CLIFFORD GROUP EQUIVARIANT DIFFUSION MODELS FOR 3D MOLECULAR GENERATION

Cong Liu*, **Sharvaree Vadgama***,
David Ruhe, Erik Bekkers, Patrick Forré
AMLab, University of Amsterdam, The Netherlands
**Equal contribution.*

ABSTRACT

This paper explores leveraging the Clifford algebra's expressive power for $\mathbb{E}(n)$-equivariant diffusion models. We utilize the geometric products between Clifford multivectors and the rich geometric information encoded in Clifford subspaces in *Clifford Diffusion Models* (CDMs). We extend the diffusion process beyond just Clifford one-vectors to incorporate all higher-grade multivector subspaces. The data is embedded in grade-k subspaces, allowing us to apply latent diffusion across complete multivectors. This enables CDMs to capture the joint distribution across different subspaces of the algebra, incorporating richer geometric information through higher-order features. We provide empirical results for unconditional molecular generation on the QM9 dataset, showing that CDMs provide a promising avenue for generative modeling.

1 INTRODUCTION

Deep generative models have revolutionized molecular science, enabling significant progress in molecular design and drug discovery (Abramson et al. (2024); Bose et al. (2024); Yim et al. (2023); Watson et al. (2023)). Learning the structure of molecules, these AI models help accelerate drug discovery by replacing costly lab experiments and help streamline the process of designing new drugs or proteins.

Since molecules exist in 3-dimensional space, thus the group of 3D symmetries $\mathrm{E}(3)$, including translations, rotations, and reflections, determine how they transform. To ensure physical validity, generative models for molecular design must respect these symmetries, so that a molecule and any of its symmetric transformations are equally likely according to the learned distribution p_θ, i.e., p_θ is invariant to $\mathrm{E}(3)$ (Xu et al. (2022)). Most recent work addresses this requirement using generative models with denoising neural networks composed of layers that are equivariant to the orthogonal group $\mathrm{O}(3)$, generated by rotations and reflections, e.g., equivariant graph neural network (EGNN) layers (Satorras et al. (2021)), which enable the $\mathrm{O}(3)$ invariance of the learned distributions (Garcia Satorras et al. (2021); Hoogeboom et al. (2022); Song et al. (2024)). However, these methods typically focus only on scalar and scaled euclidean vector representations, limiting their ability to capture richer geometric structures inherent in molecular systems.

Clifford Group Equivariant Neural Networks (CGENNs) (Ruhe et al. (2023a)) establish neural networks operating on the Clifford algebra's rich geometric representations (called *multivectors*) while maintaining $\mathrm{O}(n)$- or $\mathrm{SO}(n)$-equivariance (Liu et al. (2024a); Brandstetter et al. (2022); Brehmer et al. (2023); Zhdanov et al. (2024); Spinner et al. (2024); Ruhe et al. (2023b)). To improve their efficiency, lightweight variants of CGENNs have been introduced (Liu et al. (2024b)), which make them suitable for large-scale applications.

In this work, we extend diffusion processes into grade-k subspaces, moving beyond conventional vector representations. By harnessing multivector structures, our model captures diverse higher-order geometric features of molecular systems by learning on Clifford subspaces. This design enables parallel latent diffusion across distinct Clifford subspaces, facilitating richer symmetry encoding. We demonstrate the efficacy of our approach in unconditional molecular generation tasks on the QM9 dataset, achieving favorable performance compared to existing methods.

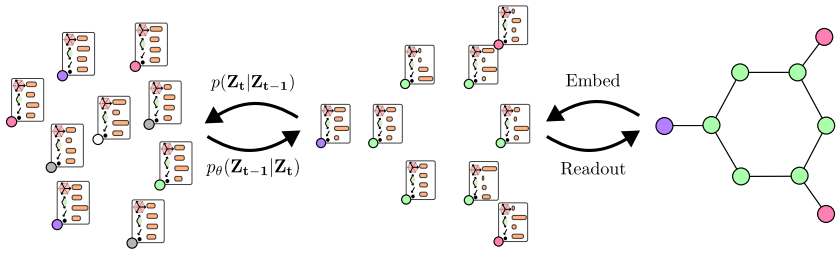

Figure 1: During sampling of Clifford all-grade diffusion models, subspace features are initialized from a Gaussian distribution, and CGENNs are used as the denoising model ϕ_θ for each subspaces. At time step zero, the molecular structure is readout by projecting a one-vector from Clifford space.

