Improving Local Connectivity in Sparse and Rigid Graphs

Anonymous submission

Abstract

Geometric graph neural networks (GNNs) depend critically on the construction of the underlying graph structure for molecular modeling, protein structure prediction, and 3D shape analysis. Recent rigidity-theory-inspired graph constructions aim to enhance the expressivity of geometric GNNs–separating geometric graphs up to isometries, while remaining sparse. However, their theoretical and empirical behavior remains poorly understood, and—crucially—they often fail to preserve local structure, introducing distortions that hinder learning of neighborhood-scale geometry. We introduce the Fundamental Forms Graph Construction (FFGC), a construction method that generates sparse, invariant, and architecture-agnostic graphs by pairing genus-0 projections with curvature-aware geometry and iterative optimization for reconciling global separation with local fidelity.

Introduction

Geometric deep learning leverages physical symmetry and structure for data efficiency and accuracy on manifolds, graphs, and groups (Bronstein et al. 2021; Battiloro et al. 2024; Liu et al. 2023), achieving remarkable performance in molecular property prediction, protein folding, and shape registration (Schütt et al. 2018; Jing et al. 2020; Igashov et al. 2021). Equivariant GNNs now anchor many state-of-the-art models (Schütt et al. 2018; Thomas et al. 2018; Satorras, Hoogeboom, and Welling 2021; Brandstetter et al. 2021; Du et al. 2023) each with varying degrees of expressivity (Joshi et al. 2023; Sverdlov and Dym 2025). However, until recently, the impact on expressivity from the underlying geometric graph received comparatively little attention (Morris et al. 2024; Sverdlov and Dym 2025; Wang et al. 2025). In practice, when graphs are selected for convenience rather than principle two recurrent issues arise. (1) over-smoothed neighborhoods (large-k kNN or large-radius ε-graphs) blur anisotropy, suppress informative spectral content, and induce near-harmonic smoothing that erases directional signals crucial for physical inference (Topping et al. 2021) and (2) task-driven rewirings (random shortcuts, small-world/ring augmentations, hierarchical/fragmentation graphs) can misalign local frames and discard anisotropy, offering ad-hoc gains without being metric- or symmetry-aware (Barbero et al. 2024; Qian et al. 2024; Sonthalia, Gilbert, and Durham 2023; Damke, Melnikov, and Hüllermeier 2020; Fey, Yuen, and Weichert 2020; Tahmasebi, Lim, and Jegelka 2023). Empirically, across benchmarks, performance tracks neighborhood quality more than generic rewiring (Dwivedi et al. 2022; Hu et al. 2020).

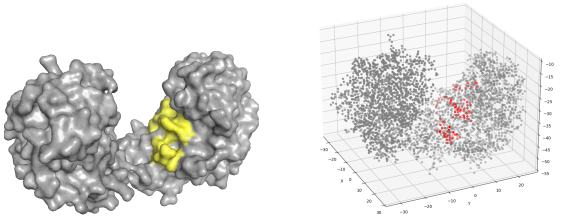

(a) Molecular Surface (b) All Atom Point Cloud

Figure 1: Molecular surface and all atom representations of the apo structure of the 4pfs protien prior to ligand binding. The cryptic binding region is highlighted as yellow on the surface (a) and red in the pointcloud (b).

The expressive power of geometric GNNs has long been tied to separability criteria, from Weisfeiler–Lehman tests (Xu et al. 2019) to more recent equivariant extension (Joshi et al. 2023; Hordan, Amir, and Dym 2024). While fragmented graphs have been shown to be expressive (Wollschläger et al. 2024), a principled path to separability emerges from rigidity theory (Sverdlov and Dym 2025)–when edge sets are sufficiently constraining, a point cloud's embedding is determined (up to global E(3) motions) by pairwise relations. Recent work operationalizes this idea, showing that carefully constructed sparse, connected graphs can enhance the distinguishability of non-isometric structures (Wang et al. 2025). However, existing approaches still face critical limitations, where (Sverdlov and Dym 2025) relies on power graphs, which need an initial connected graph, and (Wang et al. 2025)-style constructions are globally sparse, but then often rely on potentially dense radial cutoffs to reintroduce local connections.

