
Efficient Generation of Molecular Clusters with Dual-Scale Equivariant Flow Matching

Akshay Subramanian
Massachusetts Institute of Technology
akshay_s@mit.edu

Shuhui Qu
Samsung Display America Lab

Cheol Woo Park
Samsung Display America Lab

Sulin Liu
Massachusetts Institute of Technology

Janghwan Lee
Samsung Display America Lab

Rafael Gómez-Bombarelli
Massachusetts Institute of Technology
rafagb@mit.edu

Abstract

Amorphous molecular solids offer a promising alternative to inorganic semiconductors, owing to their mechanical flexibility and solution processability. The packing structure of these materials plays a crucial role in determining their electronic and transport properties, which are key to enhancing the efficiency of devices like organic solar cells (OSCs). However, obtaining these optoelectronic properties computationally requires molecular dynamics (MD) simulations to generate a conformational ensemble, a process that can be computationally expensive due to the large system sizes involved. Recent advances have focused on using generative models, particularly flow-based models as Boltzmann generators, to improve the efficiency of MD sampling. In this work, we developed a dual-scale flow matching method that separates training and inference into coarse-grained and all-atom stages and enhances both the accuracy and efficiency of standard flow matching samplers. We demonstrate the effectiveness of this method on a dataset of Y6 molecular clusters obtained through MD simulations, and we benchmark its efficiency and accuracy against single-scale flow matching methods.

1 Introduction

Amorphous molecular solids are disordered organic systems with several applications in the field of organic optoelectronics. Their mechanical flexibility and solution processability make them attractive alternatives to inorganic semiconductors [22, 14, 19]. Different molecules tend to pack differently in amorphous solids, which lead to different network extensivity, electronic properties, and transport rates [13]. The influence of packing structure on optoelectronic properties has led to significant interest in computationally simulating these materials with molecular dynamics (MD) simulations [13, 6] to enhance our understanding of the correlations between structural design and corresponding packing and optoelectronic properties. This can inform better design strategies for the next generation of organic electronics devices.

Having the ability to generate amorphous solids in a computationally efficient manner can enable faster sampling from the multi-molecule Boltzmann distribution, as well as property-guided generation of optimally packed configurations. There have been several works developing efficient methods to sample from Boltzmann distributions or their emulated (non re-weighted) versions [17, 11]. Scaling up to larger systems (such as amorphous clusters) is however computationally expensive if generation is to be performed at the all-atom (AA) level. In the field of protein generation, coarse-grained (CG) representations are used to collapse amino acids into single beads, which drastically reduces system size while still retaining higher-level structural features. There have been works that generate proteins in their coarse-grained representations, and other works that back-map coarse-grained proteins into their all-atom structure [12, 4, 24, 10].

In this paper, we develop a unified dual-scale flow matching method that improves accuracy of all-atom flow matching methods by separating training and inference into different stages for coarse-grained and all-atom resolutions. Moreover, we show that by performing a majority of the inference compute at the coarse-grained resolution, one can reduce the time taken for inference integration without sacrificing accuracy, paving the way for more efficient generation of larger systems. As a demonstrative example, in this work, we used a dataset of amorphously packed clusters consisting of 5 Y6 [25] molecules that we obtained from MD simulations. As the coarse-grained representation for our experiments, we pre-defined a mapping scheme that collapses the 505 atoms system into 65 beads (see Figure 1(a)). With gains in efficiency obtained from this method, we plan to scale up to larger systems in the future that are more representative of thin films in devices, as well as test other choices of coarse-graining mappings.

2 Background

Conditional Flow Matching (CFM). This is a new paradigm of training continuous normalizing flows (CNFs) efficiently in a simulation free manner [3, 2, 15]. An ODE parameterized by a learnable vector field $v_\theta(\mathbf{x}, t)$ transforms an easy-to-sample prior distribution p_0 , into the more complex target distribution p_1 . Parameters θ are learned through the regression objective

$$\mathcal{L}_{CFM} = \mathbb{E}_{t \sim U[0,1], \mathbf{z} \sim p_0(\mathbf{x}_0)p_1(\mathbf{x}_1), \mathbf{x} \sim p_t(\mathbf{x}|\mathbf{z})} \|v_\theta(\mathbf{x}, t) - u(\mathbf{x}|\mathbf{z})\|^2, \quad (1)$$