Related Work Equivariant Diffusion Models (EDMs) (Hoogeboom et al., 2022) leverage $\mathrm{E}(3)$ symmetries using EGNN as backbone (Satorras et al., 2021) and denoising diffusion models (Ho et al., 2020) to unconditionally generate molecules, atom positions as well as atom types by treating both as continuous variables. Prior to these, works like (Gebauer et al., 2019; Simonovsky & Komodakis, 2018; Simm et al., 2021) have shown that incorporating symmetries helps generalization for molecular generation. Bekkers et al. (2024) and Vadgama et al. (2025) use $\mathrm{SE}(3)$ equivariant graph backbone and use preconditioned diffusion models (Karras et al., 2022) to improve sampling efficiency. GeoLDMs (Xu et al., 2023) perform latent diffusion with E(3) symmetries in the latent space. Recent works have focused on a unified representation of discrete and continuous features. Vignac et al. (2023) generates both molecular graphs and 3D atomic arrangements, while in joint diffusion 2D-3D model, JODO (Huang et al., 2023) and flow matching model for generation, FLowMol (Dunn & Koes, 2024), produces complete molecules with atom types, charges, bonds, and 3D coordinates. For the scope of this work, we focus only on using atom positions and atom types as continuous variables and only compare with methods that do so for unconditional generation of molecules.

2 Background

2.1 $\mathrm{E}(3)$ Equivariant Diffusion Models

$\mathrm{E}(3)$ Equivariant Diffusion Models (EDMs) are generative models that learn to transform a simple Gaussian distribution into learned molecular coordinate distribution and atomic type distribution using Denoising Diffusion Probabilistic Models (Ho et al. (2020)).Given a data sample x_1, during the forward process, the distribution of Gaussian-corrupted sample x_t at timestep t is defined with noise schedule β_t:

$$q(x_t|x_1) = \mathcal{N}(x_t|\sqrt{\bar{\alpha}_t}x_1, (1-\bar{\alpha}_t)I),$$

where $\bar{\alpha}_t = \prod_{s=1}^{t}(1-\beta_s)$. The model predicts the noise ϵ_t added to the data during this process and is trained to minimize the denoising loss:

$$\mathcal{L}_{\text{denoise}} = \mathbb{E}_{t\sim U, x_1\sim q(x), \epsilon\sim\mathcal{N}(0,I)}\left[\|\phi_\theta(x_t(x_1), t) - \epsilon\|^2\right],$$

where U the uniform distribution across predefined time-step range and $q(x_1)$ represents the data distribution, ϕ_θ the neural network that predicts noise. To make the learned distribution p_θ invariant to $\mathrm{E}(3)$, the noise ϵ and data sample x_1 are zero-centered to their mass point. In addition, the network ϕ_θ needs to be equivariant to $\mathrm{O}(3)$ transformations (Xu et al. (2022)). Once trained, the model generates samples by reversing the forward process, starting from Gaussian noise and iteratively applying the learned denoising steps.

2.2 Clifford Group Equivariant Neural Networks

CGENNs are a class of neural networks that operate on multivectors, rather than traditional scalars or vectors. Built on the mathematical foundation of Clifford (*geometric*) algebra, CGENNs leverage its rich geometric properties to model complex relationships in 3D space, making them well suited

for tasks like molecular modeling. The Clifford algebra, denoted as $\mathrm{Cl}(V, \mathfrak{q})$, is defined over an n-dimensional vector space V paired with a quadratic form $\mathfrak{q} : V \to \mathbb{R}$. In this work, we stick to $V := \mathbb{R}^3$ and use for $\mathfrak{q}$ the standard positive definite Euclidean norm, related to the ordinary Euclidean scalar product. For simplicitly, we write $\mathrm{Cl}(\mathbb{R}^3)$. In this case, the algebra provides a unified framework to represent scalars, vectors, and higher-dimensional geometric objects, such as bivectors (oriented planes) and trivectors (oriented volumes). The fundamental operation in Clifford algebra is the *geometric product*, which allows vectors to be combined to form higher-grade objects. A *multivector* is a general element of Clifford algebra, expressed as a direct sum of components from different *grades*, i.e., for $\mathbf{v} \in \mathrm{Cl}(\mathbb{R}^3)$: $\mathbf{v} = a_0\mathbf{v}^{(0)} + a_1\mathbf{v}^{(1)} + a_2\mathbf{v}^{(2)} + a_3\mathbf{v}^{(3)}$, where $\mathbf{v}^{(m)}$ corresponds to the grade-m component, and $a_i \in \mathbb{R}$ are scalar coefficients.

