
Complete Neural Networks for Complete Euclidean Graphs

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Neural networks for point clouds, which respect their natural invariance to per-
2 mutation and rigid motion, have enjoyed recent success in modeling geometric
3 phenomena, from molecular dynamics Reiser et al. [2022] to recommender systems
4 Yi et al. [2023]. Yet, to date, no architecture with polynomial complexity is known
5 to be *complete*, that is, able to distinguish between any pair of non-isomorphic
6 point clouds. We fill this theoretical gap by showing that point clouds can be
7 completely determined, up to permutation and rigid motion, by applying the 3-WL
8 graph isomorphism test to the point cloud’s centralized Gram matrix. Moreover, we
9 formulate a Euclidean variant of the 2-WL test and show that it is also sufficient to
10 achieve completeness. We then show how our complete Euclidean WL tests can be
11 simulated by a Euclidean graph neural network of moderate size and demonstrate
12 their separation capability on highly-symmetrical point clouds.

13 1 Introduction

14 A point cloud is a collection of n points in \mathbb{R}^d , where typically in applications $d = 3$. Machine
15 learning on point clouds is an important task with applications in chemistry Gilmer et al. [2017],
16 Wang et al. [2022], physical systems Finzi et al. [2021] and image processing Ma et al. [2023]. Many
17 successful architectures for point clouds are invariant by construction to the natural symmetries of
18 point clouds: permutations and rigid motions.

19 The rapidly increasing literature on point-cloud networks with permutation and rigid-motion sym-
20 metries has motivated research aimed at theoretically understanding the expressive power of the
21 various architectures. This analysis typically focuses on two closely related concepts: *Separation*
22 and *Universality*. We say an invariant architecture is *separating*, or *complete*, if it can assign distinct
23 values to any pair of point clouds that are not related by symmetry. An invariant architecture is
24 *universal* if it can approximate all continuous invariant functions on compact sets. Generally speaking,
25 these two concepts are essentially equivalent, as discussed in Villar et al. [2021], Joshi et al. [2022],
26 Chen et al. [2019], and in our context, in Appendix A.

27 Dym and Maron [2020] proved that the well-known Tensor Field Network Thomas et al. [2018]
28 invariant architecture is universal, but the construction in their proof requires arbitrarily high-order
29 representations of the rotation group. Similarly, universality can be obtained using high-order
30 representations of the permutation group Lim et al. [2022]. However, prior to this work, it was not
31 known whether the same theoretical guarantees can be achieved by realistic point-cloud architectures
32 that use low-dimensional representations, and whose complexity has a mild polynomial dependency
33 on the data dimension. In the words of Pozdnyakov and Ceriotti [2022]: "...provably universal
34 equivariant frameworks are such in the limit in which they generate high-order correlations... It is
35 an interesting, and open, question whether a given order suffices to guarantee complete resolving
36 power." (p. 6). We note that it is known that separation of point clouds in polynomial time in n

37 is possible, assuming that d is fixed (e.g., $d = 3$) Arvind and Rattan [2014], Dym and Kovalsky
38 [2019], Kurlin [2022]. What still remains to be established is whether separation is achievable for
39 common invariant machine learning models, and more generally, whether separation can be achieved
40 by computing a continuous invariant feature that is piecewise differentiable.

41 In this paper, we give what seems to be the first positive answer to this question. We focus on analyzing
42 a popular method for the construction of invariant point-cloud networks via *Graph Neural Networks*
43 (*GNNs*). This is done in two steps: first, point clouds are represented as a *Euclidean graph*- which we
44 define to be a complete weighted graph whose edge features are simple, rotation-invariant features:
45 the inner products between pairs of (centralized) points. We then apply permutation-invariant *Graph*
46 *Neural Networks (GNNs)* to the Euclidean graphs to obtain a rotation- and permutation-invariant
47 global point-cloud feature. This leads to a rich family of invariant point-cloud architectures, which is
48 determined by the type of GNN chosen.

49 The most straightforward implementation of this idea would be to apply the popular message passing
50 GNNs to the Euclidean graphs. One could also consider applying more expressive GNNs. For
51 combinatorial graphs, it is known that message-passing GNNs are only as expressive as the 1-WL
52 graph isomorphism test. There exists a hierarchy of k -WL graph isomorphism tests, where larger
53 values of k correspond to more expressive, and more expensive, graph isomorphism tests. There
54 are also corresponding GNNs that simulate the k -WL tests and have an equivalent separation power
55 Morris et al. [2018], Maron et al. [2019]. One could then consider applying these more expressive
56 architectures to Euclidean graphs, as suggested in Lim et al. [2022]. Accordingly, we aim to answer
57 the following questions:

58 **Question 1** For which k is the k -WL test, when applied to Euclidean graphs, complete?

