AI for Scalable Material Design and Property Prediction

Kishalay Das
Indian Institute of Technology (IIT), Kharagpur, India

Introduction & Motivation

My research focuses on AI for Scalable Material Design and Property Prediction, an interdisciplinary area at the intersection of Computer Science, Chemistry, and Geometry. Screening 3D periodic structures and their atomic compositions to identify novel crystal materials with specific chemical properties remains a long-standing challenge in the materials design community. These materials have been fundamental to key innovations such as the development of batteries, solar cells, semiconductors, quantum computers, etc. [Butler et al., 2018, Desiraju, 2002, Kohn and Sham, 1965]. A major challenge in materials design is discovering novel crystal structures with desired chemical properties that are both stable and physically realizable. Typically, this is done through a two-step process: first, the generation of 3D periodic structures for new materials (Generative Task), followed by the evaluation of various chemical properties to confirm their validity (Discriminative Task).
Against this backdrop, my current research focuses on developing advanced AI frameworks targeting both aforementioned generative and discriminative tasks within the domain of materials science. Specifically, for the generative task, I explored text-guided joint diffusion models designed to generate valid and stable 3D crystal structures from natural language descriptions. This approach enables intuitive material design by bridging human language and complex material representations. In contrast, for the discriminative task, my research concentrated on building efficient and accurate crystal property prediction models. I developed frameworks based on pretrained Graph Neural Networks (GNNs) as well as multi-modal architectures that integrate diverse sources of information. These models significantly accelerate the process of property prediction, aiding rapid materials discovery and screening.

Text Guided Joint Diffusion Model for Periodic Material Generation [ICLR 2025]

Crystal material can be modeled by a minimal *unit cell*, which gets repeated infinite times in 3D space on a regular lattice to form the periodic crystal structure. Formally, a crystal material $\mathbf{M}$ can be represented using three core components: the atom type matrix ($\mathbf{A}$), the coordinate matrix ($\mathbf{X}$), and the lattice matrix ($\mathbf{L}$). In our work **TGDMat [ICLR-2025]** [Das et al., 2025], we propose a novel generative framework that jointly models and generates $\mathbf{A}$, $\mathbf{X}$, and $\mathbf{L}$ using a periodic E(3)-equivariant denoising diffusion model. This design ensures that the learned generative process adheres to the periodic E(3) invariance properties inherent to crystal structures. Building on this foundation, we further incorporate textual conditioning into the reverse diffusion process. This allows the model to generate crystal structures that align with given textual descriptions, effectively guiding the denoising trajectory using semantic cues. Such text guidance introduces a powerful modality for controllable generation and opens up new possibilities for materials discovery based on natural language specifications. Comprehensive experiments on standard benchmark datasets across multiple tasks demonstrate that TGDMat consistently outperforms existing baseline methods by a significant margin. In particular, for the structure prediction task (CSP), TGDMat achieves superior performance with just a single generated sample, highlighting the strength of text-guided generation. Moreover, in the random/de novo generation task (Gen), TGDMat outperforms all baselines—including their text-fusion variants—underscoring the effectiveness of our joint diffusion paradigm. Importantly, the integration of textual information not only improves generative performance but also leads to reduced training and inference costs, showcasing the practical benefits of our approach in terms of both efficiency and accuracy.

LLM-Guided Diffusion Models for Material Generation [NeurIPS 2025]

Recent generative models have emerged as promising alternatives, falling into two categories: **Large Language Models (LLMs):** [Antunes et al., 2024, Gruver et al., 2024] Strong at predicting discrete atomic types and ensuring compositional validity but weak in handling continuous variables like coordinates and lattices. **Denoising-based frameworks (diffusion, flow matching):** [Jiao et al., 2023, Luo et al., 2023, Xie et al., 2021, Zeni et al., 2023] Good at generating continuous variables with high structural validity but poor at capturing discrete atomic compositions. To address these complementary weaknesses, we introduces CrysLLMGen, a hybrid framework that combines LLMs and diffusion models. The LLM first predicts atom types, atomic coordinates, and lattice parameters. The atom types are retained, while coordinates and lattices are refined by a pre-trained equivariant diffusion model. This design leverages the LLM's strength in composition and the diffusion model's ability to refine geometry. The approach yields balanced improvements in structural and compositional validity, producing significantly more stable crystal materials than either standalone LLMs or diffusion models. It also supports conditional generation (e.g., user-defined compositions or space groups). The framework is architecture-agnostic and can adapt to future LLM and diffusion model advances.

GNN Pretraining Model for Property Prediction [NPJ Comp. Mat. 2022, AAAI 2023]

One fundamental problem in materials science is fast and accurate prediction of different properties of crystal materials. Although existing GNN models predict different crystal properties with high precision, they suffer from the following major limitations like *Scarcity of Labeled Data, Lack of Interpretability, Dependency on Domain Knowledge, DFT Error Bias and Lack of Pre-trained GNN Model.*To mitigate these issues and learning more robust and enriched crystal representation for property prediction, in our first work, we leverage a transfer learning-based unsupervised framework to develop an explainable property predictor, **CrysXPP [NPJ Computational-2022] Das et al. [2022]**. We further scale this idea by developing **CrysGNN [AAAI-2023]** Das et al. [2023b], a deep pre-trained GNN framework for crystalline materials trained on a large corpus of unlabeled data. Our graph pre-training approach captures (a) atomic connectivity, (b) atomic properties, and (c) graph-level similarity. To support this, we curate a dataset of 800K unlabeled crystal graphs and pre-train CrysGNN on it. CrysGNN is designed to be integrated into any state-of-the-art property predictor. Using knowledge distillation, we transfer essential structural and chemical information from the pre-trained model to downstream tasks, significantly boosting both accuracy and efficiency in property prediction.