We chose the probability path definition to be $p_t(\mathbf{x}|\mathbf{z}) = \mathcal{N}(\mathbf{x}|t\mathbf{x}_1 + (1-t)\mathbf{x}_0, \sigma^2)$, which gives $u(\mathbf{x}|\mathbf{z}) = \mathbf{x}_1 - \mathbf{x}_0$ as the conditional vector field. Similar to work by Stark et al. [20], we trained v_θ to predict \mathbf{x}_1 rather than $\mathbf{x}_1 - \mathbf{x}_0$, and parameterized v_θ by SE(3) equivariant refinement tensor field network (TFN) [23] layers. We present a comparison against other parameterizations in Section A.1.

Harmonic Prior. Besides the standard gaussian prior, another common choice for prior in molecule generation tasks is the harmonic prior. The harmonic prior was introduced in Jing et al. [9] and used within the flow-matching framework by Stark et al. [20] as a more chemically plausible prior. The prior distribution $p_0(\mathbf{x}_0)$ is interpreted as a Boltzmann distribution with a quadratic potential energy function, i.e., $p_0(\mathbf{x}_0) \propto e^{-E(\mathbf{x})} = e^{-\mathbf{x}^T \mathbf{H} \mathbf{x}}$, where \mathbf{H} is chosen such that $E(\mathbf{x}) = \sum_{i,j \in \mathcal{E}} \|\mathbf{x}_i - \mathbf{x}_j\|^2$. We need to have the bond connectivity information \mathcal{E} a-priori to sample from the harmonic prior.

3 Methods

3.1 Model

Dual-Scale CFM. This method decomposes the task of predicting all-atom coordinates \mathbf{x}_1 , into two sequential steps of predicting coarse-grained coordinates $\mathbf{c}_1 \in \mathbb{R}^{M \times 3}$, and predicting all-atom coordinates $\mathbf{x}_1 \in \mathbb{R}^{N \times 3}$, where M and N are dimensions of coarse-grained and all-atom coordinates respectively. Two separate vector field networks (VFNs) are learned: 1) $v_\theta(\mathbf{c}, t)$ to generate coarse-grained bead positions \mathbf{c}_1 , and 2) $v_\phi(\mathbf{x}, t|\hat{\mathbf{c}}_1)$ to generate all-atom positions \mathbf{x}_1 conditioned on predicted bead positions $\hat{\mathbf{c}}_1$. Separate priors $p_0(\mathbf{c}_0)$ and $p_0(\mathbf{x}_0)$ are used for the two flows since they act on different input dimensionalities.

Training and Inference. Both flows are trained using Eq. 1 with the relevant distributions and vector fields being substituted in. To be precise, the two objectives are:

$$\mathcal{L}_{CG} = \mathbb{E}_{t \sim U[0,1], \mathbf{z} \sim p_0(\mathbf{c}_0)p_1(\mathbf{c}_1), \mathbf{c} \sim p_t(\mathbf{c}|\mathbf{z})} \|v_\theta(\mathbf{c}, t) - u(\mathbf{c}|\mathbf{z})\|^2, \quad (2)$$

$$\mathcal{L}_{AA} = \mathbb{E}_{t \sim U[0,1], \mathbf{z} \sim p_0(\mathbf{x}_0)p_1(\mathbf{x}_1), \mathbf{x} \sim p_t(\mathbf{x}|\mathbf{z})} \|v_\phi(\mathbf{x}, t|\hat{\mathbf{c}}_1) - u(\mathbf{x}|\mathbf{z})\|^2, \quad (3)$$

v_ϕ was trained using the ground-truth coarse-grained coordinates \mathbf{c}_1 as the condition instead of predicted coordinates $\hat{\mathbf{c}}_1$, which decoupled the two flows. Inference was performed using an Euler ODE solver to integrate from $t = 0$ to 1 for each of the two flows. For the coarse-grained flow, the bead positions were obtained with $\hat{\mathbf{c}}_1 = v_\theta(\mathbf{c}, t)$, and an integration timestep was performed with $\mathbf{c}_{t+\Delta t} = \mathbf{c}_t + \Delta t(\hat{\mathbf{c}}_1 - \mathbf{c}_0)$. Similarly, for the all-atom flow, $\hat{\mathbf{x}}_1 = v_\phi(\mathbf{x}, t|\hat{\mathbf{c}}_1)$, and $\mathbf{x}_{t+\Delta t} = \mathbf{x}_t + \Delta t(\hat{\mathbf{x}}_1 - \mathbf{x}_0)$ were used. A total of 40 integration steps were taken to go from coarse-grained prior sample \mathbf{c}_0 to all-atom target \mathbf{x}_1 .