59 **Question 2** Can this test be implemented in polynomial time by a continuous, piecewise-differentiable
60 architecture?

61 We begin by addressing Question 1. First, we consider a variation of the WL-test adapted for point
62 clouds, which we refer to as 1-*EWL* ('E' for Euclidean). This test was first proposed by Pozdnyakov
63 and Ceriotti [2022], where it was shown that it cannot distinguish between all 3-dimensional point
64 clouds, and consequently, neither can GNNs like Victor Garcia Satorras [2021], Schütt et al. [2017],
65 which can be shown to simulate it. Our first result, described in Section 2.1, balances this by showing
66 that two iterations of 1-*EWL* are enough to separate *almost any* pair of point clouds.

67 To achieve complete separation for *all* point clouds, we consider higher-order k -*EWL* tests. We first
68 consider a natural adaptation of k -WL for Euclidean graphs, which we name the *Vanilla- k -EWL* test.
69 In this test, the standard k -WL is applied to the Euclidean graph induced by the point clouds. We
70 show that when $k = 3$, this test is complete for 3-dimensional point clouds. Additionally, we propose
71 a variant of the Vanilla 2-*EWL*, which incorporates additional geometric information while having
72 the same complexity. We call this test the 2-*EWL* test, and show that it is complete on 3D point
73 clouds. We also propose a natural variation of 2-*EWL* called 2-*SEWL*, which can distinguish between
74 point clouds that are related by a reflection. This ability is important for chemical applications, as
75 most biological molecules that are related by a reflection are *not* chemically identical Kapon et al.
76 [2021] (this molecular property is called *chirality*).

77 We next address the second question of how to construct a GNN for Euclidean data with the same
78 separation power as that of the various k -*EWL* tests we describe. For combinatorial graphs, such
79 equivalence results rely on injective functions defined on multisets of discrete features Xu et al. [2018].
80 For Euclidean graphs, one can similarly rely on injective functions for multisets with continuous
81 features, such as those proposed in Dym and Gortler [2023]. However, a naive application of this
82 approach leads to a very large number of hidden features, which grows exponentially with the number
83 of message-passing iterations (see Figure 2). We show how this problem can be remedied, so that the
84 number of features needed depends only linearly on the number of message-passing iterations.

85 To summarize, our main results in this paper are:

- 86 1. We show that two iterations of 1-*EWL* can separate *almost all* point clouds in any dimension.
- 87 2. We prove the completeness of a single iteration of the vanilla 3-*EWL* for point clouds in \mathbb{R}^3 .
- 88 3. We formulate the 2-*SEWL* and 2-*EWL* tests, and prove their completeness for point clouds
89 in \mathbb{R}^3 .

90 4. We explain how to build differentiable architectures for point clouds with the same separation
91 power as Euclidean k -WL tests, with reasonable complexity.

92 **Experiments** In Section 5 we present synthetic experiments that demonstrate that 2-SEWL can
93 separate challenging point-cloud pairs that cannot be separated by several popular architectures.

94 **Disambiguation: Euclidean Graphs** In this paper we use a simple definition of a Euclidean graph
95 as the centralized Gram matrix of a point cloud, and focus on a fundamental theoretical question
96 related to this representation. In the learning literature, terms like ‘geometric graphs’ (not used here)
97 could refer to graphs that have both geometric and non-geometric edge and vertex features, or graphs
98 where pairwise distances are only available for specific point pairs (edges in an incomplete graph).

99 1.1 Related Work

100 **Euclidean WL** Pozdnyakov and Ceriotti [2022] showed that 1-EWL is incomplete for 3-
101 dimensional point clouds. Joshi et al. [2022] defines separation for a more general definition
102 of geometric graph, which combines geometric and combinatorial features. This work holds various
103 interesting insights for this more general problem but they do not prove completeness as we do here.

104 **Other complete constructions** As mentioned earlier, Dym and Maron [2020] proved universality
105 with respect to permutations and rigid motions for architectures using high-dimensional represen-
106 tations of the rotation group. Similar results were obtained in Finkelshtein et al. [2022], Gasteiger
107 et al. [2021]. In Lim et al. [2022] universality was proven for Euclidean GNNs with very high-order
108 permutation representations. In the planar case $d = 2$, universality using low-dimensional features
109 was achieved in Bökman et al. [2022]. For $d \geq 3$ our construction seems to be the first to achieve
110 universality using low dimensional representations.