What remains missing is a principled graph construction that simultaneously balances local and global structure while maintaining sparsity and rigidity. We introduce the Fundamental Forms Graph (FFG), which encodes the first and second fundamental forms—metric and curvature—via edge lengths and face–dihedral angles. FFGs yield sparse, con-

nected graphs with $|\mathcal{E}| = 3|\mathcal{V}| - 6$, for $|\mathcal{V}| \geq 3$, ensuring all supplied points are connected. Under mild genericity, they provide provably complete E(3) invariants, leveraging a discrete fundamental theorem of surfaces to uniquely determine the underlying point cloud up to rigid motion. FFGs are flexible and preserve local structure allowing representations that better capture neighborhood-scale geometry without resorting to dense neighborhoods or ad-hoc rewiring.

Figure 1 illustrates the importance of careful graph construction in real world applications of protein pocket discovery using CryptoBench (Škrhák et al. 2024). Cryptic pockets arise from subtle, state-dependent deformations (hinge openings, side-chain flips, transient saddles) that may be invisible on the apo surface yet are presaged by local metric/curvature signals and coordinated long-range couplings. This local–global pairing lets message passing detect nascent pocket formation: local curvature shifts that hint at a cleft are reinforced (or vetoed) by global pathways that open cooperatively across a hinge.

Background

This section reviews geometric graphs, defines expressivity, and recalls an existing approach.

Geometric Graphs Despite a variety of forms for defining a geometric graph (Joshi et al. 2023; Sverdlov and Dym 2025; Wang et al. 2025), we present a generalized version accounting for the inclusion of additional features. A geometric graph $\mathcal{G} := (\mathcal{V}, \mathcal{E}, \boldsymbol{X}, \boldsymbol{F}, \boldsymbol{E})$ contains a set of m nodes $\mathcal{V} = \{1, \ldots, m\}$, and p edges $\mathcal{E}$, along with three associated feature sets for geometric features $\boldsymbol{X}$, invariant node features $\boldsymbol{F}$ and invariant edge features $\boldsymbol{E}$. It will be critical in the construction of a geometric graph to assert that any node or edge features constructed respect their invariant and equivariant principles.

The symmetry group of geometric graphs is the product of the permutation group S_n and rigid motions $E(3)$. The action of a permutation matrix $\boldsymbol{P}$ on a geometric graph is given by $\boldsymbol{P}\mathcal{G} := (\boldsymbol{P}\boldsymbol{X}, \boldsymbol{P}\boldsymbol{F}, \boldsymbol{P}\boldsymbol{E}(\boldsymbol{X})\boldsymbol{P}^\top)$. A rigid motion in $E(3)$ is a rotation or orthogonal matrix $\boldsymbol{Q} \in O(3)$ and a translation $t \in \mathbb{R}^3$ which acts on the geometric graph as $\boldsymbol{Q}\mathcal{G} + \boldsymbol{t} := (\boldsymbol{Q}\boldsymbol{X} + \boldsymbol{t}\boldsymbol{1}^\top, \boldsymbol{F}, \boldsymbol{E})$.

Geometric GNNs Let $f_i \subseteq \boldsymbol{F}$ and $e_{ij} \subseteq \boldsymbol{X}$ be the geometric node and edge features selected by a particular geometric GNN framework. Generically a geometric GNN then propagates features as

$$f_i^{(t+1)} = \mathrm{UPD}\Big(f_i^{(t)}, \mathrm{AGG}\big(\{f_i^{(t)}, f_j^{(t)}, e_{ij} \;\mid\; j \in \mathcal{N}_i\}\big)\Big),$$

where e_{ij} are edge attributes (e.g. $x_i - x_j$ or pre-determined invariant attributes such as relative distance $\|x_i - x_j\|$), $\mathcal{N}_i$ denotes the one-hop neighbors of node i, i.e., the set of nodes in $\mathcal{V}$ that are adjacent to i in $\mathcal{E}$, and UPD and AGG are learnable update and aggregate functions. After T message-passing steps, a multiset Readout function produces graph-level features via the Readout

$$s^{\mathrm{global}} = \mathrm{Readout}\Big(\{f_i^{(T)} \;\mid\; i \in \mathcal{V}\}\Big).$$