Figure 1: **Coarse-graining mapping scheme, and influence of CG:AA ratio on metrics.** (a) Our coarse-graining mapping scheme for an individual Y6 molecule consisting of 13 beads identified with different colors. (b) JSD bond lengths, angles, and inference time on test data generation as a function of CG:AA ratio. As we increased the ratio, we observed a noticeable decreasing trend in inference time with negligible change in JSD values.

3.2 Graph Representation

Besides cartesian coordinates (\mathbf{c}/\mathbf{x}), the vector field predictors $v_{\theta/\phi}$ also take atom features $\mathbf{f}^{\text{CG}/\text{AA}}$ and bond features $\mathbf{b}^{\text{CG}/\text{AA}}$ as input. More details on features and their data-types are provided in Table 2.

We pre-defined a mapping B between node indices of all-atom atoms and coarse-grained beads, based on the basic chemical intuition that atoms present within rings have restricted degrees of freedom. The scheme is shown in Figure 1(a). We needed to define a means of collapsing the atom features \mathbf{f}^{AA} and bond features \mathbf{b}^{AA} into corresponding features \mathbf{f}^{CG} and \mathbf{b}^{CG} of coarse-grained beads so that this information was not completely lost to the flow v_{θ} .

Atom Features. We aggregated atom features with $\mathbf{f}_i^{\text{CG}} = \bigoplus_{j \in B(i)} \mathbf{f}_j^{\text{AA}}$, where \bigoplus denotes an arbitrary aggregation operator. In our case we used either the elementwise Mean or OR operations depending on the data-type of the atomic feature. We did not include `String` atom features as input to v_{θ} . We assigned the cartesian coordinates of beads to be the averaged coordinates of all atoms present within the bead, and included them as equivariant node features to v_{θ} .

Bond Features. We assigned bond connectivity to beads based on connectivity of atoms within the beads. If any two atoms belonging to different beads had a bond connecting them in the all-atom structure, we created a bond between the two beads with the corresponding bond features.

3.3 Dataset

We followed the approach developed in Kupgan et al. [13] to simulate an amorphous morphology for the molecules. We first performed geometry optimization of the Y6 molecular structure using the method described in [21]. Then we used PACKMOL [16] to pack 5 structures in a cubic box at a low density ($\sim 0.1 \text{ g/cm}^3$). The OPLS-AA force field [18] was used via the LigParGen server [5]. The system was equilibrated for 30 ns at 650 K, cooled to room temperature (300 K) at 10 K/ns, and subjected to a 30 ns production run at 300 K. All MD simulations were performed under the NPT ensemble at 1 atm with a timestep of 2 fs, using GROMACS 2023 [1]. We saved 301 frames from the production run and used a random 80:10:10 split into train, validation and test datasets. We unwrapped the molecular structures to remove periodic boundary conditions (PBC) and also removed hydrogen atoms. Each resulting frame contained 505 atoms in total in the all-atom representation.

4 Results and Discussion

We evaluated generated samples on 3 metrics (shown in Table 1). To evaluate distributional matching capabilities, we measured Jensen Shannon Divergence (JSD) between distributions of generated and

MD distributions for bond lengths and angles. We evaluated efficiency through the average time taken per integration step during inference on one NVIDIA A100 GPU. We noticed that the dual-scale flow matching method improved upon a single-scale flow matching method (with Gaussian/Harmonic priors) by 15-25% on bond length and angle JSDs, while also decreasing the inference time by $\sim 85\%$.

We also tested the influence of the ratio of integration steps taken in coarse-grained and all-atom resolutions, on the 3 metrics to see how much inference time can be saved without sacrificing accuracy (shown in Figure 1(b)). We denote the ratio of number of CG and AA integration steps by CG:AA. It is interesting to see that even for large CG:AA ratios, we significantly decreased inference time with negligible sacrifice to JSD performance.