111 For general fixed d , there do exist algorithms that can separate point clouds up to equivalence
112 in polynomial time, but they do not seem to lend themselves directly to neural architectures. In
113 Kurlin [2022], Widdowson and Kurlin [2023] complete tests are described, but they represent each
114 point cloud as a ‘multiset of multisets’ rather than as a vector as we do, and so are not suitable for
115 gradient descent based learning. Efficient tests for equivalence of Euclidean graphs were described in
116 Brass and Knauer [2000], Arvind and Rattan [2014], but they compute features that do not depend
117 continuously on the point cloud.

118 **Weaker notions of universality** In Widdowson and Kurlin [2022] the authors suggest a method
119 for distinguishing almost every point clouds up to equivalence, similar to our result here on 1-EWL.
120 Similarly, efficient separation/universality can also be obtained for point clouds with distinct principal
121 axes Puny et al. [2021], Kurlin [2022]. Another setting in which universality is easier to obtain is
122 when only rigid symmetries are considered and permutation symmetries are ignored Wang et al.
123 [2022], Villar et al. [2021], Victor Garcia Satorras [2021]. All these results do not provide universality
124 for *all* point clouds, with respect to the joint action of permutations and rigid motions.

125 Mathematical notation

126 A (finite) *multiset* $\{\{y_1, \dots, y_N\}\}$ is an unordered collection of elements where repetitions are allowed.

127 Let \mathcal{G} be a group acting on a set \mathcal{X} . For $X, Y \in \mathcal{X}$, we say that $X \stackrel{\mathcal{G}}{=} Y$ if $Y = gX$ for some $g \in \mathcal{G}$.

128 We say that a function $f : \mathcal{X} \rightarrow \mathcal{Y}$ is *invariant* if $f(gx) = f(x)$ for all $x \in \mathcal{X}, g \in G$. We say that f
129 is *equivariant* if \mathcal{Y} is also endowed with some action of G and $f(gx) = gf(x)$ for all $x \in \mathcal{X}, g \in \mathcal{G}$.

130 A separating invariant mapping is an invariant mapping that is injective, up to group equivalence:

131 **Definition 1.1** (Separating Invariant). Let \mathcal{G} be a group acting on a set \mathcal{X} . We say $F : \mathcal{X} \rightarrow \mathbb{R}^K$ is a \mathcal{G} -
132 *separating invariant* with *embedding dimension* K if for all $X, Y \in \mathcal{X}$, $F(X) = F(Y) \Leftrightarrow X \stackrel{\mathcal{G}}{=} Y$.

133 We focus on the case where \mathcal{X} is some Euclidean domain. To enable gradient-based learning, we
134 shall need separating mappings that are continuous everywhere and differentiable almost everywhere.

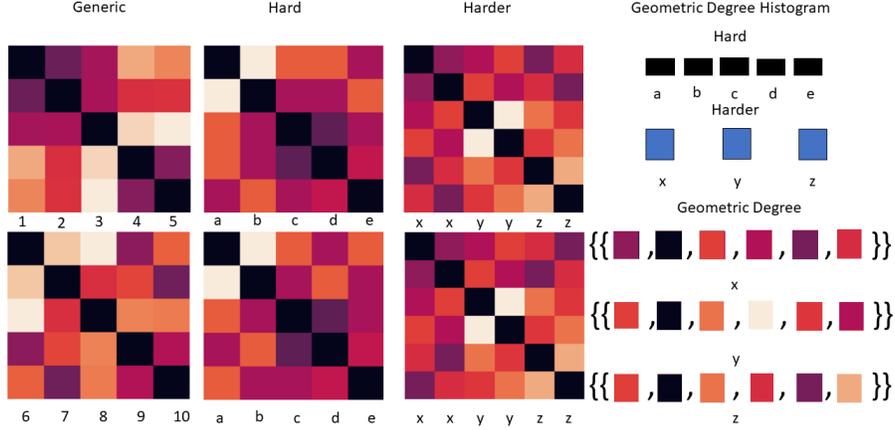


Figure 1: Distance matrices (Left), geometric degree histogram (Right) of pairs of point clouds. The *generic* pair is a randomly sampled pair of point clouds. Notice each of the nodes in each of the clouds has a distinct geometric degree. The *Hard* pair exhibits a distinct geometric degree for each node, but only within each point cloud, that is the pair shares an identical geometric degree histogram. The *Harder* example is a pair of point clouds with identical geometric degree histogram, and each point cloud is comprised of three pairs of points, with each pair having an identical geometric degree. Examples from Pozdnyakov and Ceriotti [2022] and Pozdnyakov et al. [2020].