Crystal Multi-Modal Representation for Crystal Property Prediction [UAI 2023]

In this workDas et al. [2023a], we propose to learn a more robust and enriched representation by using multi-modal data i.e graph structure and textual description of materials. One of the major advantages of using the textual description of materials is it provides a diverse set of periodic structural information which is useful to estimate different crystal properties but difficult to incorporate explicitly into a graph structure. We first curate the textual dataset of two popular materials databases (Graph-based), Material Project (MP) and JARVIS-DFT, containing textual descriptions of each material of those databases. We used a popular tool Robocrystallographer Ganose and Jain [2019] to generate descriptions for global crystal structures that include space group number, crystal symmetry, rotational information, component orientations, heterostructure information, etc.
We propose, **CrysMMNet [UAI-2023] Das et al. [2023a]**, a simple multi-modal framework for crystal materials, which has two components: Graph Encoder and Text Encoder. Given a material, a Graph Encoder uses its graph structure and applies a GNN-based approach to encode the local neighborhood structural information around a node (atom), and subsequently learn graph (crystal) representation. On the contrary, Text Encoder is a transformer-based model, which encodes the global structural knowledge from the textual description of the material and generates a textual representation. Finally, both graph structural and textual representation are fused together to generate a more enriched multimodal representation of materials, which captures both global and local structural knowledge and subsequently improves property prediction accuracy. We observe that CrysMMNet outperforms all the popular state-of-the-art baselines across ten diverse sets of properties on two popular datasets, Materials Project and JARVIS-DFT.

Future Research Directions

Robust and Efficient AI Frameworks for Scalable Material Design

In the pursuit of scalable material design, I am eager to further explore following: I aim to accelerate diffusion-based material generation by reducing the number of sampling steps, enabling large-scale and computationally efficient synthesis of crystal structures. I plan to advance inverse material generation, where models design materials based on desired properties using text-guided conditioning and large language models for property-aware generation. My long-term goal is to develop a foundation model for periodic materials that unifies generative and discriminative capabilities within a single large-scale framework for versatile material discovery and analysis.

Explore Generative and Multimodal AI for Scientific Discovery

Apart from exploring scalable material design, I am keen to extend my research to other domains of AI for Science, including molecules, proteins, and biological systems, where similar structural and relational complexities exist. I am deeply interested in exploring generative models, multimodal frameworks and and agentic AI for scientific discovery. My goal is to develop efficient and interpretable generative models for these domains and understand the mathematical foundations that make certain algorithms particularly effective for biological applications. In the long term, I aspire to build foundation models for biology—large, multimodal architectures capable of integrating molecular, structural, and phenotypic data for unified predictive and generative tasks.

References

Luis M Antunes, Keith T Butler, and Ricardo Grau-Crespo. Crystal structure generation with autoregressive large language modeling. *Nature Communications*, 15(1):1–16, 2024.

Keith T Butler, Daniel W Davies, Hugh Cartwright, Olexandr Isayev, and Aron Walsh. Machine learning for molecular and materials science. *Nature*, 559(7715):547–555, 2018.

Kishalay Das, Bidisha Samanta, Pawan Goyal, Seung-Cheol Lee, Satadeep Bhattacharjee, and Niloy Ganguly. Crysxpp: An explainable property predictor for crystalline materials. *npj Computational Materials*, 8(1):43, 2022.

Kishalay Das, Pawan Goyal, Seung-Cheol Lee, Satadeep Bhattacharjee, and Niloy Ganguly. Crysmmnet: multimodal representation for crystal property prediction. In *Uncertainty in Artificial Intelligence*, pages 507–517. PMLR, 2023a.

Kishalay Das, Bidisha Samanta, Pawan Goyal, Seung-Cheol Lee, Satadeep Bhattacharjee, and Niloy Ganguly. Crysgnn: Distilling pre-trained knowledge to enhance property prediction for crystalline materials. *arXiv preprint arXiv:2301.05852*, 2023b.

Kishalay Das, Subhojyoti Khastagir, Pawan Goyal, Seung-Cheol Lee, Satadeep Bhattacharjee, and Niloy Ganguly. Periodic materials generation using text-guided joint diffusion model. *arXiv preprint arXiv:2503.00522*, 2025.

Gautam R Desiraju. Cryptic crystallography. *Nature materials*, 1(2):77–79, 2002.

Alex M Ganose and Anubhav Jain. Robocrystallographer: automated crystal structure text descriptions and analysis. *MRS Communications*, 9(3):874–881, 2019.

Nate Gruver, Anuroop Sriram, Andrea Madotto, Andrew Gordon Wilson, C Lawrence Zitnick, and Zachary Ulissi. Fine-tuned language models generate stable inorganic materials as text. *arXiv preprint arXiv:2402.04379*, 2024.

Rui Jiao, Wenbing Huang, Peijia Lin, Jiaqi Han, Pin Chen, Yutong Lu, and Yang Liu. Crystal structure prediction by joint equivariant diffusion. *arXiv preprint arXiv:2309.04475*, 2023.

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

Youzhi Luo, Chengkai Liu, and Shuiwang Ji. Towards symmetry-aware generation of periodic materials. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL `https://openreview.net/forum?id=Jkc74vn1aZ`.

Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi Jaakkola. Crystal diffusion variational autoencoder for periodic material generation. *arXiv preprint arXiv:2110.06197*, 2021.

Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Sasha Shysheya, Jonathan Crabbé, Lixin Sun, Jake Smith, et al. Mattergen: a generative model for inorganic materials design. *arXiv preprint arXiv:2312.03687*, 2023.