3 METHODS

In this section, we introduce two approaches for extending diffusion models to the Clifford k-grade spaces. First, we demonstrate how diffusion models can be applied to 1-grade subspaces. Second, inspired by Xu et al. (2023), we describe how to perform diffusion across all k-grade subspaces of multivectors in a latent embedding, allowing the model to learn and sample the joint distribution of all grades of multivectors in Clifford space.

3.1 CLIFFORD GROUP EQUIVARIANT ONE-VECTOR DIFFUSION MODELS

Consider a molecular graph $\mathcal{G} = \{\boldsymbol{X}, \boldsymbol{H}\}$, where $\boldsymbol{X}$ represents the atomic coordinates in $\mathbb{R}^3$. In Clifford 1-vector diffusion models, we first embed $\boldsymbol{X}$ into the Clifford algebra $\mathrm{Cl}(\mathbb{R}^3)$ as a 1-vector, denoted as $\boldsymbol{X}^{\mathrm{Cl}}$. Since the grade-one subspace $\mathrm{Cl}^{(1)}(\mathbb{R}^3)$ of $\mathrm{Cl}(\mathbb{R}^3)$ is isomorphic to $\mathbb{R}^3$, this embedding corresponds to a direct identification: each atomic coordinate $\boldsymbol{x}_i \in \mathbb{R}^3$ is mapped to a 1-vector in $\mathrm{Cl}^{(1)}(\mathbb{R}^3)$, while all higher-grade components (e.g., scalars, bivectors, and trivectors) are set to zero. This ensures that the representation remains consistent with the standard Euclidean vector space while leveraging the algebraic structure of the Clifford space[1].

During the forward process, noise is progressively added to the grade-one component of the Clifford space. The distribution of the intermediate sample $\boldsymbol{X}_{\boldsymbol{t}}^{\mathrm{Cl}}$ at time step t follows:

$$p(\boldsymbol{X}_{\boldsymbol{t}}^{\mathrm{Cl}} \mid \boldsymbol{X}^{\mathrm{Cl}}) = \mathcal{N}(\boldsymbol{X}_{\boldsymbol{t}}^{\mathrm{Cl}}; \sqrt{\bar{\alpha}_t}\boldsymbol{X}^{\mathrm{Cl}}, (1 - \bar{\alpha}_t)\boldsymbol{I}).$$

For the backward process, we adopt the Clifford-EGNN introduced in (Liu et al., 2024b) as the denoising model ϕ_θ. This model estimates the noise added to the grade-one components at each time step. The overall diffusion framework follows a structure similar to EDMs, where EDMs employ EGNNs (Satorras et al. (2021)) as denoising models. However, in our approach, $\boldsymbol{X}$ is treated as a multivector, allowing Clifford-EGNN to better capture geometric structures during denoising. Other details of the diffusion process on atomic type features closely follow EDMs.

3.2 CLIFFORD GROUP EQUIVARIANT ALL-GRADE DIFFUSION MODELS

In Clifford One-Vector Diffusion Models, we treat the Cartesian coordinates $\boldsymbol{X}$ as one-vectors in Clifford space, which is equivalent to viewing $\boldsymbol{X}$ as vectors in Euclidean space. Inspired by Xu et al. (2023), we extend this approach to Clifford All-grade Diffusion Models, where the diffusion process occurs across all Clifford subspaces. To facilitate this, we introduce a latent encoder $\mathcal{E}$, which lifts $\boldsymbol{X}$ into Clifford space and enriches each sample with geometrically informed features. This enables the diffusion process to leverage higher-order geometric information encoded in different grades of the Clifford algebra. By incorporating these latent geometric features, we aim to assess whether such enrichment provides any significant advantage over traditional diffusion models that operate solely in Euclidean space.

To obtain higher-order features, we employ a Clifford-EGNN as an encoder $\mathcal{E}$. This encoder takes the initial $\boldsymbol{X}^{\mathrm{Cl}}$ as input and outputs latent multivectors $\boldsymbol{Z}$ with all grades filled with meaningful geometric features. To preserve the geometric layout of the data sample, we replace the grade-1 part of $\boldsymbol{Z}$ with $\boldsymbol{X}$. Consequently, the encoder is responsible for producing only the other grades

