
On the Expressive Power of Geometric Graph Neural Networks

Anonymous Author(s)

Anonymous Affiliation

Anonymous Email

Abstract

We propose a geometric version of the Weisfeiler-Leman graph isomorphism test (GWL) for discriminating geometric graphs while respecting the underlying physical symmetries: permutations, rotation, reflection, and translation. We use GWL to characterise the expressive power of Graph Neural Networks (GNNs) that are invariant or equivariant to physical symmetries in terms of the classes of geometric graphs they can distinguish. This allows us to formalise the advantages of equivariant GNNs over invariant GNNs: **equivariant layers have greater expressive power as they enable propagating geometric information beyond local neighbourhoods, while invariant layers only reason locally via scalars and cannot discriminate geometric graphs with different non-local properties.**

1 Introduction

The graph isomorphism problem and the Weisfeiler-Leman (WL) [1] test for distinguishing non-isomorphic graphs have become powerful tools for analysing the expressive power of Graph Neural Networks (GNNs) [2, 3]. The WL framework has been a major driver of progress for more expressive GNNs [4–8]. However, WL does not directly apply to the increasingly relevant special case of *geometric graphs* – graphs embedded in Euclidean space – which come equipped with a stronger notion of isomorphism that also takes spatial symmetries into account. The lack of theoretical tools is becoming more apparent as geometric graphs are increasingly used to model systems in biochemistry [9], material science [10], physical simulations [11], and multiagent robotics [12]. Graph Neural Networks (GNNs) with Euclidean symmetries ‘baked in’ have emerged as the architecture of choice for these domains [13].

Addressed
R2.1, R3.1,
R4.1, R4.12

Geometric GNNs follow the message passing paradigm [14] where node features are updated in a permutation equivariant manner by aggregating features from local neighbourhoods. In addition to the permutation group, the geometric attributes of the nodes (*e.g.* coordinates, velocity) transform along with Euclidean transformations of the system, *i.e.* they are equivariant to a Lie group such as the group of rotations ($SO(d)$) or rotations and reflections ($O(d)$). We use $\mathfrak{G}$ as a generic symbol for such a Lie group. Based on this, we consider two classes of GNNs for geometric graphs: (1) **$\mathfrak{G}$ -equivariant models**, where the intermediate features and propagated messages are $\mathfrak{G}$ -equivariant geometric quantities such as vectors or tensors [15–19]; and (2) **$\mathfrak{G}$ -invariant models**, which only propagate $\mathfrak{G}$ -invariant scalar features such as distances and angles [20–22]. Despite promising empirical results for both classes of architectures, key theoretical questions remain unanswered: (1) How to characterise the *expressive power* of geometric GNNs? And (2) what is the tradeoff between $\mathfrak{G}$ -equivariant and $\mathfrak{G}$ -invariant GNNs?

Addressed
R2.3, R4.8

Contributions. In this work, we study the expressive power of geometric GNNs from the perspective of discriminating non-isomorphic geometric graphs. We propose a geometric version of the Weisfeiler-Leman graph isomorphism test, termed GWL. (Figure 1). We use GWL to formally characterise classes of graphs that can and cannot be distinguished by $\mathfrak{G}$ -invariant and $\mathfrak{G}$ -equivariant GNNs. We show how invariant models have limited expressive power as they only reason locally via scalar quantities, while equivariant models distinguish a larger class of graphs by propagating geometric vector quantities beyond local neighbourhoods.

For **Background and Preliminaries**, please see Appendix A.




Figure 1: Geometric Weisfeiler-Leman Test. GWL distinguishes non-isomorphic geometric graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ by injectively assigning colours to distinct neighbourhood patterns, up to global symmetries (here $\mathfrak{G} = O(d)$). Each iteration expands the neighbourhood from which geometric information can be gathered (shaded for node i). Example inspired by [23].

2 The Geometric Weisfeiler-Leman Test

Assumptions. Analogous to the WL test, the geometric and scalar features the nodes are equipped with come from countable subsets $C \subset \mathbb{R}^d$ and $C' \subset \mathbb{R}$, respectively. As a consequence, when we require functions to be injective, we require them to be injective over the countable set of $\mathfrak{G}$ -orbits that are obtained by acting with the symmetry group $\mathfrak{G}$ on the dataset. This mimics the practically relevant situation of finite datasets, in which we have a finite pool $\mathcal{P}$ of geometric graphs (and their symmetry transformations) which we would like to distinguish.

Intuition. For an intuition of how to generalise the WL test to geometric graphs, we note that WL uses a local, node-centric, procedure to update the colour of each node i using the colours of its the 1-hop neighbourhood $\mathcal{N}_i$. In the geometric setting, $\mathcal{N}_i$ is an attributed point cloud around the central node i . As a result, each neighbourhood carries two types of information: (1) neighbourhood type (invariant to $\mathfrak{G}$) and (2) neighbourhood geometric orientation (equivariant to $\mathfrak{G}$). From an axiomatic point of view, our generalisation of the WL neighbourhood aggregation procedure must meet two properties:

Property 1: Orbit injectivity of colours. If two neighbourhoods are the same up to an action of $\mathfrak{G}$ (e.g. rotation), then the colours of the corresponding central nodes should be the same. Thus, the colouring must be $\mathfrak{G}$ -orbit injective – which also makes it $\mathfrak{G}$ -invariant – over the countable set of all orbits of neighbourhoods in our dataset.

Property 2: Preservation of local geometry. A key property of WL is that the aggregation is injective. A $\mathfrak{G}$ -invariant colouring procedure that purely satisfies Property 1 is not sufficient because, by definition, it loses spatial properties of each neighbourhood such as the relative pose or orientation [24]. Thus, we must additionally update auxiliary *geometric information* variables in a way that is $\mathfrak{G}$ -equivariant and injective.

Geometric Weisfeiler-Leman (GWL). These intuitions motivate the following definition of the GWL test. At initialisation, we assign to each node $i \in \mathcal{V}$ a scalar node colour $c_i \in C'$ and an auxiliary object g_i containing the geometric information associated to it:

$$c_i^{(0)} := \text{HASH}(s_i), \quad g_i^{(0)} := (c_i^{(0)}, \vec{v}_i), \quad (1)$$

where HASH denotes an injective map over the scalar attributes s_i of node i . To define the inductive step, assume we have the colours of the nodes and the associated geometric objects at iteration $t - 1$. Then, we can aggregate the geometric information around node i into a new object as follows:

$$g_i^{(t)} := \left((c_i^{(t-1)}, g_i^{(t-1)}), \{ (c_j^{(t-1)}, g_j^{(t-1)}, \vec{x}_{ij}) \mid j \in \mathcal{N}_i \} \right), \quad (2)$$

Here $\{\cdot\}$ denotes a multiset – a set in which elements may occur more than once. Importantly, the group $\mathfrak{G}$ can act on the geometric objects above inductively by acting on the geometric information

Addressed
R4.10

Addressed
R2.1, R4.5

Addressed
R4.11

Addressed
R4.9

inside it. This amounts to rotating (or reflecting) the entire t -hop neighbourhood contained inside:

$$\mathfrak{g} \cdot \mathbf{g}_i^{(0)} := (c_i^{(0)}, \mathbf{Q}_{\mathfrak{g}} \vec{v}_i), \quad \mathfrak{g} \cdot \mathbf{g}_i^{(t)} := ((c_i^{(t-1)}, \mathfrak{g} \cdot \mathbf{g}_i^{(t-1)}), \{ (c_j^{(t-1)}, \mathfrak{g} \cdot \mathbf{g}_j^{(t-1)}, \mathbf{Q}_{\mathfrak{g}} \vec{x}_{ij}) \mid j \in \mathcal{N}_i \})$$

Clearly, the aggregation building $\mathbf{g}_i$ for any time-step t is injective and $\mathfrak{G}$ -equivariant. Finally, we can compute the node colours at iteration t for all $i \in \mathcal{V}$ by aggregating the geometric information in the neighbourhood around node i :

$$c_i^{(t)} := \text{I-HASH}^{(t)}(\mathbf{g}_i^{(t)}), \quad (3)$$

by using a $\mathfrak{G}$ -orbit injective and $\mathfrak{G}$ -invariant function that we denote by I-HASH. That is for any geometric objects $\mathbf{g}, \mathbf{g}'$, $\text{I-HASH}(\mathbf{g}) = \text{I-HASH}(\mathbf{g}')$ if and only if there exists $\mathfrak{g} \in \mathfrak{G}$ such that $\mathbf{g} = \mathfrak{g} \cdot \mathbf{g}'$.

Overview. With each iteration, $\mathbf{g}_i^{(t)}$ aggregates geometric information in progressively larger t -hop subgraph neighbourhoods $\mathcal{N}_i^{(t)}$ around the node i . The node colours summarise the structure of these t -hops via the $\mathfrak{G}$ -invariant aggregation performed by I-HASH. The procedure terminates when the partitions of the nodes induced by the colours do not change from the previous iteration. Finally, given two geometric graphs $\mathcal{G}$ and $\mathcal{H}$, if there exists some iteration t for which $\{c_i^{(t)} \mid i \in \mathcal{V}(\mathcal{G})\} \neq \{c_i^{(t)} \mid i \in \mathcal{V}(\mathcal{H})\}$, then GWL deems the two graphs as being geometrically non-isomorphic. Otherwise, we say the test cannot distinguish the two graphs.

Invariant GWL. Since we are interested in understanding the role of $\mathfrak{G}$ -equivariance, we also consider a more restrictive Invariant GWL (IGWL) that only updates node colours using the $\mathfrak{G}$ -orbit injective I-HASH function and does not propagate geometric information:

$$c_i^{(t)} := \text{I-HASH} \left((c_i^{(t-1)}, \vec{v}_i), \{ (c_j^{(t-1)}, \vec{v}_j, \vec{x}_{ij}) \mid j \in \mathcal{N}_i \} \right). \quad (4)$$

IGWL with k -body scalars. In order to further analyse the construction of the node colouring function I-HASH, we consider $\text{IGWL}_{(k)}$ based on the maximum number of nodes involved in the computation of $\mathfrak{G}$ -invariant scalars (also known as the ‘body order’ [25]):

$$c_i^{(t)} := \text{I-HASH}_{(k)} \left((c_i^{(t-1)}, \vec{v}_i), \{ (c_j^{(t-1)}, \vec{v}_j, \vec{x}_{ij}) \mid j \in \mathcal{N}_i \} \right), \quad (5)$$

and $\text{I-HASH}_{(k+1)}$ is defined as:

$$\text{HASH} \left(\{ \text{I-HASH} \left((c_i^{(t-1)}, \vec{v}_i), \{ (c_{j_1}^{(t-1)}, \vec{v}_{j_1}, \vec{x}_{ij_1}), \dots, (c_{j_k}^{(t-1)}, \vec{v}_{j_k}, \vec{x}_{ij_k}) \} \right) \mid \mathbf{j} \in (\mathcal{N}_i)^k \} \right),$$

where $\mathbf{j} = [j_1, \dots, j_k]$ are all possible k -tuples formed of elements of $\mathcal{N}_i$. Therefore, $\text{IGWL}_{(k)}$ is now constrained to extract information only from all the possible k -sized tuples of nodes (including the central node) in a neighbourhood. For instance, $\text{I-HASH}_{(2)}$ can identify neighbourhoods only up to pairwise distances among the central node and any of its neighbours (*i.e.* a 2-body scalar), while $\text{I-HASH}_{(3)}$ up to distances and angles formed by any two edges (*i.e.* a 3-body scalar). Notably, distances and angles alone are incomplete descriptors of local geometry [26, 27]. Therefore, $\text{I-HASH}_{(k)}$ with lower k makes the colouring weaker.

