model Training

116

The first stage in our methodology is the development of a robust foundation model for crystal structure generation. The primary objective of this foundational training is to create a model adept at generating a diverse range of structurally and thermodynamically plausible crystalline materials, capturing the inherent chemical and physical rules governing crystal stability.

Training was performed using the DiffCSP framework, the architecture of which is detailed in 121 Section 2.3. During this unsupervised phase, the model learns to reverse a diffusion process that 122 gradually adds noise to the lattice parameters, fractional atomic coordinates, and atom types of known 123 crystal structures. By learning to effectively denoise these inputs, the model implicitly learns the complex manifold of crystal structures. This process is guided by the need to respect fundamental physical symmetries inherent in crystalline materials, such as rotational, translational and periodic 126 boundary conditions. The resulting foundation model is therefore well-versed in the general principles 127 128 of crystal structure formation and stability, without being biased towards any specific target property, providing a strong starting point for subsequent fine-tuning towards superconductor discovery. 129

130 2.3 DiffCSP Architecture

For crystal structure generation within our work, we employ a diffusion model. This generative approach involves a forward process where Gaussian noise is progressively added to the key attributes of a crystal structure - specifically its lattice parameters, the fractional coordinates of its constituent atoms, and their elemental types. The model is then trained to reverse this noising process, learning to denoise these attributes at various noise levels to reconstruct the original, physically valid crystal structure. New crystal structures are subsequently generated by iteratively applying this learned denoising function, beginning from a state of pure noise.

Generating physically realistic crystal structures critically requires the model's adherence to fundamental crystallographic symmetries. These include O(3) equivariance, which ensures that predictions are independent of crystal orientation by requiring consistent transformation of vectorial features (like atomic positions or lattice vectors) under rotation. Additionally, proper handling of periodic boundary conditions (PBC) is essential for accurately representing fractional atomic coordinates and ensuring the continuity of structural properties across unit cell boundaries.

To implement a diffusion process that rigorously incorporates these essential symmetries, we selected the Crystal Structure Prediction by Joint Equivariant Diffusion (DiffCSP) Jiao et al. [2023] framework as the backbone denoising architecture. DiffCSP employs an E(n) equivariant graph neural network (EGNN) Satorras et al. [2021] for the denoising model. This choice ensures O(3)-equivariance for handling lattice and coordinate symmetries and effectively manages periodic translational invariance for fractional atomic coordinates.

The input features consist of atomic embeddings and sinusoidal positional encoding that represent atom types and their positions within the crystal lattice. These features are processed through multiple message-passing layers, which incorporate both spatial and periodic symmetries. A Fourier transform is applied to the relative fractional coordinates, improving the model's ability to capture periodic behavior. The denoising model outputs two components: one for lattice noise and one for fractional coordinate noise. These outputs are designed to respect both O(3)-equivariance and translational invariance. Together, these architectural properties enable the model to robustly denoise inputs Ho et al. [2020], Song et al. [2021] in a way that preserves the physical symmetries of crystal structures.

In addition to lattice and coordinate modeling, atom types are also handled through learned embeddings/denoising, though detailed discussion of atom-type noise is omitted here for brevity. We kept the model hyperparameters consistent with those used in the original DiffCSP framework.

2.4 Fine-tuning for Conditional Generation

161

The second stage of training involves fine-tuning the foundation model to guide it toward generating superconducting materials. We use a labeled dataset of 7,000 conventional ambient condition superconductors whose $T_{\rm c}$ was calculated using ab initio methods Cerqueira et al. [2024]. We introduce an adapter module Zeni et al. [2024], inspired by text-to-image generation technique of computer vision, into the equivariant graph network (CSPNet) used in DiffCSP. This module allows us to condition generation on the target property, in this case, the superconducting critical temperature $(T_{\rm c})$. The adapter is applied after each message-passing layer as follows:

$$H_j^{'(L)} = H_j^{(L)} + f_{\text{mixin}}^{(L)}(f_{\text{adapter}}^{(L)}(g)) \cdot I(\text{property is not null})$$

Here, $H_j^{(L)}$ represents the node embeddings at layer L, and g is the sinusoidally embedded Vaswani et al. [2017] property value (T_c). The adapter module $f_{\rm adapter}^{(L)}$ consists of four fully connected linear layers, and its output is passed through a mixin layer Harrison and Ossher [2001] $f_{\rm mixin}^{(L)}$, which stabilizes training by gradually increasing the adapter's influence. This mechanism prevents abrupt disruptions to the foundation model's learned features. The mixin layer initially scales the adapter output to zero and increases it progressively during fine-tuning.

