Improve Molecular Conformation Modeling with Geometric Deep Learning

Fanmeng Wang

Gaoling School of Artificial Intelligence, Renmin University of China
fanmengwang@ruc.edu.cn

Abstract

Molecular conformations, the stable three-dimensional structures corresponding to local minima on the potential energy surface, govern key molecular properties and consequently underpin a wide range of downstream tasks. However, contemporary learning-based methods often lack scalability, interpretability, and robustness, thereby significantly constraining their practical effectiveness and reliability. In this context, I will introduce my ongoing explorations and the proposed research plan to address these challenges, with the ultimate objective of developing conformation-centric universal foundation models to accelerate scientific discovery.

Introduction

While learning-based methods are developed for modeling molecular conformations, they still face several fundamental challenges that significantly limit their performance. Firstly, existing molecular conformation generation methods typically generate the entire structure at once, inexorably undermining scalability and thus becoming particularly problematic for macromolecules. Secondly, while existing molecular conformation optimization methods refine easily obtained low-quality conformations into high-quality ones to improve efficiency, they do not correspond to minimizing a physically reasonable energy function, thereby leading to poor interpretability. Finally, designing robust molecular conformation perception methods that infer molecular attributes (e.g., bond orders and formal charges) from conformations is also essential, as conformations are inherently noisy. In addition, effectively leveraging structural information embedded in conformations to improve downstream tasks remains a major concern. Given these challenges, geometric deep learning, which respects molecular structure and symmetries, has emerged as a promising avenue.

Previous and Current Work

My ultimate goal is to develop conformation-centric universal foundation models that effectively support diverse tasks, thereby significantly accelerating scientific discovery. Toward this goal, I am focusing on developing **scalable**, **interpretable**, and **robust** algorithms to improve molecular conformation modeling, while leveraging conformational information to boost downstream tasks.

Scalable Molecular Conformation Generation. While conformation generation has been widely explored over the past few years, existing methods primarily focused on small molecules (with fewer than 100 atoms) and typically generate the entire conformation through a single process, which leads to poor scalability and rapid performance degradation for macromolecules. In this context, I developed PolyConf (Wang et al. 2025), which leverages hierarchical generative models to achieve scalable conformation generation.

Specifically, since the complete conformation can be decomposed into a sequence of local conformations, PolyConf generates these local conformations by integrating a masked autoregressive model with a diffusion model design for conformation generation, and then further generates the required orientation transformations for all local conformations to assemble the complete conformation.

Comprehensive experiments demonstrate that PolyConf not only achieves superior scalability but also significantly outperforms existing conformation generation methods in both quality and efficiency, thereby enabling conformation generation for macromolecules.

In addition, building on PolyConf, I am currently designing conformation-centric generative pretraining paradigms to unify property prediction and generation, thereby establishing a unified backbone for diverse downstream tasks.

Interpretable Molecular Conformation Optimization. Due to the ready availability of low-quality conformations, there is growing interest in improving efficiency by optimizing them rather than generating conformations from scratch. However, existing methods, which employ neural networks to predict energy gradients and then emulate gradient descent to update input conformations, lack alignment with the minimization of a physically reasonable energy function, resulting in reduced interpretability and suboptimal performance. In this context, I developed WGFormer (Wang, Cheng, and Xu 2025), which leverages Wasserstein gradient flow to achieve interpretable conformation optimization.

Specifically, WGFormer modifies the SE(3)-Transformer by adjusting the query, key, and value matrices and applying the Sinkhorn-scaling algorithm to compute attention maps. Furthermore, conformation optimization is performed within an auto-encoding framework, where WGFormer en-

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

codes low-quality conformations and an MLP decodes corresponding high-quality conformations.

Theoretical analysis demonstrates that WGFormer operates as Wasserstein gradient flows, corresponding to the minimization of an explicit energy function defined on the latent mixture models of atoms. Experimental results further validate that minimizing this energy function indeed drives potential energy minimization, thereby significantly improving performance and interpretability.

Robust Molecular Conformation Perception. As conformations are essentially sets of atom types and 3D coordinates, the capability of conformation perception methods, serving as the post-processing module that infers additional molecular attributes (e.g., bond orders and formal charges) from conformations, becomes particularly important. However, existing methods that rely on predefined chemical rules are susceptible to conformational precision, and conformations (particularly those generated by deep generative models) are often noisy, resulting in substantial performance degradation. In this context, I developed MPerformer (Wang et al. 2023), which leverages SE(3)-Transformer under the multi-task pretraining-and-finetuning learning paradigm to achieve robust conformation perception.