3 Characterising the Expressive Power of GWL

3.1 What Geometric Graphs can GWL and IGWL Distinguish?

In order to formalise the expressive power of GWL and IGWL, let us consider what geometric graphs can and cannot be distinguished by the tests. As a simple first observation, we note that when all coordinates and vectors are set equal to zero GWL coincides with the standard 1-WL. In this edge case, GWL has the same expressive power as 1-WL.

Next, let us consider the simplified setting of two geometric graphs $\mathcal{G}_1 = (\mathbf{A}_1, \mathbf{S}_1, \vec{V}_1, \vec{X}_1)$ and $\mathcal{G}_2 = (\mathbf{A}_2, \mathbf{S}_2, \vec{V}_2, \vec{X}_2)$ such that the underlying attributed graphs $(\mathbf{A}_1, \mathbf{S}_1)$ and $(\mathbf{A}_2, \mathbf{S}_2)$ are isomorphic. This case frequently occurs in (bio)chemical modelling, where molecules occur in different conformations, but with the same graph topology given by the covalent bonding structure.

Addressed
R4.6

Addressed
R4.13

Recall that each iteration of GWL aggregates geometric information $g_i^{(k)}$ from progressively larger neighbourhoods $\mathcal{N}_i^{(k)}$ around the node i , and distinguishes (sub-)graphs via comparing $\mathfrak{G}$ -orbit injective colouring of $g_i^{(k)}$. We say $\mathcal{G}_1$ and $\mathcal{G}_2$ are k -hop distinct if for all graph isomorphisms b , there is some node $i \in \mathcal{V}_1, b(i) \in \mathcal{V}_2$ such that the corresponding k -hop subgraphs $\mathcal{N}_i^{(k)}$ and $\mathcal{N}_{b(i)}^{(k)}$ are distinct. Otherwise, we say $\mathcal{G}_1$ and $\mathcal{G}_2$ are k -hop identical if all $\mathcal{N}_i^{(k)}$ and $\mathcal{N}_{b(i)}^{(k)}$ are identical up to group actions. We can now formalise what geometric graphs can and cannot be distinguished by GWL.

Proposition 1. *GWL can distinguish any k -hop distinct geometric graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ where the underlying attributed graphs are isomorphic, and k iterations are sufficient.*

Proposition 2. *Up to k iterations, GWL cannot distinguish any k -hop identical geometric graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ where the underlying attributed graphs are isomorphic.*

Additionally, we can state the following results about the more constrained IGWL.

Proposition 3. *IGWL can distinguish any 1-hop distinct geometric graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ where the underlying attributed graphs are isomorphic, and 1 iteration is sufficient.*

Proposition 4. *Any number of iterations of IGWL cannot distinguish any 1-hop identical geometric graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ where the underlying attributed graphs are isomorphic.*

We can now consider the more general case where the underlying attributed graphs for $\mathcal{G}_1 = (\mathcal{A}_1, \mathcal{S}_1, \vec{\mathcal{V}}_1, \vec{\mathcal{X}}_1)$ and $\mathcal{G}_2 = (\mathcal{A}_2, \mathcal{S}_2, \vec{\mathcal{V}}_2, \vec{\mathcal{X}}_2)$ are non-isomorphic and constructed from point clouds using radial cutoffs, as conventional for biochemistry and material science applications.

Proposition 5. *Assuming geometric graphs are constructed from point clouds using radial cutoffs, GWL can distinguish any geometric graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ where the underlying attributed graphs are non-isomorphic. At most k_{Max} iterations are sufficient, where k_{Max} is the maximum graph diameter among $\mathcal{G}_1$ and $\mathcal{G}_2$.*

These results enable us to compare the expressive powers of GWL and IGWL.

Theorem 6. *GWL is strictly more powerful than IGWL.*

This statement formalises the advantage of $\mathfrak{G}$ -equivariant intermediate layers for graphs and geometric data, as prescribed in the Geometric Deep Learning blueprint [28], in addition to echoing similar intuitions in the computer vision community. As remarked by [24], translation invariant models do not understand the relationship between the various parts of an image (colloquially called the ‘‘Picasso problem’’). Similarly, our results explain how IGWL fails to understand how the various 1-hops of a graph are stitched together. Finally, we identify a setting where this distinction between the two approaches disappears.

Proposition 7. *IGWL has the same expressive power as GWL for fully connected geometric graphs.*

3.2 Characterising the Expressive Power of Geometric GNNs

We would like to characterise the maximum expressive power of geometric GNNs based on the GWL test. Firstly, we show that any message passing $\mathfrak{G}$ -equivariant GNN can be at most as powerful as GWL in distinguishing non-isomorphic geometric (sub-)graphs.

Theorem 8. *Any pair of geometric graphs distinguishable by a $\mathfrak{G}$ -equivariant GNN is also distinguishable by GWL.*

With a sufficient number of iterations, the output of $\mathfrak{G}$ -equivariant GNNs can be equivalent to GWL if certain conditions are met regarding the aggregate, update and readout functions.

Proposition 9. *$\mathfrak{G}$ -equivariant GNNs have the same expressive power as GWL if the following conditions hold: (1) The aggregation AGG is an injective, $\mathfrak{G}$ -equivariant multiset function. (2) The scalar part of the update UPD_s is a $\mathfrak{G}$ -orbit injective, $\mathfrak{G}$ -invariant multiset function. (3) The vector part of the update UPD_v is an injective, $\mathfrak{G}$ -equivariant multiset function. (4) The graph-level readout f is an injective multiset function.*

Similar statements can be made for $\mathfrak{G}$ -invariant GNNs and IGWL. Thus, we can directly transfer our results about GWL and IGWL to the class of GNNs bounded by the respective tests. This has several interesting practical implications, discussed in [Appendix E](#).

Addressed
R4.2

Addressed
R2.2, R2.6,
R4.4

References

- [1] Boris Weisfeiler and Andrei Leman. The reduction of a graph to canonical form and the algebra which appears therein. *NTI, Series*, 1968. 1, 7
- [2] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *ICLR*, 2019. 1, 7, 12
- [3] Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks. In *AAAI*, 2019. 1, 7, 12
- [4] Zhengdao Chen, Soledad Villar, Lei Chen, and Joan Bruna. On the equivalence between graph isomorphism testing and function approximation with gnns. *NeurIPS*, 2019. 1, 7
- [5] Hagga Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably powerful graph networks. *NeurIPS*, 2019.
- [6] Christopher Morris, Gaurav Rattan, and Petra Mutzel. Weisfeiler and leman go sparse: Towards scalable higher-order graph embeddings. *NeurIPS*, 2020.
- [7] Cristian Bodnar, Fabrizio Frasca, Yuguang Wang, Nina Otter, Guido F Montufar, Pietro Lio, and Michael Bronstein. Weisfeiler and lehman go topological: Message passing simplicial networks. In *ICML*, 2021.
- [8] Cristian Bodnar, Fabrizio Frasca, Nina Otter, Yuguang Wang, Pietro Lio, Guido F Montufar, and Michael Bronstein. Weisfeiler and lehman go cellular: Cw networks. *NeurIPS*, 2021. 1, 7
- [9] Arian Rokkum Jamasb, Ramon Viñas Torné, Eric J Ma, Yuanqi Du, Charles Harris, Kexin Huang, Dominic Hall, Pietro Lio, and Tom Leon Blundell. Graphein - a python library for geometric deep learning and network analysis on biomolecular structures and interaction networks. In *NeurIPS*, 2022. 1, 8
- [10] Lowik Chanussot, Abhishek Das, Siddharth Goyal, Thibaut Lavril, Muhammed Shuaibi, Morgane Riviere, Kevin Tran, Javier Heras-Domingo, Caleb Ho, Weihua Hu, Aini Palizhati, Anuroop Sriram, Brandon Wood, Junwoong Yoon, Devi Parikh, C. Lawrence Zitnick, and Zachary Ulissi. Open catalyst 2020 (oc20) dataset and community challenges. *ACS Catalysis*, 2021. 1, 8
- [11] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *ICML*, 2020. 1, 8
- [12] Qingbiao Li, Fernando Gama, Alejandro Ribeiro, and Amanda Prorok. Graph neural networks for decentralized multi-robot path planning. In *IROS*, 2020. 1, 8
- [13] Mario Geiger and Tess Smidt. e3nn: Euclidean neural networks. *arXiv preprint*, 2022. 1
- [14] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *ICML*, 2017. 1
- [15] Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint*, 2018. 1, 8, 15
- [16] Brandon Anderson, Truong Son Hy, and Risi Kondor. Cormorant: Covariant molecular neural networks. *NeurIPS*, 2019.
- [17] Bowen Jing, Stephan Eismann, Patricia Suriana, Raphael John Lamarre Townshend, and Ron Dror. Learning from protein structure with geometric vector perceptrons. In *ICLR*, 2020. 9
- [18] Victor Garcia Satorras, Emiel Hoogetboom, and Max Welling. E (n) equivariant graph neural networks. In *ICML*, 2021. 15
- [19] 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 *ICLR*, 2022. 1, 8
- [20] Kristof T Schütt, Huziel E Sauceda, P-J Kindermans, Alexandre Tkatchenko, and K-R Müller. Schnet—a deep learning architecture for molecules and materials. *The Journal of Chemical Physics*, 148(24):241722, 2018. 1, 8, 15
- [21] Tian Xie and Jeffrey C. Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.*, 2018. 15