We retain the original loss function from DiffCSP, which ensures that generated structures remain

We retain the original loss function from DiffCSP, which ensures that generated structures remain close to the training distribution. The adapter module allows the model to incorporate property information during generation without compromising the stability learned during the foundation model training.

179 2.5 Classifier-free Guided Generation

To generate candidate superconducting materials, we implemented classifier-free diffusion Ho and Salimans [2021] guidance. This method enables conditional generation by interpolating between the unconditional foundation model and the fine-tuned, property-aware model. Specifically, we combine the denoising predictions from both models as follows:

$$\tilde{\epsilon}_{\theta}(z_{\lambda}, c) = (1 + w)\epsilon_{\theta}(z_{\lambda}, c) - w\epsilon_{\theta}(z_{\lambda})$$

Here, $\epsilon_{\theta}(z_{\lambda},c)$ is the denoising output of the conditional model guided by the target property c, and $\epsilon_{\theta}(z_{\lambda})$ is the output of the unconditional model. The guidance weight w controls the strength of the conditioning — higher values bias the generation more strongly toward the target property. We found w=2 to be a stable and effective choice in practice.

Using this guided generation strategy, we sampled 200,000 crystal structures by conditioning on different $T_{\rm c}$ values. These generated candidates were then passed through our structure analysis workflow to identify stable, high- $T_{\rm c}$ superconductors.

2.6 Structural Analysis Workflow

191

We follow a rigorous multi-step workflow to identify stable, high- $T_{\rm c}$ superconducting candidates from the set of generated structures. This pipeline ensures that final materials are metallic, thermodynamically and dynamically stable, and synthesizable. Before these generated structures enter the detailed screening pipeline, an initial filtering step is performed to ensure a focus on novel candidates. This involves removing any generated structures that are identical to those present in our fine-tuning dataset. Additionally, duplicate structures within the generated set itself are identified and reduced to unique instances. This pre-screening ensures that the subsequent computationally intensive analyses are concentrated on genuinely new potential superconductors.

First, we relax all remaining unique, novel generated structures using the M3GNet machine learning interatomic potential (MLIP) Chen and Ong [2022]. After relaxation, we use MEGNet Chen et al. [2019] to compute the bandgap $(E_{\rm g})$ and formation energy (E_f) . We retain only metallic structures $(E_{\rm g}=0)$ with negative E_f . Next, we estimate the superconducting critical temperature $(T_{\rm c})$ of these

filtered structures using BEE-Net Gibson et al. [2025]. We only keep materials with predicted $T_c > 5$ K for further analysis. To assess thermodynamic stability, we calculate the energy above the convex hull $(E_{\rm hull})$ using M3GNet and compare against the Materials Project database Jain et al. [2013]. Structures with $E_{\rm hull} > 200$ meV/atom are discarded. We further refine the surviving structures by recalculating $E_{\rm hull}$ using DFT-relaxed geometries and reapply the $T_c > 5$ K filter using BEE-Net. At this point, we incorporate phonon calculations to improve both dynamic stability assessment

At this point, we incorporate phonon calculations to improve both dynamic stability assessment and $T_{\rm c}$ accuracy. For each structure, we compute the phonon density of states (PhDOS) and use BEE-Net with coarse phonon density (CPD) embeddings to predict a more refined $T_{\rm c}$. Structures with imaginary phonon modes are eliminated. In the final stage, we perform electron-phonon coupling calculations using Quantum ESPRESSO to obtain the electron-phonon spectral function $\alpha^2 F(\omega)$ and compute the final DFT-based $T_{\rm c}$ values using the Allen-Dynes equation Allen and Dynes [1975].