[1] We refer Clifford space to any product of Clifford algebras: $\mathrm{Cl}(\mathbb{R}^3)^k$

components of the multivector. During the forward process, we add Gaussian noise to each subspace of the multivector. Therefore, the distribution of $\boldsymbol{Z}_{\boldsymbol{t}}^{(\boldsymbol{m})}$ remains Gaussian:

$$p\big(\boldsymbol{Z}_{\boldsymbol{t}}^{(\boldsymbol{m})} \mid \boldsymbol{Z}^{(\boldsymbol{m})}\big) = \mathcal{N}\big(\boldsymbol{Z}_{\boldsymbol{t}}^{(\boldsymbol{m})}; \sqrt{\bar{\alpha}_t}\boldsymbol{Z}^{(\boldsymbol{m})}, (1-\bar{\alpha}_t)\boldsymbol{I}\big),$$

with $m \in \{0, 1, 2, 3\}$. This means that we add noise independently to each of the Clifford subspaces. The joint distribution $p\big(\boldsymbol{Z}_{\boldsymbol{t}} \mid \boldsymbol{Z}\big)$ then is given by:

$$p\big(\boldsymbol{Z}_{\boldsymbol{t}} \mid \boldsymbol{Z}\big) = \prod_{m=0}^{3} p\big(\boldsymbol{Z}_{\boldsymbol{t}}^{(\boldsymbol{m})} \mid \boldsymbol{Z}^{(\boldsymbol{m})}\big).$$

During the backward process, the denoising model ϕ_θ learns to output and approximation of the noise $\boldsymbol{\epsilon}$ in Clifford space, i.e. $\boldsymbol{\epsilon} \in \mathrm{Cl}(\mathbb{R}^3)$. Thus, both the diffusion (forward) and denoising (backward) steps are fully carried out in Clifford space, leveraging the enriched geometric features at multiple grades.

4 Experiments

In this section, we present empirical results for unconditional molecular generation on the QM9 dataset (Ramakrishnan et al., 2014). QM9 consists of 3D molecular structures, where each atom is annotated with its atomic type and charge. We evaluate our generation samples on *atomic stability* percentage of generated atoms that satisfy the correct valency, *molecular stability* percentage of molecules in which all constituent atoms are stable, *uniqueness* number of unique samples generated as well as *validity* percentage of valid samples as defined by RdKit.

Models	Atom Stability (%)	Mol Stability (%)	Validity (%)	Valid & Unique (%)
GDM-AUG	97.6	71.6	90.4	89.5
EDM (Hoogeboom et al. (2022))	98.7 ± 0.1	82.0 ± 0.4	91.9 ± 0.5	83.3 ± 0.6
GeoLDM (Xu et al. (2023))	98.9 ± 0.1	89.4 ± 0.5	93.8 ± 0.4	92.7 ± 0.5
PΘNITA (Bekkers et al. (2024))	98.9	87.8	-	-
MUDiff† (Hua et al. (2024))	98.8 ± 0.2	89.9 ± 1.1	95.3 ± 1.5	94.4 ± 0.5
END† (Cornet et al. (2024))	98.9 ± 0.2	89.1 ± 0.1	94.8 ± 1.5	87.8 ± 0.4
EquiFM† (Song et al. (2024))	98.9 ± 0.1	88.3 ± 0.3	94.7 ± 0.4	93.5 ± 0.3
Rapidash† (Vadgama et al. (2025))	$\mathbf{99.38 \pm 0.02}$	$\mathbf{92.91 \pm 0.41}$	$\mathbf{98.12 \pm .003}$	$95.35 \pm .003$
CDM (one-vector)	98.9 ± 0.0	89.6 ± 0.2	96.0 ± 0.3	95.8 ± 0.3
CDM (all-grade)	99.0 ± 0.2	89.7 ± 1.4	96.4 ± 1.0	$\mathbf{96.3 \pm 1.0}$
Data	99.00	95.20	97.8	97.8

Table 1: Results for unconditional generation task on QM9 dataset. † represents works that have a different generative model than the rest.

We compare our proposed models, CDMs[2] with Graph Diffusion Models without symmetries (GDM -AUG), E(3) equivariant diffusion models EDMs, (Hoogeboom et al., 2022), latent diffusion models, GeoLDMs (Xu et al., 2023), MuDiff (Hua et al., 2024), END (Cornet et al., 2024), EquiFM (Song et al., 2024), and Rapidash (Vadgama et al., 2025). During evaluation, 10000 samples are generated from each models to evaluate the standard metrics. From Table 1, we can see that our CDMs perform competitively across all evaluated metrics compared to baseline models, which validate our models' effectiveness. CDMs with all-grade diffusion in general generate molecules with higher quality. This result indicates strong potential of applying generative models on other grades in Clifford algebra. All CDMs results are calculated based on three runs with different seeds.