Expressivity of Geometric GNNs Invariant expressivity captures what a geometric GNN can distinguish and represent. Given two point clouds that differ up to $E(3)$, typically by leveraging injective message passing, informative geometric features (distances, angles, dihedrals), all to avoid collapsing distinct structures under the limitation of WL-1.

A maximally expressive GNN F, where the update, aggregation, and readout functions are all injective, can, after depth T, distinguish two geometric graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ as distinct:

$$s_1^{\mathrm{global}} =: F(\mathcal{G}_1) \;\neq\; F(\mathcal{G}_2) := s_2^{\mathrm{global}}.$$

As shown by (Xu et al. 2019), expressivity is limited to that of the WL-1 algorithm (Weisfeiler and Leman 1968).

Fundamental Forms Graph Family

Rather than rewiring graph connectivity, we re-establish the graph connectivity. Moreover, unlike other base principles graph construction methods, we utilize graphs exclusively derived from a genus-0 mesh.

Definition 0.1. *(Closed Genus-0 Triangle Mesh)* $\mathcal{M} := (\mathcal{V}, \mathcal{E}, \mathcal{T}, \boldsymbol{X})$ *is given by a set of vertices* $\mathcal{V}$*, a set of associated vertex locations* $\boldsymbol{X} \subset \mathbb{R}^{|V|\times 3}$*,a set of edges* $\mathcal{E}$*, and a set of triplets corresponding to faces* $\mathcal{T}$*. In order for* $\mathcal{M}$ *to be genus-0,* $(\mathcal{V}, \mathcal{E})$ *must be a maximally planar graph.*

Fundamental Forms on Meshes In the smooth setting the *Fundamental Theorem of Surfaces* states that when the two fundamental forms of a 2-manifold (metric and curvature) are known, then they uniquely determine the surface up to $E(3)$. The discrete version of this theorem comes from (Wang, Liu, and Tong 2012), which states that a closed surface mesh with fixed length edges and dihedrals, can under certain conditions determine a unique immersion of the mesh up to $E(3)$.

Let $\mathcal{M} = (\mathcal{V}, \mathcal{E}, \mathcal{T}, X)$ be a closed genus-0 mesh. An *orientation* of $\mathcal{M}$ is a choice of ordered faces $\tilde{\mathcal{T}} \subset \{(i,j,k) : \{i,j,k\} \in \mathcal{T}\}$ such that along every interior edge $e = \{i,j\} \in \mathcal{E}$ with incident (unordered) faces $\{i,j,k\}$ and $\{i,j,\ell\}$, the two ordered faces appear with opposite edge directions, i.e., $(i,j,k) \in \tilde{\mathcal{T}}$ implies $(j,i,\ell) \in \tilde{\mathcal{T}}$. On a connected, 2-manifold mesh such a global orientation exists and is unique up to a global flip (Meyer et al. 2003).

For an oriented triangle $\tau = (i,j,k) \in \tilde{\mathcal{T}}$, the unit face normal is given by

$$\hat{\mathbf{n}}_\tau \;=\; \frac{(x_j - x_i) \times (x_k - x_i)}{\|(x_j - x_i) \times (x_k - x_i)\|}.$$

For an interior edge $e = \{p,q\} \in \mathcal{E}$ shared by faces $\tau, \tau' \in \mathcal{T}$, define the (unit) edge tangent by choosing the orientation (p,q) induced by τ and setting

$$\hat{\mathbf{t}}_{e,\tau} \;=\; \frac{x_q - x_p}{\|x_q - x_p\|}.$$

The signed dihedral angle at e is then

$$\phi_e \;=\; \mathrm{atan2}\big((\hat{\mathbf{n}}_\tau \times \hat{\mathbf{n}}_{\tau'})\cdot\hat{\mathbf{t}}_e,\; \hat{\mathbf{n}}_\tau \cdot \hat{\mathbf{n}}_{\tau'}\big).$$

This convention makes the set $\{\phi_e\}_{e\in\mathcal{E}}$ well defined, and always exist because we assume a closed mesh. We explain how to handle degenerate faces in **??**.