- [22] Johannes Gasteiger, Janek Groß, and Stephan Günnemann. Directional message passing for molecular graphs. In *ICLR*, 2020. 1, 8, 15
- [23] Kristof Schütt, Oliver Unke, and Michael Gastegger. Equivariant message passing for the prediction of tensorial properties and molecular spectra. In *ICML*, 2021. 2, 8, 14, 16
- [24] Geoffrey Hinton, Alex Krizhevsky, and Sida D Wang. Transforming auto-encoders. In *ICANN*, 2011. 2, 4
- [25] Ilyes Batatia, Simon Batzner, Dávid Péter Kovács, Albert Musaelian, Gregor NC Simm, Ralf Drautz, Christoph Ortner, Boris Kozinsky, and Gábor Csányi. The design space of e (3)-equivariant atom-centered interatomic potentials. *arXiv preprint*, 2022. 3, 14
- [26] Albert P Bartók, Risi Kondor, and Gábor Csányi. On representing chemical environments. *Physical Review B*, 2013. 3, 8
- [27] Sergey N Pozdnyakov, Michael J Willatt, Albert P Bartók, Christoph Ortner, Gábor Csányi, and Michele Ceriotti. Incompleteness of atomic structure representations. *Physical Review Letters*, 2020. 3, 8, 9, 16, 17
- [28] Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint*, 2021. 4
- [29] Ronald C Read and Derek G Corneil. The graph isomorphism disease. *Journal of graph theory*, 1977. 7
- [30] Anthony Zee. *Group theory in a nutshell for physicists*. Princeton University Press, 2016. 7
- [31] Zuobai Zhang, Minghao Xu, Arian Jamasb, Vijil Chenthamarakshan, Aurelie Lozano, Payel Das, and Jian Tang. Protein representation learning by geometric structure pretraining. *arXiv preprint*, 2022. 8, 9, 14
- [32] Justas Dauparas, Ivan Anishchenko, Nathaniel Bennett, Hua Bai, Robert J Ragotte, Lukas F Milles, Basile IM Wicky, Alexis Coubet, Rob J de Haas, Neville Bethel, et al. Robust deep learning based protein sequence design using proteinmpnn. *bioRxiv*, 2022. 8
- [33] Johannes Gasteiger, Florian Becker, and Stephan Günnemann. Gemnet: Universal directional graph neural networks for molecules. In *NeurIPS*, 2021. 8, 9, 14
- [34] Yu Shi, Shuxin Zheng, Guolin Ke, Yifei Shen, Jiacheng You, Jiyan He, Shengjie Luo, Chang Liu, Di He, and Tie-Yan Liu. Benchmarking graphormer on large-scale molecular modeling datasets. *arXiv preprint*, 2022. 8, 9, 14
- [35] Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E Smidt, and Boris Kozinsky. E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature communications*, 2022. 8
- [36] Ilyes Batatia, Dávid Péter Kovács, Gregor NC Simm, Christoph Ortner, and Gábor Csányi. Mace: Higher order equivariant message passing neural networks for fast and accurate force fields. *arXiv preprint*, 2022. 8, 15
- [37] Chaitanya Joshi. Transformers are graph neural networks. *The Gradient*, 2020. 9, 14
- [38] Alexander V Shapeev. Moment tensor potentials: A class of systematically improvable interatomic potentials. *Multiscale Modeling & Simulation*, 2016. 9
- [39] Ralf Drautz. Atomic cluster expansion for accurate and transferable interatomic potentials. *Physical Review B*, 2019. 15
- [40] Genevieve Dussan, Markus Bachmayr, Gabor Csanyi, Ralf Drautz, Simon Etter, Cas van der Oord, and Christoph Ortner. Atomic cluster expansion: Completeness, efficiency and stability. *arXiv preprint*, 2019. 9
- [41] Nadav Dym and Haggai Maron. On the universality of rotation equivariant point cloud networks. In *ICLR*, 2020. 9
- [42] Julian J. McAuley, Teófilo de Campos, and Tiberio S. Caetano. Unified graph matching in euclidean spaces. In *CVPR*, 2010. 9
- [43] Vikas Garg, Stefanie Jegelka, and Tommi Jaakkola. Generalization and representational limits of graph neural networks. In *ICML*, 2020. 14, 16
- [44] Sergey N Pozdnyakov and Michele Ceriotti. Incompleteness of graph convolutional neural networks for points clouds in three dimensions. *arXiv preprint*, 2022. 15

A Background

A.1 Graph Isomorphism and the Weisfeiler-Leman Test

 Addressed
R3.1

An attributed graph $\mathcal{G} = (\mathcal{A}, \mathcal{S})$ with a node set $\mathcal{V}$ of size n consists of an $n \times n$ adjacency matrix $\mathcal{A}$ and a matrix of scalar features $\mathcal{S} \in \mathbb{R}^{n \times f}$. Two attributed graphs $\mathcal{G}, \mathcal{H}$ are *isomorphic* if there exists an edge-preserving bijection $b : \mathcal{V}(\mathcal{G}) \rightarrow \mathcal{V}(\mathcal{H})$ such that $s_i^{(\mathcal{G})} = s_{b(i)}^{(\mathcal{H})}$, where the subscripts index rows and columns in the corresponding matrices.

The *Weisfeiler-Leman test* (WL) is an algorithm for testing whether two (attributed) graphs are isomorphic [1, 29]. At iteration zero the algorithm assigns a *colour* $c_i^{(0)} \in C$ from a countable space of colours C to each node i . Nodes are coloured the same if their features are the same, otherwise, they are coloured differently. In subsequent iterations t , WL iteratively updates the node colouring by producing a new $c_i^{(t)} \in C$:

$$c_i^{(t)} := \text{HASH} \left(c_i^{(t-1)}, \{ \{ c_j^{(t-1)} \mid j \in \mathcal{N}_i \} \} \right), \quad (6)$$

where HASH is an injective map (i.e. a perfect hash map) that assigns a unique colour to each input and $\{ \{ \cdot \} \}$ denotes a multiset – a set that allows for repeated elements. The test terminates when the partition of the nodes induced by the colours becomes stable. Given two graphs $\mathcal{G}$ and $\mathcal{H}$, if there exists some iteration t for which $\{ \{ c_i^{(t)} \mid i \in \mathcal{V}(\mathcal{G}) \} \} \neq \{ \{ c_i^{(t)} \mid i \in \mathcal{V}(\mathcal{H}) \} \}$, then the graphs are not isomorphic. Otherwise, the WL test is inconclusive, and we say it cannot distinguish the two graphs.

Ever since Xu et al. [2], Morris et al. [3] noticed that Graph Neural Networks are at most as powerful as the Weisfeiler-Leman (WL) [1] test at distinguishing non-isomorphic graphs, the WL hierarchy became a powerful tool for analysing the expressive power of GNNs and guided the search for more expressive models [4–8].

A.2 Group Theory

We assume basic familiarity with group theory, see [30] for an overview. We denote the action of the group $\mathfrak{G}$ on a space X by $\mathfrak{g} \cdot x$. If $\mathfrak{G}$ acts on spaces X and Y , we say a function $f : X \rightarrow Y$ is $\mathfrak{G}$ -*equivariant* if $f(\mathfrak{g} \cdot x) = \mathfrak{g} \cdot f(x)$. A function $f : X \rightarrow Y$ is $\mathfrak{G}$ -*invariant* if $f(\mathfrak{g} \cdot x) = f(x)$. The $\mathfrak{G}$ -*orbit* of $x \in X$ is $\mathcal{O}_{\mathfrak{G}}(x) = \{ \mathfrak{g} \cdot x \mid \mathfrak{g} \in \mathfrak{G} \} \subseteq X$. When x and x' are part of the same orbit, we write $x \simeq x'$. We say a function $f : X \rightarrow Y$ is $\mathfrak{G}$ -*orbit injective* if we have $f(x_1) = f(x_2)$ if and only if $x_1 \simeq x_2$ for any $x_1, x_2 \in X$. Necessarily, such a function is $\mathfrak{G}$ -invariant, since $f(\mathfrak{g} \cdot x) = f(x)$.

We work with the permutation group over n elements S_n and the Lie groups $\mathfrak{G} = SO(d)$ or $\mathfrak{G} = O(d)$. Invariance to the translation group $T(d)$ is conventionally handled using relative positions or by subtracting the centre of mass from all nodes positions. Given one of the standard groups above, for an element $\mathfrak{g}$ we denote by $M_{\mathfrak{g}}$ (or another capital letter) its standard matrix representation.

 Addressed
R2.3, R2.8,
R4.8

A.3 Geometric Graphs

A geometric graph $\mathcal{G} = (\mathcal{A}, \mathcal{S}, \vec{\mathcal{V}}, \vec{\mathcal{X}})$ with a node set $\mathcal{V}$ is an attributed graph that is also decorated with geometric attributes: node coordinates $\vec{\mathcal{X}} \in \mathbb{R}^{n \times d}$ and (optionally) vector features $\vec{\mathcal{V}} \in \mathbb{R}^{n \times d}$ (e.g. velocity, acceleration). **Without loss of generality, we work with a single vector feature per node. Our results generalise to multiple vector features or higher-order geometric tensors per node.**

 Addressed
R2.7

The geometric attributes transform as follows under the action of the relevant groups: (1) S_n acts on the graph via $P_{\sigma} \mathcal{G} := (P_{\sigma} \mathcal{A} P_{\sigma}^{\top}, P_{\sigma} \mathcal{S}, P_{\sigma} \vec{\mathcal{V}}, P_{\sigma} \vec{\mathcal{X}})$; (2) Orthogonal transformations $\mathbf{Q}_{\mathfrak{g}} \in \mathfrak{G}$ act on $\vec{\mathcal{V}}, \vec{\mathcal{X}}$ via $\vec{\mathcal{V}} \mathbf{Q}_{\mathfrak{g}}, \vec{\mathcal{X}} \mathbf{Q}_{\mathfrak{g}}$; and (3) Translations $\vec{t} \in T(d)$ act on the coordinates $\vec{\mathcal{X}}$ via $\vec{x}_i + \vec{t}$ for all nodes i .

Two geometric graphs $\mathcal{G}$ and $\mathcal{H}$ are *geometrically isomorphic* (denoted $\mathcal{G} \simeq \mathcal{H}$) if there exists an attributed graph isomorphism b such that the geometric attributes are equivalent, up to global group actions $\mathbf{Q}_{\mathfrak{g}} \in \mathfrak{G}$ and $\vec{t} \in T(d)$:

 Addressed
R2.3

$$(s_i^{(\mathcal{G})}, \vec{v}_i^{(\mathcal{G})}, \vec{x}_i^{(\mathcal{G})}) = (s_{b(i)}^{(\mathcal{H})}, \mathbf{Q}_{\mathfrak{g}} \vec{v}_{b(i)}^{(\mathcal{H})}, \mathbf{Q}_{\mathfrak{g}} (\vec{x}_{b(i)}^{(\mathcal{H})} + \vec{t})) \quad \text{for all } i \in \mathcal{V}(\mathcal{G}). \quad (7)$$

Note that if two geometric graphs are geometrically isomorphic, they are also isomorphic as attributed graphs. However, the converse is not true.