5 Conclusion

We introduce Clifford Group Equivariant Diffusion Models (CDMs), a diffusion framework operating in Clifford grade-k (sub)spaces for unconditional molecular generation. Among the two variants, Clifford one-vector diffusion effectively captures molecular geometry while maintaining

[2]To ensure a fair comparison, our CDMs have equal or less parameter counts compared to baseline models.

a direct correspondence with the Euclidean space. To compare with Xu et al. (2023), we explored Clifford all-grade diffusion, where a geometric encoder maps molecular structures into a latent Clifford representation, initializing latent Clifford features – scalars, bivectors, and trivectors before diffusion. While its effectiveness compared to Clifford one-vector diffusion remains under evaluation, we believe encoding higher-order geometric structures could enhance generative modeling in more complex scenarios.

6 ACKNOWLEDGEMENTS

C. Liu is supported by Health-Holland, Top Sector Life Sciences & Health (LSH-TKI), project number LSHM22023, which realizes a public-private partnership (PPP) between the University of Amsterdam and Janssen Vaccines and Prevention B.V. S. Vadgama is supported by the Hybrid Intelligence Center, a 10-year program funded by the Dutch Ministry of Education, Culture and Science through the Netherlands Organisation for Scientific Research (NWO). We thank SURF (www.surf.nl) for the support in using the National Supercomputer Snellius.

REFERENCES

Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J. Ballard, Joshua Bambrick, Sebastian W. Bodenstein, David A. Evans, Chia-Chun Hung, Michael O'Neill, David Reiman, Kathryn Tunyasuvunakool, Zachary Wu, Akvilė Žemgulytė, Eirini Arvaniti, Charles Beattie, Ottavia Bertolli, Alex Bridgland, Alexey Cherepanov, Miles Congreve, Alexander I. Cowen-Rivers, Andrew Cowie, Michael Figurnov, Fabian B. Fuchs, Hannah Gladman, Rishub Jain, Yousuf A. Khan, Caroline M. R. Low, Kuba Perlin, Anna Potapenko, Pascal Savy, Sukhdeep Singh, Adrian Stecula, Ashok Thillaisundaram, Catherine Tong, Sergei Yakneen, Ellen D. Zhong, Michal Zielinski, Augustin Žídek, Victor Bapst, Pushmeet Kohli, Max Jaderberg, Demis Hassabis, and John M. Jumper. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, 630(8016), 2024. doi: 10.1038/s41586-024-07487-w.

Erik J Bekkers, Sharvaree Vadgama, Rob Hesselink, Putri A Van der Linden, and David W. Romero. Fast, expressive se(n) equivariant networks through weight-sharing in position-orientation space. In *The Twelfth International Conference on Learning Representations*, 2024.

Joey Bose, Tara Akhound-Sadegh, Guillaume Huguet, Kilian FATRAS, Jarrid Rector-Brooks, Cheng-Hao Liu, Andrei Cristian Nica, Maksym Korablyov, Michael M Bronstein, and Alexander Tong. Se (3)-stochastic flow matching for protein backbone generation. In *The Twelfth International Conference on Learning Representations*, 2024.

Johannes Brandstetter, Rob Hesselink, Elise van der Pol, Erik J Bekkers, and Max Welling. Geometric and Physical Quantities Improve E(3) Equivariant Message Passing. In *International Conference on Learning Representations*, 2022.

Johann Brehmer, Pim de Haan, Sönke Behrends, and Taco Cohen. Geometric Algebra Transformer. In *Neural Information Processing Systems*, 2023.

François Cornet, Grigory Bartosh, Mikkel N Schmidt, and Christian A Naesseth. Equivariant neural diffusion for molecule generation. In *38th Conference on Neural Information Processing Systems*, 2024.

Ian Dunn and David Ryan Koes. Mixed continuous and categorical flow matching for 3d de novo molecule generation. *ArXiv*, 2024.

Victor Garcia Satorras, Emiel Hoogeboom, Fabian Fuchs, Ingmar Posner, and Max Welling. E (n) equivariant normalizing flows. *Advances in Neural Information Processing Systems*, 2021.

Niklas Gebauer, Michael Gastegger, and Kristof Schütt. Symmetry-adapted generation of 3d point sets for the targeted discovery of molecules. *Advances in neural information processing systems*, 2019.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 2020.

Emiel Hoogeboom, Vıctor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3d. In *International conference on machine learning*, 2022.