Table 1: **Distribution and inference time metrics.** Comparison of single-scale CFM with Gaussian and Harmonic priors, and dual-scale CFM with Gaussian prior. We used 2 distributional metrics (on bond lengths and bond angles) that quantify the Jensen Shannon divergence (JSD) between generated and MD distributions on the test dataset. Inference time is the average time taken per integration step on one NVIDIA A100 GPU.

Model	Bond Lengths (JSD)	Bond Angles (JSD)	Inference Time (s)
SINGLE SCALE	0.6563	0.6316	0.2949
SINGLE SCALE HARMONIC	0.6298	0.6066	0.3039
DUAL SCALE 30:10 (OURS)	0.5472	0.4610	0.0496

The preliminary results are encouraging, opening up several avenues of further exploration for the future. We plan to scale to larger systems (~ 200 Y6 molecules) which is more representative of active layers in devices, and also test our approach on different chemistries and system sizes. Secondly, while we pre-defined a coarse-graining mapping scheme in this paper, we plan to test other schemes in the future to identify correlations between the chosen schemes and performance. Finally, we plan on expanding to other metrics that are also typically used in Boltzmann generators and conformer generation such as effective sample size (ESS), average minimum RMSD (AMR) and coverage [11, 7, 8].

5 Conclusions

In this work, we developed a dual-scale flow matching method that improves upon single-scale flow matching-based samplers in accuracy as well as inference efficiency, as demonstrated on a dataset of amorphous Y6 clusters. Moreover, it is interesting to see that we can push the ratio of CG:AA during inference integration to large values without sacrificing the prediction accuracy, while boosting inference speed. While we tested on clusters of size 5 in this paper, we plan to expand to larger clusters in the future whose packing properties are more representative of active layers in optoelectronic devices. We also plan on testing the influence of coarse-graining mapping scheme on performance since we tested a single pre-defined scheme in this work.

References

- [1] M. J. Abraham, T. Murtola, R. Schulz, S. Páll, J. C. Smith, B. Hess, and E. Lindahl. Gromacs: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX*, 1:19–25, 2015.
- [2] M. S. Albergo and E. Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022.
- [3] M. S. Albergo, N. M. Boffi, and E. Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.
- [4] N. E. Charron, F. Musil, A. Guljas, Y. Chen, K. Bonneau, A. S. Pasos-Trejo, J. Venturin, D. Gusew, I. Zaporozhets, A. Krämer, et al. Navigating protein landscapes with a machine-learned transferable coarse-grained model. *arXiv preprint arXiv:2310.18278*, 2023.