215 3 Results

220

242

We first validate our computational approach by demonstrating that guidance allows for the generation of materials with desired superconducting critical temperatures. We then characterize the full set of generated structures, followed by the experimental synthesis and characterization of selected candidates to verify our theoretical predictions.

3.1 Validation of the Guided Diffusion model

We validated the performance of the guided diffusion model by assessing its ability to generate 221 structures conditioned on a target T_c . The guidance mechanism, based on classifier-free guidance Ho and Salimans [2021] (detailed in Sec. 2.5), is designed to balance this drive toward a specific T_c with the principles of structural stability learned by the foundation model (Sec. 2.2 and Sec. 2.4). 224 To efficiently analyze the distributions of predicted T_c values to validate our generative pipeline, 225 we estimated T_c using our screening workflow until the input of the crystal structure only (CSO) to 226 BEE-NET (Fig. 1 (b)). This approach provides near-immediate feedback without the computational 227 cost of ab initio calculations, which are reserved for the more comprehensive analysis workflow 228 described in Sec. 2.6. 229

For the analysis shown in Fig. 2 we generated 1000 structures for each configuration. Our test result illustrated in Fig. 2a first confirms that with guidance disabled (w = -1), the T_c distribution of 231 generated samples closely follows that of the fine-tuning dataset, establishing a baseline. In contrast, enabling classifier-free guidance (e.g., w=2) to target a specific $T_{\rm c}$ of 10 K resulted in a clear and 233 effective shift of the distribution towards the desired value, demonstrating property-driven control 234 (Fig. 2b). We further probed the model's behavior with an out-of-distribution target of $T_c = 110 \text{ K}$, a 235 regime absent from the fine-tuning data. Even with strong guidance, the model consistently failed 236 to produce stable high- T_c structures. Instead, as shown in Fig. 2c (w=2), the model preferentially 237 generates stable low- T_c structures. This behavior demonstrates that while guidance is effective 238 for targeted design within the training domain, the foundation model acts as a crucial prerequisite 239 for stability, constraining the generation to physically plausible structures for extreme property 240 extrapolation. 241

3.2 Prediction of Superconducting Candidates

The application of our full generative pipeline starts with 200,000 initial structures and incorporates the rigorous multi-stage structure analysis workflow (detailed in Section 2.6), yielding 773 unique candidate superconductors with DFT-calculated $T_c > 5$ K, indicating promising superconducting properties and thermodynamic stability conducive to synthesis. The relationship between predicted T_c and $E_{\rm hull}$ for these candidates is illustrated in Figure 3, showing a diverse set of materials spanning T_c values up to approximately 35 K.

A compositional analysis of these 773 candidates revealed a strong trend towards multi-component compounds of 133 binaries (17%), 455 ternaries (59%), 178 quaternaries (23%), and 7 pentanaries (1%). Among the final set of computationally identified superconducting candidates, we observed a significantly higher proportion of ternary structures compared to binary ones. We attribute this trend to the combined effect of our model's generative capabilities and the current state of superconductor research. The landscape of binary superconductors has been extensively investigated. In essence, for

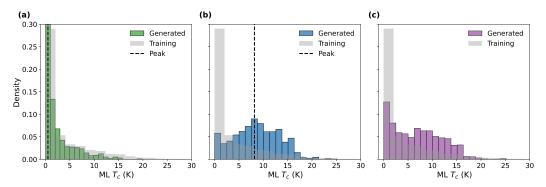


Figure 2: Effect of classifier-free guidance on $T_{\rm c}$ distribution. (a) Generated structures with guidance disabled (w=-1) closely follow the fine-tuning dataset distribution. (b) With guidance enabled (w=2) and target $T_{\rm c}=10$ K, the distribution shifts toward the desired value, demonstrating the model's controllability. (c) Distribution of predicted $T_{\rm c}$ values for structures generated with a target of $T_{\rm c}=110$ K using guidance weight w=2. Although the target lies far outside the training distribution, the model predominantly generates low- $T_{\rm c}$ structures, reflecting a preference for physically plausible outputs.