Fundamental Forms Graph

Definition 0.2. *(Fundamental Forms Graph) We define the* fundamental forms graph *associated with a genus-0 mesh $\mathcal{M}$ as*

$$FFG(\mathcal{M}) = (\mathcal{V}, \mathcal{E}, \{ l_e, \phi_e \}_{e \in \mathcal{E}}),$$

where l_e and ϕ_e denote the edge lengths and oriented dihedral angles, respectively.

Definition 0.3. *(Fundamental Forms Graph Constructor) Given a point cloud $\boldsymbol{X}$ a* fundamental forms graph constructor *returns an FFG, by first building a mesh and then calculating its geometric features.*

These edge features alone suffice to use the discrete fundamental theorem of surfaces, guaranteeing that there exists a unique embedding $F : \mathcal{V} \to X$, up to a global rigid motion $E(3)$.

Proposition 0.1. *(Complete Invariant) Given two FFGs built on any two distinct-distance point clouds (no two points are the same distance apart), $\mathcal{G} = FFGC(X)$ and $\mathcal{G}' = FFGC(X')$ Assuming a maximally expressive GNN F*

$$F(\mathcal{G}) = F(\mathcal{G}') \iff X \cong X' \text{ up to } E(3)$$

Proof. Suppose a maximally expressive one-layer GNN F assigns the same global feature to $FFG(\mathcal{M})$ and $FFG(\mathcal{M}')$. Then the multisets of accumulated edge features coincide, i.e.

$$\begin{aligned} &\{\!\{ \ \{\!\{ e_{ij} : j \in \mathcal{N}_i \}\!\} : i \in V \ \}\!\} \\ &\qquad = \ \{\!\{ \ \{\!\{ e'_{i'j'} : j' \in \mathcal{N}'_{i'} \}\!\} : i' \in V' \ \}\!\}. \end{aligned}$$

This implies $|\mathcal{V}| = |\mathcal{V}'|$, and after relabeling if necessary, the multisets of incident edge features at each vertex also match. Hence each node $i \in \mathcal{V}$ has the same degree as its counterpart in V', so $|\mathcal{E}| = |\mathcal{E}'|$.

Under the assumption of a distinct-distance point cloud, we assume that each edge length is unique, for any edge $l_{i,j} \in E$ the feature $e_{ij} = \|x_i - x_j\|$ appears exactly once in the incident multisets of i and j. Therefore, i, j must also be present in $\mathcal{E}'$, and thus the underlying graphs $(\mathcal{V}, \mathcal{E})$ and $(\mathcal{V}', \mathcal{E}')$ are equivalent.

Since the graphs share the same connectivity and are both maximally planar by construction, they share the unique set of associated faces $\mathcal{T} = \mathcal{T}'$ (Whitney 1932).

Therefore, both FFGs $\mathcal{M} = (\mathcal{V}, \mathcal{E}, \mathcal{T}; \{e_{ij}\})$ and $\mathcal{M}' = (\mathcal{V}', \mathcal{E}', \mathcal{T}'; \{e'_{ij}\})$ are identical. Now by the local integrability condition we can invoke the *Discrete Fundamental Theorem of Surfaces* to realize a unique immersion. Hence the reconstructed point clouds X and X' coincide up to $E(3)$. The reverse is trivial: if the point clouds coincide, an invariant construction method will result in the same graph, meaning even a maximally expressive GNN F will map to the same global feature.