Chenqing Hua, Sitao Luan, Minkai Xu, Zhitao Ying, Jie Fu, Stefano Ermon, and Doina Precup. Mudiff: Unified diffusion for complete molecule generation. In *Learning on Graphs Conference*, pp. 33–1. PMLR, 2024.

Han Huang, Leilei Sun, Bowen Du, and Weifeng Lv. Learning joint 2d & 3d diffusion models for complete molecule generation, 2023.

Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models, 2022.

Cong Liu, David Ruhe, Floor Eijkelboom, and Patrick Forré. Clifford Group Equivariant Simplicial Message Passing Networks. In *The Twelfth International Conference on Learning Representations*, 2024a.

Cong Liu, David Ruhe, and Patrick Forré. Multivector neurons: Better and faster o(n)-equivariant clifford GNNs. In *ICML 2024 Workshop on Geometry-grounded Representation Learning and Generative Modeling*, 2024b.

Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, (1), 2014.

David Ruhe, Johannes Brandstetter, and Patrick Forré. Clifford Group Equivariant Neural Networks. In *Neural Information Processing Systems*, 2023a.

David Ruhe, Jayesh K. Gupta, Steven De Keninck, Max Welling, and Johannes Brandstetter. Geometric Clifford Algebra Networks. In *International Conference on Machine Learning*, 2023b.

Vıctor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E (n) equivariant graph neural networks. In *International conference on machine learning*, 2021.

Gregor N. C. Simm, Robert Pinsler, Gábor Csányi, and José Miguel Hernández-Lobato. Symmetry-aware actor-critic for 3d molecular design. In *International Conference on Learning Representations*, 2021.

Martin Simonovsky and Nikos Komodakis. Graphvae: Towards generation of small graphs using variational autoencoders. In *International Conference on Artificial Neural Networks*, 2018.

Yuxuan Song, Jingjing Gong, Minkai Xu, Ziyao Cao, Yanyan Lan, Stefano Ermon, Hao Zhou, and Wei-Ying Ma. Equivariant flow matching with hybrid probability transport for 3d molecule generation. *Advances in Neural Information Processing Systems*, 36, 2024.

Jonas Spinner, Victor Bresó, Pim de Haan, Tilman Plehn, Jesse Thaler, and Johann Brehmer. Lorentz-Equivariant Geometric Algebra Transformers for High-Energy Physics. 2024.

Sharvaree Vadgama, Mohammad Mohaiminul Islam, Domas Buracus, Christian Shewmake, and Erik Bekkers. On the utility of equivariance and symmetry breaking in deep learning architectures on point clouds, 2025.

Clement Vignac, Nagham Osman, Laura Toni, and Pascal Frossard. Midi: Mixed graph and 3d denoising diffusion for molecule generation. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2023.

James L. Watson, David Juergens, Nathaniel Bennett, Brian Trippe, Jaekyung Yim, Helen Eisenach, William Ahern, Andrew Borst, Robert Ragotte, Lukas Milles, Basile I. M. Wicky, Nikita Hanikel, Samuel J. Pellock, Alexis Courbet, William Sheffler, Jianyi Wang, Keenan Bett, Asim Bera, Ambarish Roy, Christopher Savile, Yifan Xu, James Dou, Rebecca R. Eguchi, Taylor Powers, Ranjani Ravichandran, Patrick Madden, Bruno Correia, William R. Schief, David Marquis, Christine M. Chow, T. J. Brunette, Frank DiMaio, Minkyung Baek, and David Baker. De novo design of protein structure and function with rfdiffusion. *Nature*, 620:687–696, 2023. doi: 10.1038/s41586-023-06426-4.

Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geometric diffusion model for molecular conformation generation. In *International Conference on Learning Representations*, 2022.

Minkai Xu, Alexander S Powers, Ron O Dror, Stefano Ermon, and Jure Leskovec. Geometric latent diffusion models for 3d molecule generation. In *International Conference on Machine Learning*, 2023.

Jason Yim, Brian L Trippe, Valentin De Bortoli, Emile Mathieu, Arnaud Doucet, Regina Barzilay, and Tommi Jaakkola. Se (3) diffusion model with application to protein backbone generation. In *International Conference on Machine Learning*, 2023.

Maksim Zhdanov, David Ruhe, Maurice Weiler, Ana Lucic, Johannes Brandstetter, and Patrick Forré. Clifford-Steerable Convolutional Neural Networks. In *International Conference on Machine Learning*, 2024.