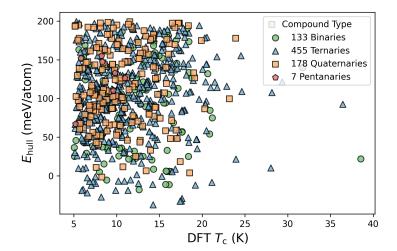


Figure 3: Distribution of the 773 predicted superconducting candidates. The plot shows the energy above the convex hull ($E_{\rm hull}$) versus the DFT-calculated critical temperature ($T_{\rm c}$) for each candidate. The markers are colored by compound type, as shown in the inset legend.

the workflow to identify new and viable superconductors, it is more probable that these will emerge from these more complex systems where a greater potential for undiscovered materials exists, rather than from the more saturated binary space.

4 Discussion

In this study, generative AI was used to create 200,000 crystal structures, which were subsequently screened using an ML driven workflow to arrive at 773 good candidates for superconductivity. Although the final superconducting candidates are in an acceptable regime of $E_{hull} < 200$ meV. This distribution of stability can likely be attributed to the training and fine-tuning data sets, which both consist of mainly theoretical materials that are not all on the convex hull. We believe that training on experimentally synthesized materials, and/or only materials with E_{hull} near zero, would guide the model to generate candidates that are, on average, closer to the hull.

Analysis of the compositions of these 773 candidates revealed a significantly higher proportion of ternary structures compared to binary ones. This outcome is attributed to the combined effect of our

pipeline search in the underexplored regions of the materials space and the prior investigation of the superconductor landscape. The comprehensive screening (as detailed in Section 2.6) naturally identified more novel candidates within the less-charted ternary, a higher-order compositional spaces where the potential for undiscovered materials is greater. This exploration of less-investigated chemical spaces was also reflected in the finding that many of the identified ternary candidates incorporate elements typically considered rare and more expensive. Such compositions may have been historically under-explored in experimental studies due to material costs, thus representing a richer domain for computational discovery.

However, a critical analysis suggests that this exploration of less-charted chemical spaces comes with a significant caveat that may explain our experimental outcomes. The model's prediction of stable phases in these regions could be an artifact of data scarcity, not an indication of true physical plausibility. For a predicted structure to be deemed stable, its energy must be low relative to a convex hull constructed from known, competing phases. In chemical spaces where data is sparse – for instance, among ternary systems whose constituent elements do not readily form any known binary compounds – the reference convex hull is incomplete or poorly defined. Consequently, a generated structure may appear to have a low energy above the hull ($E_{\rm hull}$) simply because there are no known, more stable structures to compete with it in the database. Our generative model, therefore, may be expertly identifying gaps in existing materials data rather than discovering genuinely synthesizable, stable phases. This highlights a fundamental challenge for AI-driven materials discovery: distinguishing true, synthesizable novelty from artifacts of an incomplete reference dataset.

Our current computational pipeline does not take into account the type of elements that are used in the predicted compounds. As we pointed out earlier, many of the ternary predicted superconductors have rare and expensive elements, which may pose some practical difficulties with the availability of the elements. Another problem is element that are either radioactive or toxic, making it difficult to work with within experimental laboratories such as Technetium.

294 5 Conclusion

In this work, we developed and deployed an end-to-end workflow that combines a guided diffusion generative model with high-throughput computational screening to accelerate the inverse design of novel superconductors. Our AI framework proved highly effective, successfully predicting over 700 unique and promising superconducting candidate materials with $T_{\rm c} \geq 5~{\rm K}$, by exploring less-charted regions of chemical space, particularly among ternary and higher-order compounds. However, our studies have revealed a critical challenge highlighting a fundamental insight: the generative model, while powerful, may be expertly identifying gaps in existing materials databases where the convex hull is sparse, rather than discovering genuinely synthesizable metastable phases. Therefore, while our work validates the immense potential of generative AI to accelerate materials discovery, it also underscores the need for future developments to focus on improving predictions of synthesizability and integrating experimental feedback into a true active learning loop to bridge the gap between computational prediction and laboratory realization.