□

Fundamental Forms Graph Construction

To make FFGs tunable toward preferred representations, we introduce a *graph-point cloud fitting* step posed as a quadratic assignment over permutations. The first step is to choose a template graph. The vertices of this graph are then assigned to a given point cloud with by matching their intrinsic distances. After the assignment is found, the geometric features are calculated.

Template Graph We use meshes obtained from the convex hull of Fibonacci spheres (González 2010). Their near–equal-area node placement yields uniform triangulations with low node-degree variance, and—crucially—supports any prescribed vertex count N. By contrast, grid-based schemes (e.g., equiangular, HEALPix, cube, icosahedral) are only available at specific resolutions; to match a given N requires oversampling to the nearest resolution and then downsampling before retriangulating. Random sampling provides the least regularity.

Graph-point cloud assignment problem To match intrisic distances we consider the following. Let $D \in \mathbb{R}^{n \times n}_{\geq 0}$ be a matrix derived from X (e.g., Euclidean distances $D_{ij} = \|x_i - x_j\|_2$, or any pairwise weighting that emphasizes locality), and let $S \in \mathbb{Z}^{n \times n}_{\geq 0}$ be the graph–geodesic (shortest-path) distance matrix of a template graph. The *graph–point cloud assignment* problem seeks a bijection $P \in \Pi_n$ such that the permuted point-cloud metric $PDP^\top$ best aligns with the template metric S. We formulate this as

$$\min_{P \in \Pi_n} \ \| S - PDP^\top \|_F,$$

which explicitly aligns two metric spaces—graph geodesic and Euclidean—by searching for an (approximate) isometry between them. Intuitively, pairs that are close in D are encouraged to be close in S, i.e., mapped to nearby nodes on the template.

Degenerate Faces Even though the calculation of face normals necessitates that a face $\tau \in \mathcal{T}$, must have a non-zero area, we can still assign it a normal such that our resulting dihedrals ensure local integrability. The only caveat is that this will only work for degenerate faces with unique points.

For any three collinear points $p, q, r \in \mathcal{V}$, choose any unit vector orthogonal to,

$$\pm(x_q - x_p) \text{ or } \pm (x_p - x_r) \text{ or } \pm (x_r - x_q)$$

as the normal for the face. The additional degree of freedom does not impact the uniqueness of the immersion. This can be understood by considering an in-process reconstruction where $p, q \in \mathcal{V}$ have already been placed in $\mathbb{R}^3$, then r will be placed somewhere in the span of $x_p - x_q$, and therefore its position is fixed no matter what transitionary rotation is applied. Now with these normals in place we can calculate dihedrals as before.

Graph Construction Properties

Permutation invariance. Because the objective depends on $\mathbf{X}$ only through D, any unbiased optimizer (or an exact

solver) yields a mapping that is invariant to relabelings of the point cloud: if $Q \in \Pi_n$ reindexes X (so $D \mapsto QDQ^\top$), then $P^\star$ transforms to $P^\star Q^\top$ without changing the value of equation . Thus the resulting fitted graph construction remains permutation invariant to input labeling.

Characterization of Prior Approaches

SCHull Graphs In (Wang et al. 2025) spherical convex-hull (SCHull) graphs are introduced, they are generically complete and simple to build. Let $X \subset \mathbb{R}^{3\times N}$ and $\hat{x}_i = u_i/\|u_i\| \in \mathbb{S}^2$. Take the convex hull of $\{\hat{x}_i\}_{i=1}^n$ in $\mathbb{R}^3$; its boundary induces edges between pairs $(\hat{x}_i, \hat{x}_j)$. The SCHull graph uses edge features $\ell_{ij} = \|x_i - x_j\|$ and ϕ_{ij} equal to the dihedral angle between the two boundary faces meeting along $\hat{x}_i, \hat{x}_j$, and node feature $r_i = \|u_i\|$.

The limitations come from a single static construction per cloud: collisions when $\hat{x}_i = \hat{x}_j$ for $i \neq j$; loss of local Euclidean structure from the direct projection $x_i \mapsto \hat{x}_i$, which motivates added radial-cutoff connections; and as n grows, the dihedrals begin to conform, which weakens local discrimination.