Systems in biochemistry [9], material science [10], physical simulations [11], and multiagent robotics [12] are conventionally modelled as geometric graphs. For example, molecules are represented as a set of nodes corresponding to atoms, which contain information about the atom type as well as its 3D spatial coordinates and other geometric quantities such as velocity or acceleration. The geometric attributes transform along with Euclidean transformations of the system. In biochemistry and material science, the conventional procedure for constructing the geometric graph $\mathcal{G} = (\mathbf{A}, \mathbf{S}, \mathbf{\bar{V}}, \mathbf{\bar{X}})$ is via the underlying point cloud $(\mathbf{S}, \mathbf{\bar{V}}, \mathbf{\bar{X}})$ using a predetermined radial cutoff r . Thus, the adjacency matrix is defined as $a_{ij} = 1$ if $\|\vec{x}_i - \vec{x}_j\|_2 \leq r$, or 0 otherwise, for all $a_{ij} \in \mathbf{A}$.

 Addressed
R2.7, R3.1

Geometric graph isomorphism and distinguishing (sub-)graph geometries has important practical implications for representation learning. For *e.g.*, in molecular systems, an ideal architecture should map distinct local structural environments around atoms to distinct embeddings in representation space [26, 27].

 Addressed
R4.2

A.4 Geometric Graph Neural Networks

We consider two broad classes of geometric GNN architectures. $\mathfrak{G}$ -equivariant GNN layers [15–19] update scalar and vector features from iteration t to $t + 1$ via learnable aggregate and update functions, AGG and UPD, respectively:

$$\mathbf{m}_i^{(t)}, \vec{\mathbf{m}}_i^{(t)} := \text{AGG} \left(\{ (s_j^{(t)}, \vec{s}_j^{(t)}, \vec{v}_j^{(t)}, \vec{x}_{ij}) \mid j \in \mathcal{N}_i \} \right) \quad (\text{Aggregate}) \quad (8)$$

$$\mathbf{s}_i^{(t+1)}, \vec{v}_i^{(t+1)} := \text{UPD} \left((s_i^{(t)}, \vec{v}_i^{(t)}), (\mathbf{m}_i^{(t)}, \vec{\mathbf{m}}_i^{(t)}) \right) \quad (\text{Update}) \quad (9)$$

For *e.g.*, PaiNN [23] interaction layers aggregate scalar and vector features via learnt radial filters:

 Addressed
R2.2

$$s_i^{(t+1)} := s_i^{(t)} + \sum_{j \in \mathcal{N}_i} f_1 \left(s_j^{(t)}, \|\vec{x}_{ij}\| \right) \quad (10)$$

$$\vec{v}_i^{(t+1)} := \vec{v}_i^{(t)} + \sum_{j \in \mathcal{N}_i} f_2 \left(s_j^{(t)}, \|\vec{x}_{ij}\| \right) \odot \vec{v}_j^{(t)} + \sum_{j \in \mathcal{N}_i} f_3 \left(s_j^{(t)}, \|\vec{x}_{ij}\| \right) \odot \vec{x}_{ij} \quad (11)$$

Alternatively, $\mathfrak{G}$ -invariant layers [20–22] do not update vector features and only aggregate scalar quantities from local neighbourhoods:

$$s_i^{(t+1)} := \text{UPD} \left(s_i^{(t)}, \text{AGG} \left(\{ (s_j^{(t)}, \vec{s}_j^{(t)}, \vec{v}_j^{(t)}, \vec{x}_{ij}) \mid j \in \mathcal{N}_i \} \right) \right). \quad (12)$$

For *e.g.*, SchNet [20] uses relative distances to scalarise local geometric information, while DimeNet [22] uses both distances and angles, as follows:

 Addressed
R2.2

$$s_i^{(t+1)} := s_i^{(t)} + \sum_{j \in \mathcal{N}_i} f_1 \left(s_j^{(t)}, \|\vec{x}_{ij}\| \right) \quad (\text{SchNet}) \quad (13)$$

$$s_i^{(t+1)} := \sum_{j \in \mathcal{N}_i} f_1 \left(s_i^{(t)}, s_j^{(t)}, \sum_{k \in \mathcal{N}_i \setminus \{j\}} f_2 \left(s_j^{(t)}, s_k^{(t)}, \|\vec{x}_{ij}\|, \vec{x}_{ij} \cdot \vec{x}_{ik} \right) \right) \quad (\text{DimeNet}) \quad (14)$$

For both $\mathfrak{G}$ -invariant and $\mathfrak{G}$ -equivariant architectures, the scalar features $\{s_i^{(T)}\}$ at the final iteration T are mapped to graph-level features via a permutation-invariant readout $f: \mathbb{R}^{n \times f} \rightarrow \mathbb{R}^{f'}$.

Invariant GNNs have shown strong performance for protein design [31, 32] and electrocatalysis [33, 34], while equivariant GNNs are being used within learnt interatomic potentials for molecular dynamics [23, 35, 36].

B Discussion

Practical Implications. Proposition 10, together with Propositions 1 and 4, highlight critical theoretical limitations of $\mathfrak{G}$ -invariant GNNs. Our results suggest that $\mathfrak{G}$ -equivariant GNNs should be preferred when working with large geometric graphs such as macromolecules with thousands of nodes, where message passing is restricted to local radial neighbourhoods around each node.

Motivated by these limitations, two straightforward approaches to improving $\mathfrak{G}$ -invariant GNNs may be: (1) pre-computing non-local geometric properties as input features, *e.g.* models such as GemNet [33] and GearNet [31] successfully use two-hop dihedral angles. And (2) working with fully connected geometric graphs, as Proposition 7 suggests that $\mathfrak{G}$ -equivariant and $\mathfrak{G}$ -invariant GNNs can be made equally powerful when performing all-to-all message passing. This is supported by the empirical success of recent $\mathfrak{G}$ -invariant ‘Graph Transformers’ [34, 37] for small molecules with tens of nodes, where working with full graphs is tractable.

Related Work. Literature on the *completeness* of atom-centred interatomic potentials has focused on distinguishing 1-hop local neighbourhoods (point clouds) around atoms by building spanning sets for continuous, $\mathfrak{G}$ -equivariant multiset functions [27, 38–40]. Recent theoretical work on geometric GNNs and their universality has shown that Tensor Field Networks [41], GemNet [33], and GVP [17] can be universal approximators of continuous, $\mathfrak{G}$ -equivariant or $\mathfrak{G}$ -invariant multiset function over point clouds (not sparse graphs). In contrast, the GWL framework studies the expressive power of geometric GNNs from the perspective of geometric graph isomorphism. *Geometric graph matching has also been studied from the perspective of finding global isometries in the computer vision community*[42]. Our notion of geometric graph isomorphism is more general as it considers local message passing procedures as well as both scalar and geometric node attributes. Overall, our work formalises what classes of geometric graphs can and cannot be distinguished by message passing $\mathfrak{G}$ -invariant/equivariant GNNs while abstracting away implementation details.

 Addressed
R4.3

Future Work. GWL provides an abstraction to study the limits of geometric GNNs, but in practice it is challenging to build maximally powerful GNNs that satisfy the conditions of Proposition 9 as GWL relies on $\mathfrak{G}$ -orbit injective colouring and $\mathfrak{G}$ -equivariant propagation of auxiliary geometric information. Based on the intuitions gained from GWL, future work will explore building provably powerful, *practical* geometric GNNs for applications in biochemistry, material science, and multiagent robotics, and better characterise the trade-offs related to practical implementation choices.

C Proofs for What GWL and IGWL can Distinguish

The following results are a consequence of the construction of GWL as well as the definitions of k -hop distinct and k -hop identical geometric graphs. Note that k -hop distinct geometric graphs are also $(k + 1)$ -hop distinct. Similarly, k -hop identical geometric graphs are also $(k - 1)$ -hop identical, but not necessarily $(k + 1)$ -hop distinct.

Given two distinct neighbourhoods $\mathcal{N}_1$ and $\mathcal{N}_2$, the $\mathfrak{G}$ -orbits of the corresponding geometric multisets $\mathbf{g}_1$ and $\mathbf{g}_2$ are mutually exclusive, *i.e.* $\mathcal{O}_{\mathfrak{G}}(\mathbf{g}_1) \cap \mathcal{O}_{\mathfrak{G}}(\mathbf{g}_2) \equiv \emptyset$. By the properties of I-HASH this implies $c_1 \neq c_2$. Conversely, if $\mathcal{N}_1$ and $\mathcal{N}_2$ were identical up to group actions, their $\mathfrak{G}$ -orbits would overlap, *i.e.* $\mathbf{g}_1 = \mathbf{g} \mathbf{g}_2$ for some $\mathbf{g} \in \mathfrak{G}$ and $\mathcal{O}_{\mathfrak{G}}(\mathbf{g}_1) = \mathcal{O}_{\mathfrak{G}}(\mathbf{g}_2) \Rightarrow c_1 = c_2$.

Proposition 1. *GWL can distinguish any k -hop distinct geometric graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ where the underlying attributed graphs are isomorphic, and k iterations are sufficient.*

Proof of Proposition 1. The k -th iteration of GWL identifies the $\mathfrak{G}$ -orbit of the k -hop subgraph $\mathcal{N}_i^{(k)}$ at each node i via the geometric multiset $\mathbf{g}_i^{(k)}$. $\mathcal{G}_1$ and $\mathcal{G}_2$ being k -hop distinct implies that there exists some bijection b and some node $i \in \mathcal{V}_1, b(i) \in \mathcal{V}_2$ such that the corresponding k -hop subgraphs $\mathcal{N}_i^{(k)}$ and $\mathcal{N}_{b(i)}^{(k)}$ are distinct. Thus, the $\mathfrak{G}$ -orbits of the corresponding geometric multisets $\mathbf{g}_i^{(k)}$ and $\mathbf{g}_{b(i)}^{(k)}$ are mutually exclusive, *i.e.* $\mathcal{O}_{\mathfrak{G}}(\mathbf{g}_i^{(k)}) \cap \mathcal{O}_{\mathfrak{G}}(\mathbf{g}_{b(i)}^{(k)}) \equiv \emptyset \Rightarrow c_i^{(k)} \neq c_{b(i)}^{(k)}$. Thus, k iterations of GWL are sufficient to distinguish $\mathcal{G}_1$ and $\mathcal{G}_2$. $\square$

Proposition 2. *Up to k iterations, GWL cannot distinguish any k -hop identical geometric graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ where the underlying attributed graphs are isomorphic.*

Proof of Proposition 2. The k -th iteration of GWL identifies the $\mathfrak{G}$ -orbit of the k -hop subgraph $\mathcal{N}_i^{(k)}$ at each node i via the geometric multiset $\mathbf{g}_i^{(k)}$. $\mathcal{G}_1$ and $\mathcal{G}_2$ being k -hop identical implies that for all bijections b and all nodes $i \in \mathcal{V}_1, b(i) \in \mathcal{V}_2$, the corresponding k -hop subgraphs $\mathcal{N}_i^{(k)}$ and $\mathcal{N}_{b(i)}^{(k)}$ are identical up to group actions. Thus, the $\mathfrak{G}$ -orbits of the corresponding geometric multisets $\mathbf{g}_i^{(k)}$ and




Figure 2: Invariant GWL Test. IGWL cannot distinguish $\mathcal{G}_1$ and $\mathcal{G}_2$ as they are 1-hop identical: The $\mathfrak{G}$ -orbit of the 1-hop neighbourhood around each node is the same, and IGWL cannot propagate geometric orientation information beyond 1-hop (here $\mathfrak{G} = O(d)$).