The symmetry group we consider for point clouds $(x_1, \dots, x_n) \in \mathbb{R}^{d \times n}$ is generated by a rotation matrix $R \in \mathcal{SO}(d)$, and a permutation $\sigma \in S_n$. These act on a point cloud by

$$(R, \sigma)_*(x_1, \dots, x_n) = (Rx_{\sigma^{-1}(1)}, \dots, Rx_{\sigma^{-1}(n)}).$$

135 We denote this group by $\mathcal{SO}[d, n]$. In some instances, reflections $R \in \mathcal{O}(d)$ are also permitted,
 136 leading to a slightly larger symmetry group, which we denote by $\mathcal{O}[d, n]$. Our goal shall be to
 137 construct separating invariants for these groups. For the sake of brevity, we do not discuss *translation*
 138 invariance and separation, as these can easily be achieved by centering the input point clouds, once
 139 $\mathcal{SO}[d, n]$ (or $\mathcal{O}[d, n]$) separating invariants are constructed, see Dym and Gortler [2023].

140 For simplicity of notation, throughout this paper, we focus on the case $d = 3$. In Appendix C we
 141 explain how our constructions and theorems can be generalized to $d > 3$.

142 2 Euclidean Graph Isomorphism Tests

143 The k -WL Graph Isomorphism Test Weisfeiler and Leman [1968] is a classical paradigm for testing
 144 the isomorphism of combinatorial graphs, which we shall now briefly describe. Let \mathcal{G} be a graph with
 145 vertices indexed by $[n] = \{1, 2, \dots, n\}$. We denote each ordered k -tuple of vertices by a multi-index
 146 $\mathbf{i} = (i_1, \dots, i_k) \in [n]^k$. Essentially, for each such k -tuple \mathbf{i} , the test maintains a *coloring* $\mathbf{C}(\mathbf{i})$ that
 147 belongs to a discrete set, and updates it iteratively. First, the coloring of each k -tuple is assigned an
 148 initial value that encodes the *isomorphism type* of the corresponding k -dimensional subgraph:

$$\mathbf{C}_{(0)} = \mathbf{C}_{(0)}(\mathbf{i}), \mathbf{i} \in [n]^k. \quad (1)$$

149 Then the color of each k -tuple \mathbf{i} is iteratively refined according to the colors of its ‘neighboring’
 150 k -tuples. The update rule is given by

$$\mathbf{C}_{(t+1)}(\mathbf{i}) = \mathbf{Embed}^{(t+1)} \left(\mathbf{C}_{(t)}(\mathbf{i}), \{ \{ \mathbf{C}_{(t)}(\mathbf{i}[j \setminus 1]), \dots, \mathbf{C}_{(t)}(\mathbf{i}[j \setminus k]) \} \mid j \in [n] \} \right), \quad (2)$$

where $\mathbf{i}[j \setminus t]$ is the multi-index \mathbf{i} with its t -th coordinate replaced by j ; e.g. for $j = 1$, $\mathbf{i}[j \setminus 1] = (j, i_2, \dots, i_k)$. **Embed** is a function that maps its input injectively to some discrete set. This process is repeated T times to obtain a final coloring $\{ \{ \mathbf{C}_{(T)}(\mathbf{i}) \} \}_{\mathbf{i} \in [n]^k}$. A global label is then calculated by

$$\mathbf{C}_{\mathcal{G}} = \mathbf{Embed}^{(T+1)} \left(\{ \{ \mathbf{C}_{(T)}(\mathbf{i}) \} \mid \mathbf{i} \in [n]^k \} \right),$$

151 where $\mathbf{Embed}^{(T+1)}$ is a function that maps label-multisets injectively to some discrete set.

152 To test whether two graphs \mathcal{G} and \mathcal{G}' are isomorphic, the k -WL test computes the corresponding
 153 colorings $\mathbf{C}_{\mathcal{G}}$ and $\mathbf{C}_{\mathcal{G}'}$ for some chosen T . If $\mathbf{C}_{\mathcal{G}} \neq \mathbf{C}_{\mathcal{G}'}$, then \mathcal{G} and \mathcal{G}' are guaranteed not to be
 154 isomorphic, whereas if $\mathbf{C}_{\mathcal{G}} = \mathbf{C}_{\mathcal{G}'}$, then \mathcal{G} and \mathcal{G}' may either be isomorphic or not, and the test does
 155 not, in general, provide a decisive answer for combinatorial graphs. It is known that this test is able
 156 to distinguish a strictly larger class of combinatorial graphs for every strict increase in the value of k ,
 157 i.e. it is a strict hierarchy of tests in terms of distinguishing power Cai et al. [1992], Grohe [2017].