Power Graphs Power graphs are also generically complete (Sverdlov and Dym 2025), but rely on an initial starting graph. In tasks where they are available chemical bonds can fulfill this use, but in such tasks atom positions are not typically generic. Otherwise neither radial cutoff nor K-NN guarantee connectedness, and $\mathcal{G} \mapsto \mathcal{G}^p$ means that sparsity may be compromised depending on the power and sparsity of the underlying graph.

ComENet ComENet (Wang et al. 2022) offers another sparse-complete style by assigning to each edge a dihedral computed from its 2-hop neighborhood, but without a global consistency prescription entire 2-hop neighborhoods can be permuted while preserving these local dihedrals, so global completeness is not guaranteed.

Graph Construction

The FFG optimization is implemented in JAX using OTT-JAX for entropic Gromov–Wasserstein. Each inner step uses Sinkhorn (max 200 iters, tol 10^{-5}) and an outer GW loop (max 50 iters, threshold 10^{-3}) with warm starts. The solver returns a soft transport T that we normalize and a hard permutation P via Hungarian matching on T. The resulting node mapping p induces a standard edge list (and/or cross-graph correspondences) that is directly amenable to PyTorch Geometric loaders and batching.

Protein Cryptic Docking

CryptoBench (Škrhák et al. 2024) is a large, curated benchmark for cryptic binding sites (CBSs)—sites that are not visible/accessible in the apo state but become druggable in the holo state. It's built from matched apo–holo pairs, grouped by UniProt, clustered by sequence identity, and filtered for substantial binding-site rearrangements; the public release reports 1,107 structures with predefined splits to train/evaluate CBS predictors. The benchmark standardizes provides region identifies for the cryptic regime of pocket detection.

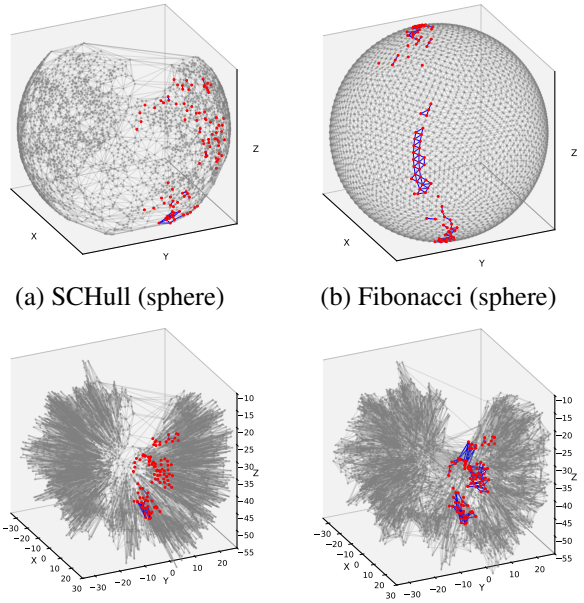

(a) SCHull (sphere) (b) Fibonacci (sphere)

(c) SHull (point cloud) (d) Fibonacci (point cloud)

Figure 2: Graph construction comparison with a protein cryptic pocket atoms highlighted in red and local connections on the pocket in blue. We observe that the SCHull graph appears to have local structure in the point cloud projection, but the projection often separates points in the convex hull graph. By comparison all of the FFGC graphs have better local structure preservation.

Figure 2 illustrates the improved local structure preservation of our FFG construction in comparison the baseline SCHull. In particular, each atom identified in the pocket region is colored red and any connected binding atoms are connected by a blue edge. We observe that although the SCHull graph appears to have local structure, it is primarily an artifact of the projection method. By contrast, the Fibonacci method maintains local structure of the cryptic pocket region. The molecule and its cryptic pocket region are illustrated in Figure 1.