Improvements to the generative model itself present a key direction for future work. We plan to improve the conditioning of the properties by implementing learnable embeddings for target properties, rather than the current fixed embedding approach, and exploring alternative guidance techniques Karras et al. [2024], Tang et al. [2025] that may offer better generalization and spread over the guided region than classifier-free guidance. To better reflect the discrete nature of elemental identities, we also intend to investigate guided discrete diffusion Nisonoff et al. [2025] models to generate atom types, which could offer improvements over denoising in a continuous latent space for atom types. Finally, we are working to integrate our guidance framework with entirely new generative backbones, such as Flow Matching Miller et al. [2024] and stochastic interpolants Höllmer et al. [2025], to assess their performance in discovering novel and synthesizable superconductors.

Our findings highlight several key avenues for future work, in order to address the challenge of sparse reference data in uncharted chemical spaces, the stability analysis could be improved by implementing on-the-fly DFT calculations to refine the local convex hull around promising candidates. Beyond refining the screening process, the most impactful advancement would be to establish a true active learning loop. By incorporating experimental synthesis as new training data, the generative model

could be systematically retrained to better discriminate between synthesizable and non-synthesizable 322 structures. This iterative process could be further enhanced by expanding the training datasets with 323 more experimentally verified compounds and by extending the guided diffusion framework to a 324 multi-objective optimization that balances the search for high T_c with synthesizability and practical 325 constraints, such as low material cost. A significant enhancement to our screening workflow would be 326 to incorporate predictions of disorder, disorder is completely ignored in the current pipeline, allowing 327 the model to assess the relative stability of a predicted ordered compound against a more likely 328 disordered solid solution. 329

330 References

- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In
 Advances in Neural Information Processing Systems, volume 33, pages 6840–6851. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper/2020/hash/
 4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html.
- Rui Jiao, Wenbing Huang, Peijia Lin, Jiaqi Han, Pin Chen, Yutong Lu, and Yang Liu. Crystal structure prediction by joint equivariant diffusion. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=DNdN26m2Jk.
- Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Sasha Shysheya, Jonathan Crabbé, Lixin Sun, Jake Smith, Bichlien Nguyen, Hannes Schulz, Sarah Lewis, Chin-Wei Huang, Ziheng Lu, Yichi Zhou, Han Yang, Hongxia Hao, Jielan Li, Ryota Tomioka, and Tian Xie. Mattergen: a generative model for inorganic materials design, 2024. URL https://arxiv.org/abs/2312.03687.
- Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=PqvMRDCJT9t.
- Benjamin Kurt Miller, Ricky T. Q. Chen, Anuroop Sriram, and Brandon M Wood. FlowMM:
 Generating materials with riemannian flow matching. In *Forty-first International Conference on Machine Learning*, 2024. URL https://openreview.net/forum?id=W4pB7VbzZI.
- Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions, 2023. URL https://arxiv.org/abs/2303.08797.
- Philipp Höllmer, Thomas Egg, Maya Martirossyan, Eric Fuemmeler, Zeren Shui, Amit Gupta, Pawan Prakash, Adrian Roitberg, Mingjie Liu, George Karypis, Mark Transtrum, Richard Hennig, Ellad B. Tadmor, and Stefano Martiniani. Open materials generation with stochastic interpolants. In Forty-second International Conference on Machine Learning, 2025. URL https://openreview.net/forum?id=gHGrzxFujU.
- Lilia Boeri, Richard Hennig, Peter Hirschfeld, Gianni Profeta, Antonio Sanna, Eva Zurek, War-356 ren E Pickett, Maximilian Amsler, Ranga Dias, Mikhail I Eremets, Christoph Heil, Russell J 357 Hemley, Hanyu Liu, Yanming Ma, Carlo Pierleoni, Aleksey N Kolmogorov, Nikita Rybin, Dmitry 358 Novoselov, Vladimir Anisimov, Artem R Oganov, Chris J Pickard, Tiange Bi, Ryotaro Arita, Ion Er-359 rea, Camilla Pellegrini, Ryan Requist, E K U Gross, Elena Roxana Margine, Stephen R Xie, Yundi 360 Quan, Ajinkya Hire, Laura Fanfarillo, G R Stewart, J J Hamlin, Valentin Stanev, Renato S Gonnelli, 361 Erik Piatti, Davide Romanin, Dario Daghero, and Roser Valenti. The 2021 room-temperature 362 superconductivity roadmap. Journal of Physics: Condensed Matter, 34(18):183002, mar 2022. 363 doi: 10.1088/1361-648X/ac2864. URL https://dx.doi.org/10.1088/1361-648X/ac2864. 364
- David Larbalestier, Alex Gurevich, D. Matthew Feldmann, and Anatoly Polyanskii. High- $T_{\rm c}$ superconducting materials for electric power applications. *Nature*, 414(6861):368–377, 2001. doi: 10.1038/35104654.
- Alexis P. Malozemoff, Jochen Mannhart, and Douglas Scalapino. High-temperature cuprate superconductors get to work. *Physics Today*, 58(4):41–47, 04 2005. ISSN 0031-9228. doi: 10.1063/1.1955478. URL https://doi.org/10.1063/1.1955478.