Specifically, recognizing that all molecular attributes are SE(3)-invariant, MPerformer adopts SE(3)-Transformer as its backbone module to align inductive biases with physical symmetries. Furthermore, MPerformer jointly learns to denoise conformations and predict molecular attributes during pretraining, and focuses solely on attribute prediction during finetuning, yielding the final model.

Experiments consistently indicate that MPerformer exhibits excellent robustness while achieving state-of-the-art performance. Moreover, integrating MPerformer as the post-processing module for 3D molecular generators significantly improves the quality of generated molecules, thereby underscoring its great potential for practical applications.

Conformation-Enabled Improvement of Downstream Tasks. Since molecular function is essentially governed by three-dimensional structure, many works are incorporating conformational information into downstream tasks to improve performance. However, this also increases computational complexity, memory demands, and stability concerns, particularly for macromolecules such as polymers. In this context, I developed MMPolymer (Wang et al. 2024), which leverages local conformations to approximate entire conformations through the multimodal multitask pretraining framework to improve downstream tasks.

Specifically, considering the regularity and complexity of entire conformations, MMPolymer converts corresponding local conformations (i.e., conformations of repeating units) through the specially-designed "Star Substitution" strategy, where the "*" symbol is replaced by the neighboring atom symbol of another "*" symbol, to approximate entire conformations. Furthermore, MMPolymer not only predicts masked tokens and recovers clear conformations but also achieves cross-modal alignment during pretraining.

Downstream experiments demonstrate that MMPolymer consistently outperforms existing methods on various property prediction tasks. Moreover, relying on single-modal information during fine-tuning, MMPolymer can still achieve state-of-the-art performance, thereby substantiating the intrinsic value of conformations for downstream tasks.

In addition to MMPolymer, I was also the principal contributor to ToxScan (Wu et al. 2025), an SE(3)-equivariant multiscale model that functions as a universal framework for toxicity prediction. It incorporates conformational information through a two-level molecular and atomic representation learning paradigm, achieving robust generalization and high-accuracy predictions across diverse toxicity endpoints.

Future Work

My research focuses on molecular conformation modeling, developing a series of methods that improve modeling and downstream tasks. Building on these foundations, I will push forward along two complementary axes:

Model Architecture. Existing model architectures designed for conformations typically account for information between atom pairs, resulting in substantial computational complexity that limits scalability. In this context, I will pursue lightweight equivariant architectures with sparsified or hierarchical neighborhoods that approximate all-pairs interactions to retain 3D fidelity while reducing the nominal $O(N^2)$ cost toward linear or near-linear scaling.

Learning Paradigm. In the absence of universal foundation models that unify both molecular understanding and generation, I will pursue conformation-centric unified backbones that learns geometry-aware representations and support a broad spectrum of downstream tasks.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (92270110), the Fundamental Research Funds for the Central Universities, and the Research Funds of Renmin University of China.

References

Wang, F.; Cheng, M.; and Xu, H. 2025. WGFormer: An SE (3)-Transformer Driven by Wasserstein Gradient Flows for Molecular Ground-State Conformation Prediction. In *Forty-second International Conference on Machine Learning*.

Wang, F.; Guo, W.; Cheng, M.; Yuan, S.; Xu, H.; and Gao, Z. 2024. MMPolymer: A Multimodal Multitask Pretraining Framework for Polymer Property Prediction. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*.

Wang, F.; Guo, W.; Ou, Q.; Wang, H.; et al. 2025. PolyConf: Unlocking Polymer Conformation Generation through Hierarchical Generative Models. In *Forty-second International Conference on Machine Learning*.

Wang, F.; Xu, H.; Chen, X.; Lu, S.; et al. 2023. MPerformer: An SE(3) Transformer-based Molecular Perceptron. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*.

Wu, L.; Wang, F.; Zhang, Y.; Li, R.; Zhao, Y.; et al. 2025. Multiscale-Information-Embedded Universal Toxicity Prediction Framework. *Environmental Science & Technology*.