Concluding Remarks

We introduced the Fundamental Forms Graph Construction (FFGC) as a principled framework for graph construction. By encoding edge lengths and dihedral angles subject to local integrability, our approach guarantees sparsity and connectivity while providing features that uniquely determine embeddings up to $E(3)$. This stands in contrast to prevailing constructions such as radial cutoff, k–nearest neighbors, or SCHull graphs, which do not simultaneously ensure these properties and admit direct geometric tunability. We leave to future work the inclusion of a comprehensive empirical evaluation including contact-preservation, Euclidean locality preservation, and downstream cryptic-pocket benchmarks.

References

Barbero, F.; Velingker, A.; Saberi, A.; Bronstein, M.; and Giovanni, F. D. 2024. Locality-Aware Graph-Rewiring in GNNs. arXiv:2310.01668.

Battiloro, C.; Karaismailoglu, E.; Tec, M.; Dasoulas, G.; Audirac, M.; and Dominici, F. 2024. E(n) Equivariant Topological Neural Networks. In *The Thirteenth International Conference on Learning Representations*.

Brandstetter, J.; Hesselink, R.; Pol, E. v. d.; Bekkers, E. J.; and Welling, M. 2021. Geometric and Physical Quantities improve E(3) Equivariant Message Passing. In *International Conference on Learning Representations*.

Bronstein, M. M.; Bruna, J.; Cohen, T.; and Veličković, P. 2021. Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges. ArXiv:2104.13478 [cs].

Damke, C.; Melnikov, V.; and Hüllermeier, E. 2020. A Novel Higher-order Weisfeiler-Lehman Graph Convolution. ArXiv:2007.00346 [cs].

Du, W.; Du, Y.; Wang, L.; Feng, D.; Wang, G.; Ji, S.; Gomes, C. P.; and Ma, Z.-M. 2023. A new perspective on building efficient and expressive 3D equivariant graph neural networks. In *Thirty-seventh Conference on Neural Information Processing Systems*.

Dwivedi, V. P.; Rampášek, L.; Galkin, M.; Parviz, A.; Wolf, G.; Luu, A. T.; and Beaini, D. 2022. Long Range Graph Benchmark. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Fey, M.; Yuen, J.-G.; and Weichert, F. 2020. Hierarchical Inter-Message Passing for Learning on Molecular Graphs. ArXiv:2006.12179 [cs].

González, Á. 2010. Measurement of areas on a sphere using Fibonacci and latitude–longitude lattices. *Mathematical geosciences*, 42(1): 49–64.

Hordan, S.; Amir, T.; and Dym, N. 2024. Weisfeiler Leman for Euclidean Equivariant Machine Learning. *CoRR*, abs/2402.02484.

Hu, W.; Fey, M.; Zitnik, M.; Dong, Y.; Ren, H.; Liu, B.; Catasta, M.; and Leskovec, J. 2020. Open Graph Benchmark: Datasets for Machine Learning on Graphs. *Advances in Neural Information Processing Systems*, 33: 22118–22133.

Igashov, I.; Olechnovic, K.; Kadukova, M.; Venclovas, C.; and Grudinin, S. 2021. VoroCNN: deep convolutional neural network built on 3D Voronoi tessellation of protein structures. *Bioinformatics*, 37(16): 2332–2339.

Jing, B.; Eismann, S.; Suriana, P.; Townshend, R. J. L.; and Dror, R. 2020. Learning from Protein Structure with Geometric Vector Perceptrons. In *International Conference on Learning Representations*.

Joshi, C. K.; Bodnar, C.; Mathis, S. V.; Cohen, T.; and Lio, P. 2023. On the expressive power of geometric graph neural networks. In *International conference on machine learning*, 15330–15355. PMLR.

Liu, C.; Ruhe, D.; Eijkelboom, F.; and Forré, P. 2023. Clifford Group Equivariant Simplicial Message Passing Networks. In *The Twelfth International Conference on Learning Representations*.

Meyer, M.; Desbrun, M.; Schröder, P.; and Barr, A. H. 2003. Discrete Differential-Geometry Operators for Triangulated 2-Manifolds. In *Visualization and Mathematics III*.