- Rene Flükiger. Overview of superconductivity and challenges in applications. *Reviews of Accelerator Science and Technology*, 05:1–23, 2012. doi: 10.1142/S1793626812300010. URL https://doi.org/10.1142/S1793626812300010.
- J. Bardeen, L. N. Cooper, and J. R. Schrieffer. Theory of superconductivity. *Phys. Rev.*, 108:
 1175-1204, Dec 1957. doi: 10.1103/PhysRev.108.1175. URL https://link.aps.org/doi/10.1103/PhysRev.108.1175.
- G M Eliashberg. Interactions between electrons and lattice vibrations in a superconductor. *Sov. Phys. JETP (Engl. Transl.); (United States)*, 9 1960. URL https://www.osti.gov/biblio/7354388.
- Paolo Giannozzi, Oscar Baseggio, Pietro Bonfà, Davide Brunato, Roberto Car, Ivan Carnimeo,
 Carlo Cavazzoni, Stefano de Gironcoli, Pietro Delugas, Fabrizio Ferrari Ruffino, Andrea Ferretti,
 Nicola Marzari, Iurii Timrov, Andrea Urru, and Stefano Baroni. Quantum espresso toward the
 exascale. *The Journal of Chemical Physics*, 152(15):154105, 04 2020. ISSN 0021-9606. doi:
 10.1063/5.0005082. URL https://doi.org/10.1063/5.0005082.
- Feliciano Giustino, Marvin L. Cohen, and Steven G. Louie. Electron-phonon interaction using wannier functions. *Phys. Rev. B*, 76:165108, Oct 2007. doi: 10.1103/PhysRevB.76.165108. URL https://link.aps.org/doi/10.1103/PhysRevB.76.165108.
- Daniel Wines, Tian Xie, and Kamal Choudhary. Inverse design of next-generation superconductors using data-driven deep generative models. *The Journal of Physical Chemistry Letters*, 14(29): 6630–6638, July 2023. ISSN 1948-7185. doi: 10.1021/acs.jpclett.3c01260. URL http://dx.doi.org/10.1021/acs.jpclett.3c01260.
- Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi S. Jaakkola. Crystal diffusion variational autoencoder for periodic material generation. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=03RLpj-tc_.
- Kamal Choudhary and Brian DeCost. Atomistic line graph neural network for improved materials property predictions. *npj Computational Materials*, 7(1):185, 2021. doi: 10.1038/s41524-021-00650-1. URL https://www.nature.com/articles/s41524-021-00650-1.
- Samuel Yuan and S. V. Dordevic. Diffusion models for conditional generation of hypothetical new families of superconductors. *Scientific Reports*, 14(1):10275, 2024. doi: 10.1038/s41598-024-61040-3. URL https://www.nature.com/articles/s41598-024-61040-3.
- Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr:
 Conditioning method for denoising diffusion probabilistic models. In 2021 IEEE/CVF International
 Conference on Computer Vision (ICCV), pages 14347–14356, 2021. doi: 10.1109/ICCV48922.
 2021.01410.
- Jonathan Schmidt, Hai-Chen Wang, Tiago F. T. Cerqueira, Silvana Botti, and Miguel A. L. Marques.