Figure 3: Geometric Computation Trees for GWL and IGWL. Unlike GWL, geometric orientation information cannot flow from the leaves to the root in IGWL, restricting its expressive power. IGWL cannot distinguish $\mathcal{G}_1$ and $\mathcal{G}_2$ as all 1-hop neighbourhoods are computationally identical.

392 $\mathbf{g}_{b(i)}^{(k)}$ overlap, i.e. $\mathcal{O}_{\mathfrak{G}}(\mathbf{g}_i^{(k)}) = \mathcal{O}_{\mathfrak{G}}(\mathbf{g}_{b(i)}^{(k)}) \Rightarrow c_i^{(k)} = c_{b(i)}^{(k)}$. Thus, up to k iterations of GWL cannot
 393 distinguish $\mathcal{G}_1$ and $\mathcal{G}_2$. $\square$

394 **Proposition 3.** IGWL can distinguish any 1-hop distinct geometric graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ where the
 395 underlying attributed graphs are isomorphic, and 1 iteration is sufficient.

396 *Proof of Proposition 3.* Each iteration of IGWL identifies the $\mathfrak{G}$ -orbit of the 1-hop local neighbour-
 397 hood $\mathcal{N}_i^{(k=1)}$ at each node i . $\mathcal{G}_1$ and $\mathcal{G}_2$ being 1-hop distinct implies that there exists some bijection
 398 b and some node $i \in \mathcal{V}_1, b(i) \in \mathcal{V}_2$ such that the corresponding 1-hop local neighbourhoods $\mathcal{N}_i^{(1)}$
 399 and $\mathcal{N}_{b(i)}^{(1)}$ are distinct. Thus, the $\mathfrak{G}$ -orbits of the corresponding geometric multisets $\mathbf{g}_i^{(1)}$ and $\mathbf{g}_{b(i)}^{(1)}$ are
 400 mutually exclusive, i.e. $\mathcal{O}_{\mathfrak{G}}(\mathbf{g}_i^{(1)}) \cap \mathcal{O}_{\mathfrak{G}}(\mathbf{g}_{b(i)}^{(1)}) \equiv \emptyset \Rightarrow c_i^{(1)} \neq c_{b(i)}^{(1)}$. Thus, 1 iteration of IGWL is
 401 sufficient to distinguish $\mathcal{G}_1$ and $\mathcal{G}_2$. $\square$

402 **Proposition 4.** Any number of iterations of IGWL cannot distinguish any 1-hop identical geometric
 403 graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ where the underlying attributed graphs are isomorphic.

404 *Proof of Proposition 4.* Each iteration of IGWL identifies the $\mathfrak{G}$ -orbit of the 1-hop local neighbour-
 405 hood $\mathcal{N}_i^{(k=1)}$ at each node i , but cannot identify $\mathfrak{G}$ -orbits beyond 1-hop by the construction of IGWL
 406 as no geometric information is propagated. $\mathcal{G}_1$ and $\mathcal{G}_2$ being 1-hop identical implies that for all
 407 bijections b and all nodes $i \in \mathcal{V}_1, b(i) \in \mathcal{V}_2$, the corresponding 1-hop local neighbourhoods $\mathcal{N}_i^{(k)}$ and
 408 $\mathcal{N}_{b(i)}^{(k)}$ are identical up to group actions. Thus, the $\mathfrak{G}$ -orbits of the corresponding geometric multisets

409 $\mathbf{g}_i^{(1)}$ and $\mathbf{g}_{b(i)}^{(1)}$ overlap, i.e. $\mathcal{O}_{\mathfrak{G}}(\mathbf{g}_i^{(1)}) = \mathcal{O}_{\mathfrak{G}}(\mathbf{g}_{b(i)}^{(1)}) \Rightarrow c_i^{(k)} = c_{b(i)}^{(k)}$. Thus, any number of IGWL
 410 iterations cannot distinguish $\mathcal{G}_1$ and $\mathcal{G}_2$. $\square$

411 **Proposition 5.** *Assuming geometric graphs are constructed from point clouds using radial cutoffs,*
 412 *GWL can distinguish any geometric graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ where the underlying attributed graphs are*
 413 *non-isomorphic. At most k_{Max} iterations are sufficient, where k_{Max} is the maximum graph diameter*
 414 *among $\mathcal{G}_1$ and $\mathcal{G}_2$.*

415 *Proof of Proposition 5.* We assume that a geometric graph $\mathcal{G} = (\mathbf{A}, \mathbf{S}, \vec{\mathbf{V}}, \vec{\mathbf{X}})$ is constructed from a
 416 point cloud $(\mathbf{S}, \vec{\mathbf{V}}, \vec{\mathbf{X}})$ using a predetermined radial cutoff r . Thus, the adjacency matrix is defined
 417 as $a_{ij} = 1$ if $\|\vec{\mathbf{x}}_i - \vec{\mathbf{x}}_j\|_2 \leq r$, or 0 otherwise, for all $a_{ij} \in \mathbf{A}$. Such construction procedures are
 418 conventional for geometric graphs in biochemistry and material science.

419 Given geometric graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ where the underlying attributed graphs are non-isomorphic,
 420 identify k_{Max} the maximum of the graph diameters of $\mathcal{G}_1$ and $\mathcal{G}_2$, and chose any arbitrary nodes
 421 $i \in \mathcal{V}_1, j \in \mathcal{V}_2$. We can define the k_{Max} -hop subgraphs $\mathcal{N}_i^{(k_{\text{Max}})}$ and $\mathcal{N}_j^{(k_{\text{Max}})}$ at i and j , respectively.
 422 Thus, $\mathcal{N}_i^{(k_{\text{Max}})} = \mathcal{V}_1$ for all $i \in \mathcal{V}_1$, and $\mathcal{N}_j^{(k_{\text{Max}})} = \mathcal{V}_2$ for all $j \in \mathcal{V}_2$. Due to the assumed construction
 423 procedure of geometric graphs, $\mathcal{N}_i^{(k_{\text{Max}})}$ and $\mathcal{N}_j^{(k_{\text{Max}})}$ must be distinct. Otherwise, if $\mathcal{N}_i^{(k_{\text{Max}})}$ and
 424 $\mathcal{N}_j^{(k_{\text{Max}})}$ were identical up to group actions, the sets $(\mathbf{S}_1, \vec{\mathbf{V}}_1, \vec{\mathbf{X}}_1)$ and $(\mathbf{S}_2, \vec{\mathbf{V}}_2, \vec{\mathbf{X}}_2)$ would have
 425 yielded isomorphic graphs.

426 The k_{Max} -th iteration of GWL identifies the $\mathfrak{G}$ -orbit of the k_{Max} -hop subgraph $\mathcal{N}_i^{(k_{\text{Max}})}$ at each node
 427 i via the geometric multiset $\mathbf{g}_i^{(k_{\text{Max}})}$. As $\mathcal{N}_i^{(k_{\text{Max}})}$ and $\mathcal{N}_j^{(k_{\text{Max}})}$ are distinct for any arbitrary nodes
 428 $i \in \mathcal{V}_1, j \in \mathcal{V}_2$, the $\mathfrak{G}$ -orbits of the corresponding geometric multisets $\mathbf{g}_i^{(k_{\text{Max}})}$ and $\mathbf{g}_j^{(k_{\text{Max}})}$ are mutually
 429 exclusive, i.e. $\mathcal{O}_{\mathfrak{G}}(\mathbf{g}_i^{(k_{\text{Max}})}) \cap \mathcal{O}_{\mathfrak{G}}(\mathbf{g}_j^{(k_{\text{Max}})}) \equiv \emptyset \Rightarrow c_i^{(k_{\text{Max}})} \neq c_j^{(k_{\text{Max}})}$. Thus, k_{Max} iterations of GWL
 430 are sufficient to distinguish $\mathcal{G}_1$ and $\mathcal{G}_2$. $\square$

431 **Theorem 6.** *GWL is strictly more powerful than IGWL.*

432 *Proof of Theorem 6.* Firstly, we can show that the GWL class contains IGWL if GWL can learn the
 433 identity when updating $\mathbf{g}_i$ for all $i \in \mathcal{V}$, i.e. $\mathbf{g}_i^{(t)} = \mathbf{g}_i^{(t-1)} = \mathbf{g}_i^{(0)} \equiv (\mathbf{s}_i, \vec{\mathbf{v}}_i)$. Thus, GWL is at least
 434 as powerful as IGWL, which does not update $\mathbf{g}_i$.

435 Secondly, to show that GWL is strictly more powerful than IGWL, it suffices to show that there exist
 436 a pair of geometric graphs that can be distinguished by GWL but not by IGWL. We may consider any
 437 k -hop distinct geometric graphs for $k > 1$, where the underlying attributed graphs are isomorphic.
 438 Proposition 1 states that GWL can distinguish any such graphs, while Proposition 4 states that IGWL
 439 cannot distinguish them. An example is the pair of graphs in Figures 1 and 2. $\square$

440 **Proposition 7.** *IGWL has the same expressive power as GWL for fully connected geometric graphs.*

441 *Proof of Proposition 7.* We will prove by contradiction. Assume that there exist a pair of fully
 442 connected geometric graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ which GWL can distinguish, but IGWL cannot.