Morris, C.; Frasca, F.; Dym, N.; Maron, H.; Ceylan, I. I.; Levie, R.; Lim, D.; Bronstein, M. M.; Grohe, M.; and Jegelka, S. 2024. Position: Future Directions in the Theory of Graph Machine Learning. In *Proceedings of the 41st International Conference on Machine Learning*, 36294–36307. PMLR. ISSN: 2640-3498.

Qian, C.; Manolache, A.; Ahmed, K.; Zeng, Z.; den Broeck, G. V.; Niepert, M.; and Morris, C. 2024. Probabilistically Rewired Message-Passing Neural Networks. arXiv:2310.02156.

Satorras, V. G.; Hoogeboom, E.; and Welling, M. 2021. E(n) Equivariant Graph Neural Networks. In *Proceedings of the 38th International Conference on Machine Learning*, 9323–9332. PMLR. ISSN: 2640-3498.

Schütt, K. T.; Sauceda, H. E.; Kindermans, P.-J.; Tkatchenko, A.; and Müller, K.-R. 2018. SchNet - A deep learning architecture for molecules and materials. *The Journal of chemical physics*, 148(24): 241722.

Sonthalia, R.; Gilbert, A.; and Durham, M. 2023. RelWire: Metric Based Graph Rewiring. In *NeurIPS 2023 Workshop on Symmetry and Geometry in Neural Representations*.

Sverdlov, Y.; and Dym, N. 2025. On the Expressive Power of Sparse Geometric MPNNs. In *The Thirteenth International Conference on Learning Representations*.

Tahmasebi, B.; Lim, D.; and Jegelka, S. 2023. The Power of Recursion in Graph Neural Networks for Counting Substructures. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, 11023–11042. PMLR. ISSN: 2640-3498.

Thomas, N.; Smidt, T.; Kearnes, S.; Yang, L.; Li, L.; Kohlhoff, K.; and Riley, P. 2018. Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds. ArXiv:1802.08219 [cs].

Topping, J.; Giovanni, F. D.; Chamberlain, B. P.; Dong, X.; and Bronstein, M. M. 2021. Understanding over-squashing and bottlenecks on graphs via curvature. In *International Conference on Learning Representations*.

Wang, L.; Liu, Y.; Lin, Y.; Liu, H.; and Ji, S. 2022. ComENet: Towards Complete and Efficient Message Passing for 3D Molecular Graphs. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 650–664. Curran Associates, Inc.

Wang, S.-H.; Huang, Y.; Baker, J. M.; Sun, Y.-E.; Tang, Q.; and Wang, B. 2025. A Theoretically-Principled Sparse, Connected, and Rigid Graph Representation of Molecules. In *The Thirteenth International Conference on Learning Representations*.

Wang, Y.; Liu, B.; and Tong, Y. 2012. Linear surface reconstruction from discrete fundamental forms on triangle meshes. In *Computer Graphics Forum*, volume 31, 2277–2287. Wiley Online Library. Issue: 8.

Weisfeiler, B. Y.; and Leman, A. A. 1968. A Reduction of a Graph to a Canonical Form and an Algebra Arising During This Reduction. *Nauchno-Technicheskaya Informatsiya*, none(9): 12–16.

Whitney, H. 1932. Congruent Graphs and the Connectivity of Graphs. *American Journal of Mathematics*, 54(1): 150–168.

Wollschläger, T.; Kemper, N.; Hetzel, L.; Sommer, J.; and Günnemann, S. 2024. Expressivity and Generalization: Fragment-Biases for Molecular GNNs. In *Proceedings of the 41st International Conference on Machine Learning*, 53113–53139. PMLR. ISSN: 2640-3498.

Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? In *International Conference on Learning Representations*.

Škrhák, V.; Novotný, M.; Feidakis, C. P.; Krivák, R.; and Hoksza, D. 2024. CryptoBench: Cryptic protein-ligand binding sites dataset and benchmark. *Bioinformatics*, btae745.