 A dataset of 175k stable and metastable materials calculated with the pbesol and scan functionals.

 Scientific Data, 9(1), March 2022. ISSN 2052-4463. doi: 10.1038/s41597-022-01177-w. URL

 http://dx.doi.org/10.1038/s41597-022-01177-w.
- Jonathan Schmidt, Noah Hoffmann, Hai-Chen Wang, Pedro Borlido, Pedro J. M. A. Carriço, Tiago
 F. T. Cerqueira, Silvana Botti, and Miguel A. L. Marques. Machine-learning-assisted determination
 of the global zero-temperature phase diagram of materials. *Advanced Materials*, 35(22), April
 2023. ISSN 1521-4095. doi: 10.1002/adma.202210788. URL http://dx.doi.org/10.1002/adma.202210788.
- Tiago F. T. Cerqueira, Antonio Sanna, and Miguel A. L. Marques. Sampling the materials space for conventional superconducting compounds. *Advanced Materials*, 36(1):2307085, 2024. doi: https://doi.org/10.1002/adma.202307085. URL https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/adma.202307085.
- Chi Chen, Weike Ye, Yunxing Zuo, Chen Zheng, and Shyue Ping Ong. Graph networks as a universal machine learning framework for molecules and crystals. *Chemistry of Materials*, 31(9):3564–3572, 2019. doi: 10.1021/acs.chemmater.9b01294. URL https://doi.org/10.1021/acs.chemmater.9b01294.

- Chi Chen and Shyue Ping Ong. A universal graph deep learning interatomic potential for the periodic
 table. *Nature Computational Science*, 2:718–728, 2022. doi: 10.1038/s43588-022-00349-3. URL
 https://www.nature.com/articles/s43588-022-00349-3.
- Jason B. Gibson, Ajinkya C. Hire, Philip M. Dee, Benjamin Geisler, Jung Soo Kim, Zhongwei Li,
 James J. Hamlin, Gregory R. Stewart, P. J. Hirschfeld, and Richard G. Hennig. Developing a
 complete ai-accelerated workflow for superconductor discovery, 2025. URL https://arxiv.org/abs/2503.20005.
- P. B. Allen and R. C. Dynes. Transition temperature of strong-coupled superconductors reanalyzed.

 Phys. Rev. B, 12:905-922, Aug 1975. doi: 10.1103/PhysRevB.12.905. URL https://link.aps.org/doi/10.1103/PhysRevB.12.905.
- Victor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E(n) equivariant graph neural networks. *CoRR*, abs/2102.09844, 2021. URL https://arxiv.org/abs/2102.09844.
- Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 Poole. Score-based generative modeling through stochastic differential equations, 2021. URL
 https://arxiv.org/abs/2011.13456.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://arxiv.org/abs/1706.03762.
- William Harrison and Harold Ossher. Mixin layers: An object-oriented implementation technique for
 refinements and collaboration-based designs. ACM Transactions on Software Engineering and
 Methodology (TOSEM), 11(2):205–255, 2001.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop*on Deep Generative Models and Downstream Applications, 2021. URL https://openreview.net/forum?id=qw8AKxfYbI.
- Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen
 Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, and Kristin A Persson.
 Commentary: The materials project: A materials genome approach to accelerating materials
 innovation. APL Materials, 1(1):011002, 2013. doi: 10.1063/1.4812323.
- Tero Karras, Miika Aittala, Tuomas Kynkäänniemi, Jaakko Lehtinen, Timo Aila, and Samuli Laine.
 Guiding a diffusion model with a bad version of itself. In *The Thirty-eighth Annual Conference*on Neural Information Processing Systems, 2024. URL https://openreview.net/forum?id=
 bg6fVPVs3s.
- Zhicong Tang, Jianmin Bao, Dong Chen, and Baining Guo. Diffusion models without classifier-free guidance, 2025. URL https://arxiv.org/abs/2502.12154.
- Hunter Nisonoff, Junhao Xiong, Stephan Allenspach, and Jennifer Listgarten. Unlocking guidance for discrete state-space diffusion and flow models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=XsgH154y07.