443 If the underlying attributed graphs of $\mathcal{G}_1$ and $\mathcal{G}_2$ are isomorphic, by Proposition 1 and Proposition
 444 4, $\mathcal{G}_1$ and $\mathcal{G}_2$ are 1-hop identical but k -hop distinct for some $k > 1$. For all bijections b and all
 445 nodes $i \in \mathcal{V}_1, b(i) \in \mathcal{V}_2$, the local neighbourhoods $\mathcal{N}_i^{(1)}$ and $\mathcal{N}_{b(i)}^{(1)}$ are identical up to group actions,
 446 and $\mathcal{O}_{\mathfrak{G}}(\mathbf{g}_i^{(1)}) = \mathcal{O}_{\mathfrak{G}}(\mathbf{g}_{b(i)}^{(1)}) \Rightarrow c_i^{(1)} = c_{b(i)}^{(1)}$. Additionally, there exists some bijection b and
 447 some nodes $i \in \mathcal{V}_1, b(i) \in \mathcal{V}_2$ such that the k -hop subgraphs $\mathcal{N}_i^{(k)}$ and $\mathcal{N}_{b(i)}^{(k)}$ are distinct, and
 448 $\mathcal{O}_{\mathfrak{G}}(\mathbf{g}_i^{(k)}) \cap \mathcal{O}_{\mathfrak{G}}(\mathbf{g}_{b(i)}^{(k)}) \equiv \emptyset \Rightarrow c_i^{(k)} \neq c_{b(i)}^{(k)}$. However, as $\mathcal{G}_1$ and $\mathcal{G}_2$ are fully connected, for any k ,
 449 $\mathcal{N}_i^{(1)} = \mathcal{N}_i^{(k)}$ and $\mathcal{N}_{b(i)}^{(1)} = \mathcal{N}_{b(i)}^{(k)}$ are identical up to group actions. Thus, $\mathcal{O}_{\mathfrak{G}}(\mathbf{g}_i^{(1)}) = \mathcal{O}_{\mathfrak{G}}(\mathbf{g}_i^{(k)}) =$
 450 $\mathcal{O}_{\mathfrak{G}}(\mathbf{g}_{b(i)}^{(1)}) = \mathcal{O}_{\mathfrak{G}}(\mathbf{g}_{b(i)}^{(k)}) \Rightarrow c_i^{(1)} = c_i^{(k)} = c_{b(i)}^{(k)} = c_{b(i)}^{(1)}$. This is a contradiction.

451 If $\mathcal{G}_1$ and $\mathcal{G}_2$ are non-isomorphic and fully connected, for any arbitrary $i \in \mathcal{V}_1, j \in \mathcal{V}_2$ and any k -hop
 452 neighbourhood, we know that $\mathcal{N}_i^{(1)} = \mathcal{N}_i^{(k)}$ and $\mathcal{N}_j^{(1)} = \mathcal{N}_j^{(k)}$. Thus, a single iteration of GWL and

IGWL identify the same $\mathfrak{G}$ -orbits and assign the same node colours, i.e. $\mathcal{O}_{\mathfrak{G}}(\mathbf{g}_i^{(1)}) = \mathcal{O}_{\mathfrak{G}}(\mathbf{g}_i^{(k)}) \Rightarrow c_i^{(1)} = c_i^{(k)}$ and $\mathcal{O}_{\mathfrak{G}}(\mathbf{g}_j^{(1)}) = \mathcal{O}_{\mathfrak{G}}(\mathbf{g}_j^{(k)}) \Rightarrow c_j^{(1)} = c_j^{(k)}$. This is a contradiction. $\square$

D Proofs for equivalence between GWL and Geometric GNNs

Our proofs adapt the techniques used in [2, 3] for connecting 1-WL with GNNs. Note that we omit including the relative position vectors $\vec{\mathbf{x}}_{ij} = \vec{\mathbf{x}}_i - \vec{\mathbf{x}}_j$ in GWL and geometric GNN updates for brevity, as relative positions vectors can be merged into the vector features.

Theorem 8. *Any pair of geometric graphs distinguishable by a $\mathfrak{G}$ -equivariant GNN is also distinguishable by GWL.*

Proof of Theorem 8. Consider two geometric graphs $\mathcal{G}$ and $\mathcal{H}$. The theorem implies that if the GNN graph-level readout outputs $f(\mathcal{G}) \neq f(\mathcal{H})$, then the GWL test will always determine $\mathcal{G}$ and $\mathcal{H}$ to be non-isomorphic, i.e. $\mathcal{G} \neq \mathcal{H}$.

We will prove by contradiction. Suppose after T iterations, a GNN graph-level readout outputs $f(\mathcal{G}) \neq f(\mathcal{H})$, but the GWL test cannot decide $\mathcal{G}$ and $\mathcal{H}$ are non-isomorphic, i.e. $\mathcal{G}$ and $\mathcal{H}$ always have the same collection of node colours for iterations 0 to T . Thus, for iteration t and $t + 1$ for any $t = 0 \dots T - 1$, $\mathcal{G}$ and $\mathcal{H}$ have the same collection of node colours $\{c_i^{(t)}\}$ as well as the same collection of neighbourhood geometric multisets $\{(c_i^{(t)}, \mathbf{g}_i^{(t)}), \{(c_j^{(t)}, \mathbf{g}_j^{(t)}) \mid j \in \mathcal{N}_i\}\}$ up to group actions. Otherwise, the GWL test would have produced different node colours at iteration $t + 1$ for $\mathcal{G}$ and $\mathcal{H}$ as different geometric multisets get unique new colours.

We will show that on the same graph for nodes i and k , if $(c_i^{(t)}, \mathbf{g}_i^{(t)}) = (c_k^{(t)}, \mathbf{g} \cdot \mathbf{g}_k^{(t)})$, we always have GNN features $(\mathbf{s}_i^{(t)}, \vec{\mathbf{v}}_i^{(t)}) = (\mathbf{s}_k^{(t)}, \mathbf{Q}_{\mathbf{g}} \vec{\mathbf{v}}_k^{(t)})$ for any iteration t . This holds for $t = 0$ because GWL and the GNN start with the same initialisation. Suppose this holds for iteration t . At iteration $t + 1$, if for any i and k , $(c_i^{(t+1)}, \mathbf{g}_i^{(t+1)}) = (c_k^{(t+1)}, \mathbf{g} \cdot \mathbf{g}_k^{(t+1)})$, then:

$$\{(c_i^{(t)}, \mathbf{g}_i^{(t)}), \{(c_j^{(t)}, \mathbf{g}_j^{(t)}) \mid j \in \mathcal{N}_i\}\} = \{(c_k^{(t)}, \mathbf{g} \cdot \mathbf{g}_k^{(t)}), \{(c_j^{(t)}, \mathbf{g} \cdot \mathbf{g}_j^{(t)}) \mid j \in \mathcal{N}_k\}\} \quad (15)$$

By our assumption on iteration t ,

$$\{(\mathbf{s}_i^{(t)}, \vec{\mathbf{v}}_i^{(t)}), \{(\mathbf{s}_j^{(t)}, \vec{\mathbf{v}}_j^{(t)}) \mid j \in \mathcal{N}_i\}\} = \{(\mathbf{s}_k^{(t)}, \mathbf{Q}_{\mathbf{g}} \vec{\mathbf{v}}_k^{(t)}), \{(\mathbf{s}_j^{(t)}, \mathbf{Q}_{\mathbf{g}} \vec{\mathbf{v}}_j^{(t)}) \mid j \in \mathcal{N}_k\}\} \quad (16)$$

As the same aggregate and update operations are applied at each node within the GNN, the same inputs, i.e. neighbourhood features, are mapped to the same output. Thus, $(\mathbf{s}_i^{(t+1)}, \vec{\mathbf{v}}_i^{(t+1)}) = (\mathbf{s}_k^{(t+1)}, \mathbf{Q}_{\mathbf{g}} \vec{\mathbf{v}}_k^{(t+1)})$. By induction, if $(c_i^{(t)}, \mathbf{g}_i^{(t)}) = (c_k^{(t)}, \mathbf{g} \cdot \mathbf{g}_k^{(t)})$, we always have GNN node features $(\mathbf{s}_i^{(t)}, \vec{\mathbf{v}}_i^{(t)}) = (\mathbf{s}_k^{(t)}, \mathbf{Q}_{\mathbf{g}} \vec{\mathbf{v}}_k^{(t)})$ for any iteration t . This creates valid mappings ϕ_s, ϕ_v such that $\mathbf{s}_i^{(t)} = \phi_s(c_i^{(t)})$ and $\vec{\mathbf{v}}_i^{(t)} = \phi_v(c_i^{(t)}, \mathbf{g}_i^{(t)})$ for any $i \in \mathcal{V}$.

Thus, if $\mathcal{G}$ and $\mathcal{H}$ have the same collection of node colours and geometric multisets, then $\mathcal{G}$ and $\mathcal{H}$ also have the same collection of GNN neighbourhood features

$$\{(\mathbf{s}_i^{(t)}, \vec{\mathbf{v}}_i^{(t)}), \{(\mathbf{s}_j^{(t)}, \vec{\mathbf{v}}_j^{(t)}) \mid j \in \mathcal{N}_i\}\} = \{(\phi_s(c_i^{(t)}), \phi_v(c_i^{(t)}, \mathbf{g}_i^{(t)})), \{(\phi_s(c_j^{(t)}), \phi_v(c_j^{(t)}, \mathbf{g}_j^{(t)})) \mid j \in \mathcal{N}_i\}\}$$

Thus, the GNN will output the same collection of node scalar features $\{\mathbf{s}_i^{(T)}\}$ for $\mathcal{G}$ and $\mathcal{H}$ and the permutation-invariant graph-level readout will output $f(\mathcal{G}) = f(\mathcal{H})$. This is a contradiction. $\square$

Proposition 9. *$\mathfrak{G}$ -equivariant GNNs have the same expressive power as GWL if the following conditions hold: (1) The aggregation AGG is an injective, $\mathfrak{G}$ -equivariant multiset function. (2) The scalar part of the update UPD_s is a $\mathfrak{G}$ -orbit injective, $\mathfrak{G}$ -invariant multiset function. (3) The vector part of the update UPD_v is an injective, $\mathfrak{G}$ -equivariant multiset function. (4) The graph-level readout f is an injective multiset function.*

Proof of Theorem 9. Consider a GNN where the conditions hold. We will show that, with a sufficient number of iterations t , the output of this GNN is equivalent to GWL, i.e. $\mathbf{s}^{(t)} \equiv c^{(t)}$.

Let $\mathcal{G}$ and $\mathcal{H}$ be any geometric graphs which the GWL test decides as non-isomorphic at iteration T . Because the graph-level readout function is injective, *i.e.* it maps distinct multiset of node scalar features into unique embeddings, it suffices to show that the GNN's neighbourhood aggregation process, with sufficient iterations, embeds $\mathcal{G}$ and $\mathcal{H}$ into different multisets of node features.

For this proof, we replace $\mathfrak{G}$ -orbit injective functions with injective functions over the equivalence class generated by the actions of $\mathfrak{G}$. Thus, all elements belonging to the same $\mathfrak{G}$ -orbit will first be mapped to the same representative of the equivalence class, denoted by the square brackets $[\dots]$, followed by an injective map. The result is $\mathfrak{G}$ -orbit injective.