- [5] L. S. Dodda, I. Cabeza de Vaca, J. Tirado-Rives, and W. L. Jorgensen. Ligpargen web server: an automatic opls-aa parameter generator for organic ligands. *Nucleic acids research*, 45(W1): W331–W336, 2017.
- [6] Y. Fu, T. H. Lee, Y.-C. Chin, R. A. Pacalaj, C. Labanti, S. Y. Park, Y. Dong, H. W. Cho, J. Y. Kim, D. Minami, et al. Molecular orientation-dependent energetic shifts in solution-processed non-fullerene acceptors and their impact on organic photovoltaic performance. *Nature Communications*, 14(1):1870, 2023.
- [7] O. Ganea, L. Pattanaik, C. Coley, R. Barzilay, K. Jensen, W. Green, and T. Jaakkola. Geomol: Torsional geometric generation of molecular 3d conformer ensembles. *Advances in Neural Information Processing Systems*, 34:13757–13769, 2021.
- [8] B. Jing, H. Stark, T. Jaakkola, and B. Berger. Generative modeling of molecular dynamics trajectories. In *ICML'24 Workshop ML for Life and Material Science: From Theory to Industry Applications*.
- [9] B. Jing, E. Erives, P. Pao-Huang, G. Corso, B. Berger, and T. Jaakkola. Eigenfold: Generative protein structure prediction with diffusion models. *arXiv preprint arXiv:2304.02198*, 2023.
- [10] B. Jing, B. Berger, and T. Jaakkola. Alphafold meets flow matching for generating protein ensembles. *arXiv preprint arXiv:2402.04845*, 2024.
- [11] L. Klein and F. Noé. Transferable boltzmann generators. *arXiv preprint arXiv:2406.14426*, 2024.
- [12] J. Kohler, Y. Chen, A. Kramer, C. Clementi, and F. Noé. Flow-matching: Efficient coarse-graining of molecular dynamics without forces. *Journal of Chemical Theory and Computation*, 19(3):942–952, 2023.
- [13] G. Kupgan, X. Chen, and J.-L. Bredas. Molecular packing of non-fullerene acceptors for organic solar cells: distinctive local morphology in y6 vs. itic derivatives. *Materials Today Advances*, 11:100154, 2021.
- [14] Y. Li, W. Huang, D. Zhao, L. Wang, Z. Jiao, Q. Huang, P. Wang, M. Sun, and G. Yuan. Recent Progress in Organic Solar Cells: A Review on Materials from Acceptor to Donor. *Molecules*, 27(6):1800, 2022.
- [15] Y. Lipman, R. T. Chen, H. Ben-Hamu, M. Nickel, and M. Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- [16] L. Martínez, R. Andrade, E. G. Birgin, and J. M. Martínez. Packmol: A package for building initial configurations for molecular dynamics simulations. *Journal of Computational Chemistry*, 30(13):2157–2164, 2009.
- [17] F. Noé, S. Olsson, J. Köhler, and H. Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, 2019.
- [18] M. J. Robertson, J. Tirado-Rives, and W. L. Jorgensen. Improved peptide and protein torsional energetics with the opls-aa force field. *Journal of chemical theory and computation*, 11(7): 3499–3509, 2015.
- [19] T. Shan, X. Hou, X. Yin, and X. Guo. Organic photodiodes: device engineering and applications. *Frontiers of Optoelectronics*, 15(1):49, 2022.
- [20] H. Stark, B. Jing, R. Barzilay, and T. Jaakkola. Harmonic self-conditioned flow matching for joint multi-ligand docking and binding site design. In *Forty-first International Conference on Machine Learning*.
- [21] A. Subramanian, K. P. Greenman, A. Gervaix, T. Yang, and R. Gómez-Bombarelli. Automated patent extraction powers generative modeling in focused chemical spaces. *Digital Discovery*, 2(4):1006–1015, 2023.
- [22] L. Sun, K. Fukuda, and T. Someya. Recent progress in solution-processed flexible organic photovoltaics. *npj Flexible Electronics*, 6(1):89, 2022.

- [23] N. Thomas, T. Smidt, S. Kearnes, L. Yang, L. Li, K. Kohlhoff, and P. Riley. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*, 2018.
- [24] S. Yang and R. Gómez-Bombarelli. Chemically transferable generative backmapping of coarse-grained proteins. *arXiv preprint arXiv:2303.01569*, 2023.
- [25] J. Yuan, Y. Zhang, L. Zhou, G. Zhang, H.-L. Yip, T.-K. Lau, X. Lu, C. Zhu, H. Peng, P. A. Johnson, et al. Single-junction organic solar cell with over 15% efficiency using fused-ring acceptor with electron-deficient core. *Joule*, 3(4):1140–1151, 2019.

A Appendix / supplemental material

A.1 Choice of VFN architecture

Our choice of VFN architecture was based on a comparison of various architectures (shown in Table 2). We observed that using TFN as the VFN achieved best JSD performance across bond lengths and angles, as well as across gaussian and harmonic prior choices.

Table 2: **Bond length and angle distributions comparison.** JSD comparison between TFN, EGNN, and Attentive FP with Gaussian and Harmonic priors across the two collective variables (CV).

CV	Prior	TFN	EGNN	Attentive FP
BOND LENGTH	GAUSSIAN	0.656	0.791	0.959
	HARMONIC	0.629	0.781	0.894
BOND ANGLE	GAUSSIAN	0.631	0.732	0.908
	HARMONIC	0.606	0.770	0.868

A.2 Graph Featurization

Table 3: **Atom and Bond Features.** Features extracted for atoms and bonds, along with their respective data types used for featurization. Same as was used by Stark et al. [20].

Feature Name	Datatype
ATOMIC_NUM	Integer
CHIRALITY	String
DEGREE	Integer
NUMRING	Integer
IMPLICIT_VALENCE	Integer
FORMAL_CHARGE	Integer
NUMH	Integer
HYBRIDIZATION	String
IS_AROMATIC	Boolean
IS_IN_RING5	Boolean
IS_IN_RING6	Boolean
BOND_TYPE	String
BOND_STEREO	String
IS_CONJUGATED	Boolean