 Addressed
R2.3

Let us assume the GNN updates node scalar and vector features as:

$$\mathbf{s}_i^{(t)} = \text{UPD}_s \left(\left[(\mathbf{s}_i^{(t-1)}, \vec{\mathbf{v}}_i^{(t-1)}) , \text{AGG} \left(\{ \{ (\mathbf{s}_j^{(t-1)}, \mathbf{s}_j^{(t-1)}, \vec{\mathbf{v}}_j^{(t-1)}, \vec{\mathbf{v}}_j^{(t-1)}) \mid j \in \mathcal{N}_i \} \} \right) \right] \right) \quad (17)$$

$$\vec{\mathbf{v}}_i^{(t)} = \text{UPD}_v \left((\mathbf{s}_i^{(t-1)}, \vec{\mathbf{v}}_i^{(t-1)}) , \text{AGG} \left(\{ \{ (\mathbf{s}_j^{(t-1)}, \mathbf{s}_j^{(t-1)}, \vec{\mathbf{v}}_j^{(t-1)}, \vec{\mathbf{v}}_j^{(t-1)}) \mid j \in \mathcal{N}_i \} \} \right) \right) \quad (18)$$

with the aggregation function AGG being $\mathfrak{G}$ -equivariant and injective, the scalar update function UPD_s being $\mathfrak{G}$ -invariant and injective, and the vector update function UPD_v being $\mathfrak{G}$ -equivariant and injective.

The GWL test updates the node colour $c_i^{(t)}$ and geometric multiset $\mathbf{g}_i^{(t)}$ as:

$$c_i^{(t)} = h_s \left(\left[(c_i^{(t-1)}, \mathbf{g}_i^{(t-1)}) , \{ \{ (c_j^{(t-1)}, \mathbf{g}_j^{(t-1)}) \mid j \in \mathcal{N}_i \} \} \right] \right), \quad (19)$$

$$\mathbf{g}_i^{(t)} = h_v \left((c_i^{(t-1)}, \mathbf{g}_i^{(t-1)}) , \{ \{ (c_j^{(t-1)}, \mathbf{g}_j^{(t-1)}) \mid j \in \mathcal{N}_i \} \} \right), \quad (20)$$

where h_s is a $\mathfrak{G}$ -invariant and injective map, and h_v is a $\mathfrak{G}$ -equivariant and injective operation (e.g. in equation 2, expanding the geometric multiset by copying).

 Addressed
R2.3

We will show by induction that at any iteration t , there always exist injective functions φ_s and φ_v such that $\mathbf{s}_i^{(t)} = \varphi_s(c_i^{(t)})$ and $\vec{\mathbf{v}}_i^{(t)} = \varphi_v(c_i^{(t)}, \mathbf{g}_i^{(t)})$. This holds for $t = 0$ because the initial node features are the same for GWL and GNN, $c_i^{(0)} \equiv \mathbf{s}_i^{(0)}$ and $\mathbf{g}_i^{(0)} \equiv (\mathbf{s}_i^{(0)}, \vec{\mathbf{v}}_i^{(0)})$ for all $i \in \mathcal{V}(\mathcal{G}), \mathcal{V}(\mathcal{H})$. Suppose this holds for iteration t . At iteration $t + 1$, substituting $\mathbf{s}_i^{(t)}$ with $\varphi_s(c_i^{(t)})$, and $\vec{\mathbf{v}}_i^{(t)}$ with $\varphi_v(c_i^{(t)}, \mathbf{g}_i^{(t)})$ gives us

$$\mathbf{s}_i^{(t+1)} = \text{UPD}_s \left(\left[(\varphi_s(c_i^{(t)}), \varphi_v(c_i^{(t)}, \mathbf{g}_i^{(t)})) , \text{AGG} \left(\{ \{ (\varphi_s(c_j^{(t)}), \varphi_s(c_j^{(t)}), \varphi_v(c_j^{(t)}, \mathbf{g}_j^{(t)}), \varphi_v(c_j^{(t)}, \mathbf{g}_j^{(t)})) \mid j \in \mathcal{N}_i \} \} \right) \right] \right)$$

$$\vec{\mathbf{v}}_i^{(t+1)} = \text{UPD}_v \left((\varphi_s(c_i^{(t)}), \varphi_v(c_i^{(t)}, \mathbf{g}_i^{(t)})) , \text{AGG} \left(\{ \{ (\varphi_s(c_j^{(t)}), \varphi_s(c_j^{(t)}), \varphi_v(c_j^{(t)}, \mathbf{g}_j^{(t)}), \varphi_v(c_j^{(t)}, \mathbf{g}_j^{(t)})) \mid j \in \mathcal{N}_i \} \} \right) \right)$$

The composition of multiple injective functions is injective. Therefore, there exist some injective functions g_s and g_v such that:

$$\mathbf{s}_i^{(t+1)} = g_s \left(\left[(c_i^{(t)}, \mathbf{g}_i^{(t)}) , \{ \{ (c_j^{(t)}, \mathbf{g}_j^{(t)}) \mid j \in \mathcal{N}_i \} \} \right] \right), \quad (21)$$

$$\vec{\mathbf{v}}_i^{(t+1)} = g_v \left((c_i^{(t)}, \mathbf{g}_i^{(t)}) , \{ \{ (c_j^{(t)}, \mathbf{g}_j^{(t)}) \mid j \in \mathcal{N}_i \} \} \right), \quad (22)$$

We can then consider:

$$\mathbf{s}_i^{(t+1)} = g_s \circ h_s^{-1} h_s \left(\left[(c_i^{(t)}, \mathbf{g}_i^{(t)}) , \{ \{ (c_j^{(t)}, \mathbf{g}_j^{(t)}) \mid j \in \mathcal{N}_i \} \} \right] \right), \quad (23)$$

$$\vec{\mathbf{v}}_i^{(t+1)} = g_v \circ h_v^{-1} h_v \left((c_i^{(t)}, \mathbf{g}_i^{(t)}) , \{ \{ (c_j^{(t)}, \mathbf{g}_j^{(t)}) \mid j \in \mathcal{N}_i \} \} \right), \quad (24)$$

Then, we can denote $\varphi_s = g_s \circ h_s^{-1}$ and $\varphi_v = g_v \circ h_v^{-1}$ as injective functions because the composition of injective functions is injective. Hence, for any iteration $t + 1$, there exist injective functions φ_s and φ_v such that $\mathbf{s}_i^{(t+1)} = \varphi_s(c_i^{(t+1)})$ and $\vec{\mathbf{v}}_i^{(t+1)} = \varphi_v(c_i^{(t+1)}, \mathbf{g}_i^{(t+1)})$.

At the T -th iteration, the GWL test decides that $\mathcal{G}$ and $\mathcal{H}$ are non-isomorphic, which means the multisets of node colours $\{c_i^{(T)}\}$ are different for $\mathcal{G}$ and $\mathcal{H}$. The GNN's node scalar features $\{\mathbf{s}_i^{(T)}\} = \{\varphi_s(c_i^{(T)})\}$ must also be different for $\mathcal{G}$ and $\mathcal{H}$ because of the injectivity of φ_s .

521

□

E Understanding the Design Space of Geometric GNNs via GWL

We can use the GWL framework to better understand key design choices for building geometric GNNs [25]: (1) Depth or number of layers; and (2) Body order of invariant scalars. In doing so, we formalise theoretical limitations of current geometric GNNs and provide practical implications.

Addressed
R2.6, R4.4

E.1 Role of Depth: Propagating Geometric Information

Each iteration of GWL expands the neighbourhood from which geometric information can be gathered. Thus, stacking multiple $\mathfrak{G}$ -equivariant GNN layers enables the computation of compositional geometric features.

This can be understood via a geometric version of computation trees [43]. A computation tree $\mathcal{T}_i^{(t)}$ represents the maximum information contained in GWL/IGWL colours or GNN features at iteration t by an ‘unrolling’ of the message passing procedure. Geometric computation trees are constructed recursively: $\mathcal{T}_i^{(0)} = (s_i, \vec{v}_i)$ for all $i \in \mathcal{V}$. For $t > 0$, we start with a root node $(s_i, \vec{v}_i)$ and add a child subtree $\mathcal{T}_j^{(t-1)}$ for all $j \in \mathcal{N}_i$ along with the relative position $\vec{x}_{ij}$ along the edge, as shown in Figure 3. To obtain the root node’s embedding or colour, both scalar and geometric information is propagated from the leaves up to the root. Thus, if two nodes have identical geometric computation trees, they will be mapped to the same node embedding or colour.

Critically, geometric orientation information cannot flow from one level to another in the computation trees for IGWL and $\mathfrak{G}$ -invariant GNNs, as they only update scalar information. In the recursive construction procedure, we must insert a connector node $(s_j, \vec{v}_j)$ before adding the child subtree $\mathcal{T}_j^{(t-1)}$ for all $j \in \mathcal{N}_i$ and prevent geometric information propagation between them. Consequently, even the most powerful $\mathfrak{G}$ -invariant GNNs are restricted in their ability to compute global and non-local geometric properties.

Proposition 10. *IGWL and $\mathfrak{G}$ -invariant GNNs cannot decide several geometric graph properties: (1) perimeter, surface area, and volume of the bounding box/sphere enclosing the geometric graph; (2) distance from the centroid or centre of mass; and (3) dihedral angles.*

Practical Implications. Proposition 10, together with Propositions 1 and 4, highlight critical theoretical limitations of $\mathfrak{G}$ -invariant GNNs. Our results suggest that $\mathfrak{G}$ -equivariant GNNs should be preferred when working with large geometric graphs such as macromolecules with thousands of nodes, where message passing is restricted to local radial neighbourhoods around each node.


Motivated by these limitations, two straightforward approaches to improving $\mathfrak{G}$ -invariant GNNs may be: (1) pre-computing non-local geometric properties as input features, *e.g.* models such as GemNet [33] and GearNet [31] successfully use two-hop dihedral angles. And (2) working with fully connected geometric graphs, as Proposition 7 suggests that $\mathfrak{G}$ -equivariant and $\mathfrak{G}$ -invariant GNNs can be made equally powerful when performing all-to-all message passing. This is supported by the empirical success of recent $\mathfrak{G}$ -invariant ‘Graph Transformers’ [34, 37] for small molecules with tens of nodes, where working with full graphs is tractable.

Synthetic Experiment: k -chains. We present a simple synthetic experiment to demonstrate the role of depth in propagating geometric information beyond local neighbourhoods. We consider k -chain geometric graphs, which generalise the examples from [23] and the pair of graphs presented in Figure 3. Each pair of k -chains consists of $k + 2$ nodes with k nodes arranged in a line and differentiated by the orientation of the 2 end points. Thus, k -chain graphs are $(k - 1)$ -hop distinguishable, and $(k - 1)$ iterations of GWL are sufficient to distinguish them. In Table 1, we find that $(k - 1)$ layers of $\mathfrak{G}$ -equivariant message passing can learn to perfectly distinguish any k -chains, while $\mathfrak{G}$ -invariant GNNs are unable to distinguish k -chains for $k \geq 2$.

Addressed
R1.1, R3.3

E.2 Role of Body Order: Distinguishing $\mathfrak{G}$ -Orbits

At each iteration of GWL and IGWL, the I-HASH function assigns a $\mathfrak{G}$ -invariant colouring to distinct geometric neighbourhood patterns. For geometric GNNs, this corresponds to scalarising local geometric information when updating the scalar features, as shown in equation 13 and equation 14. Let us analyse the construction of the I-HASH function via the k -body variations $\text{IGWL}_{(k)}$. In doing so, we will make connections between IGWL and 1-WL for non-geometric graphs.



$(k = 6\text{-chains})$		Number of layers						
GNN Layer		1	2	3	4	5	6	7
Rand. guess		50%						
Non-geom.		50%	50%	50%	50%	50%	50%	50%
$O(3)$ -Inv.		50%	50%	50%	50%	50%	50%	50%
$O(3)$ -Equiv.		50%	50%	50%	50%	100%	100%	100%
$SO(3)$ -Equiv.		100%	100%	100%	100%	100%	100%	100%

Table 1: k -chain geometric graphs. k -chain graphs are $(k - 1)$ -hop distinguishable, and $(k - 1)$ iterations of $\mathfrak{G}$ -equivariant message passing are sufficient to distinguish them, while $\mathfrak{G}$ -invariant message passing is unable to distinguish k -chains for $k \geq 2$ (for $\mathfrak{G} = O(3)$).




Figure 4: Two geometric graphs for which IGWL and $\mathfrak{G}$ -invariant GNNs cannot distinguish their geometric properties such as perimeter, surface area, and volume of the bounding box/sphere, distance from the centroid, and dihedral angles. The centroid is denoted by a red point and distances from it are denoted by dotted red lines. The bounding box enclosing the geometric graph is denoted by the dotted green lines.

Firstly, we formalise the relationship between the injectivity of $\text{I-HASH}_{(k)}$ and the maximum cardinality of local neighbourhoods in a given dataset.

Proposition 11. $\text{I-HASH}_{(m)}$ is $\mathfrak{G}$ -orbit injective for $m = \max(\{|\mathcal{N}_i| \mid i \in \mathcal{V}\})$, the maximum cardinality of all local neighbourhoods $\mathcal{N}_i$ in a given dataset.

While building provably injective $\text{I-HASH}_{(k)}$ functions may require intractably high k , the hierarchy of $\text{IGWL}_{(k)}$ tests enable us to study the expressive power of practical $\mathfrak{G}$ -invariant aggregators used in current geometric GNN layers, e.g. SchNet [20], E-GNN [18], and Tensor Field Networks [15] use distances; DimeNet [22] uses distances and angles; and MACE [36] constructs scalars up to arbitrary k via Atomic Cluster Expansion [39]. We can state the following about the $\text{IGWL}_{(k)}$ hierarchy and the corresponding GNNs.

Proposition 12. $\text{IGWL}_{(k)}$ is at least as powerful as $\text{IGWL}_{(k-1)}$. For $k \leq 5$, $\text{IGWL}_{(k)}$ is strictly more powerful than $\text{IGWL}_{(k-1)}$.

Finally, we show that $\text{IGWL}_{(2)}$ is equivalent to 1-WL when all the pairwise distances between the nodes are the same. A similar observation was recently made by [44].

Proposition 13. Let $\mathcal{G}_1 = (\mathbf{A}_1, \mathbf{S}_1, \vec{\mathbf{X}}_1)$ and $\mathcal{G}_2 = (\mathbf{A}_2, \mathbf{S}_2, \vec{\mathbf{X}}_2)$ be two geometric graphs with the property that all edges have equal length. Then, $\text{IGWL}_{(2)}$ distinguishes the two graphs if and only if 1-WL can distinguish the attributed graphs $(\mathbf{A}_1, \mathbf{S}_1)$ and $(\mathbf{A}_2, \mathbf{S}_2)$.

This equivalence also points out the limitations of $\mathfrak{G}$ -invariant models like SchNet [20] and CGCNN [21] which only rely on pairwise distances and consequently suffer from all well-known failure cases of 1-WL. For instance, such models cannot distinguish two equilateral triangles from the regular hexagon [22].

E.3 Geometric GNN Design Space Proofs

Proposition 10. *IGWL and $\mathfrak{G}$ -invariant GNNs cannot decide several geometric graph properties: (1) perimeter, surface area, and volume of the bounding box/sphere enclosing the geometric graph; (2) distance from the centroid or centre of mass; and (3) dihedral angles.*

Proof of Proposition 10. Following [43], we say that a class of models *decides* a geometric graph property if there exists a model belonging to this class such that for any two geometric graphs that differ in the property, the model is able to distinguish the two geometric graphs.

In Figure 4 we provide an example of two geometric graphs that demonstrate the proposition. $\mathcal{G}_1$ and $\mathcal{G}_2$ differ in the following geometric graph properties:

- Perimeter, surface area, and volume of the bounding box enclosing the geometric graph¹: (32 units, 40 units², 16 units³) vs. (28 units, 24 units², 8 units³).
- Multiset of distances from the centroid or centre of mass: $\{0.00, 1.00, 1.00, 2.45, 2.45\}$ vs. $\{0.40, 1.08, 1.08, 2.32, 2.32\}$.
- Dihedral angles: $\angle(ljkm) = \frac{(\vec{x}_{jk} \times \vec{x}_{lj}) \cdot (\vec{x}_{jk} \times \vec{x}_{mk})}{|\vec{x}_{jk} \times \vec{x}_{lj}| |\vec{x}_{jk} \times \vec{x}_{mk}|}$ are clearly different for the two graphs.

However, according to Proposition 4 and Theorem ??, both IGWL and $\mathfrak{G}$ -invariant GNNs cannot distinguish these two geometric graphs, and therefore, cannot decide all these properties.

We can also show this by constructing geometric computation trees for any number of IGWL or $\mathfrak{G}$ -invariant GNN iterations, as illustrated in Figure 3. We observe that the geometric computation trees of any pair of isomorphic nodes are identical, as all 1-hop neighbourhoods are computationally identical. Therefore, the set of node colours or node scalar features will also be identical, which implies that $\mathcal{G}_1$ and $\mathcal{G}_2$ cannot be distinguished. $\square$

Proposition 11. *I-HASH_(m) is $\mathfrak{G}$ -orbit injective for $m = \max(\{|\mathcal{N}_i| \mid i \in \mathcal{V}\})$, the maximum cardinality of all local neighbourhoods $\mathcal{N}_i$ in a given dataset.*

Proof of Proposition 11. As m is the maximum cardinality of all local neighbourhoods $\mathcal{N}_i$ under consideration, any distinct neighbourhoods $\mathcal{N}_1$ and $\mathcal{N}_2$ must have distinct multisets of m -body scalars. As I-HASH_(m) computes scalars involving up to m nodes, it will be able to distinguish any such $\mathcal{N}_1$ and $\mathcal{N}_2$. Thus, I-HASH_(m) is $\mathfrak{G}$ -orbit injective. $\square$

Proposition 12. *IGWL_(k) is at least as powerful as IGWL_(k-1). For $k \leq 5$, IGWL_(k) is strictly more powerful than IGWL_(k-1).*

Proof of Proposition 12. By construction, I-HASH_(k) computes $\mathfrak{G}$ -invariant scalars from all possible tuples of up to k nodes formed by the elements of a neighbourhood and the central node. Thus, the I-HASH_(k) class contains I-HASH_(k-1), and I-HASH_(k) is at least as powerful as I-HASH_(k-1). Thus, the corresponding test IGWL_(k) is at least as powerful as IGWL_(k-1).

Secondly, to show that IGWL_(k) is strictly more powerful than IGWL_(k-1) for $k \leq 5$, it suffices to show that there exist a pair of geometric neighbourhoods that can be distinguished by IGWL_(k) but not by IGWL_(k-1):

- For $k = 3$ and $\mathfrak{G} = O(3)$ or $SO(3)$, for the local neighbourhood from Figure 1 in [23], two configurations with different angles between the neighbouring nodes can be distinguished by IGWL₍₃₎ but not by IGWL₍₂₎.
- For $k = 4$ and $\mathfrak{G} = O(3)$ or $SO(3)$, the pair of local neighbourhoods from Figure 1 in [27] can be distinguished by IGWL₍₄₎ but not by IGWL₍₃₎.
- For $k = 5$ and $\mathfrak{G} = O(3)$, the pair of local neighbourhoods from Figure 2(e) in [27] can be distinguished by IGWL₍₅₎ but not by IGWL₍₄₎.

¹The same result applies for the bounding sphere, not shown in the figure.

- For $k = 5$ and $\mathfrak{G} = SO(3)$, the pair of local neighbourhoods from Figure 2(f) in [27] can be distinguished by $IGWL_{(5)}$ but not by $IGWL_{(4)}$.

638

□

639 **Proposition 13.** *Let $\mathcal{G}_1 = (\mathbf{A}_1, \mathbf{S}_1, \vec{\mathbf{X}}_1)$ and $\mathcal{G}_2 = (\mathbf{A}_2, \mathbf{S}_2, \vec{\mathbf{X}}_2)$ be two geometric graphs with the*
 640 *property that all edges have equal length. Then, $IGWL_{(2)}$ distinguishes the two graphs if and only if*
 641 *1-WL can distinguish the attributed graphs $(\mathbf{A}_1, \mathbf{S}_1)$ and $(\mathbf{A}_2, \mathbf{S}_2)$.*

642 *Proof of Proposition 13.* Let c and k the colours produced by $IGWL_{(2)}$ and 1-WL, respectively, and
 643 let i and j be two nodes belonging to any two graphs like in the statement of the result. We prove the
 644 statement inductively.

645 Clearly, $c_i^{(0)} = k_i^{(0)}$ for all nodes i and $c_j^{(0)} = k_j^{(0)}$ if and only if $k_i^{(0)} = k_j^{(0)}$. Now, assume that the
 646 statement holds for iteration t . That is $c_i^{(t)} = c_j^{(t)}$ if and only if $k_i^{(t)} = k_j^{(t)}$ holds for all i . Note that
 647 $c_i^{(t+1)} = c_j^{(t+1)}$ if and only if $c_i^{(t)} = c_j^{(t)}$ and $\{(c_p^{(t)}, \|\vec{\mathbf{x}}_{ip}\|) \mid p \in \mathcal{N}_i\} = \{(c_p^{(t)}, \|\vec{\mathbf{x}}_{jp}\|) \mid p \in \mathcal{N}_j\}$,
 648 since the norm of the relative vectors is the only injective invariant that $IGWL_{(2)}$ can compute (up to
 649 a scaling). Since all the norms are equal, by the induction hypothesis, this is equivalent to $k_i^{(t)} = k_j^{(t)}$
 650 and $\{k_p^{(t)} \mid p \in \mathcal{N}_i\} = \{k_p^{(t)} \mid p \in \mathcal{N}_j\}$. Therefore, this is equivalent to $k_i^{(t+1)} = k_j^{(t+1)}$ □