158 **Vanilla- k -WL tests** As a first step from a combinatorial to a Euclidean setting, we identify each
 159 point cloud $X = (x_1, \dots, x_n) \in \mathbb{R}^{d \times n}$ with a complete graph on n vertices, wherein each edge (i, j)
 160 is endowed with the weight $w_{ij}(X) = \langle x_i, x_j \rangle$. We name such a graph a *Euclidean graph*. Similarly
 161 to k -WL for combinatorial graphs, k -WL for Euclidean graphs maintains a coloring of the k -tuples of
 162 vertices. However, the initial color of each k -tuple \mathbf{i} is not a discrete label as in the combinatorial case,
 163 but rather a $k \times k$ matrix of continuous features, which represent all edge weights w_{ij} corresponding
 164 to pairs of indices from \mathbf{i} . We will call the k -WL test defined by this initial coloring the **vanilla k -WL**
 165 test. This test is invariant by construction to reflections, rotations, and permutations. We note that our
 166 definition of the vanilla k -EWL test via inner products follows that of Lim et al. [2022]. Another
 167 popular, and essentially equivalent, formulation, uses distances instead.

168 **k -EWL tests** An inherent limitation of the Vanilla-1-EWL test is that no pairwise Euclidean
 169 information is passed, yielding it rather uninformative. Indeed, Pozdnyakov and Ceriotti [2022]
 170 proposed a Euclidean analog of the 1-WL test, where the update rule (2) is replaced with

$$\mathbf{C}_{(t+1)}(i) = \mathbf{Embed}^{(t)}(\mathbf{C}_{(t)}(i), \{\{\mathbf{C}_{(t)}(j), \|x_i - x_j\|, j \neq i\}\}). \quad (3)$$

171 We call this test the **1-EWL** test. This formulation is motivated by the fact that many symmetry-
 172 preserving networks for point clouds are in fact a realization of it, though they use **Embed** functions
 173 that are continuous and, in general, may assign the same value to different multisets. Consequently,
 174 the separation power of these architectures is at most that of 1-EWL with discrete injective hash
 175 functions. Moreover, the separation power will be equivalent if continuous injective multiset functions
 176 are used for embedding, as we discuss in Section 4.

177 The 1-EWL test strengthens the Vanilla-1-EWL test by allowing the messages passed to a node
 178 in each step to contain not only previous colorings but also geometric information in the form of
 179 pairwise distances. More generally, we shall use the term **k -EWL** to refer to tests that follow the
 180 Euclidean k -WL paradigm, but incorporate geometric invariants into the message-passing procedure.
 181 In particular, for point clouds with dimension 3, we define the 2-SEWL test ('SE' for *Special*
 182 *Euclidean*) by replacing the update step (2) with

$$\mathbf{C}_{(t+1)}(i, j) = \mathbf{Embed}^{(t)}(\mathbf{C}_{(t)}(i, j), \{\{\mathbf{C}_{(t)}(k, j), \mathbf{C}_{(t)}(i, k), \langle x_i \times x_j, x_k \rangle\}\}_{k=1}^n}). \quad (4)$$

183 Note that $\langle x_i \times x_j, x_k \rangle$ is equal to the determinant of the 3×3 matrix whose rows are the three vectors
 184 x_i, x_j, x_k , which makes this a natural choice as all polynomial invariants of $\mathcal{SO}(3)$ are generated by
 185 these determinants and the inner products we use for the initial coloring Kraft and Procesi [1996].

186 We note that, Using the fact that $O(3)$ is just two copies of $SO(3)$, it is not difficult to generalize
 187 2-SEWL to a complete $\mathcal{O}[3, n]$ test, which we name 2-EWL. for general d , similar complete $(d - 1)$ -
 188 SEWL and $(d - 1)$ -EWL tests can be formulated for point clouds in \mathbb{R}^d via the Hodge-star operator;
 189 see Appendix C for more details.

190 In the rest of this section, we shall prove that the 2-SEWL, 2-EWL and vanilla 3-EWL tests are
 191 complete when applied to $\mathbb{R}^{3 \times n}$, even when using a single iteration ($T = 1$). We shall also show that
 192 two iterations of the 1-EWL test is complete, except on a set of measure zero.

193 2.1 Generic completeness of 1-EWL

The separation power of 1-EWL is closely linked to the notion of *geometric degree*: For a point
 cloud $X = (x_1, \dots, x_n)$, we define the geometric degree $d(i, X)$ of the i th point, and the induced
 geometric degree histogram $d_H(X)$, to be the multisets

$$d(i, X) = \{\{\|x_1 - x_i\|, \dots, \|x_n - x_i\|\}\}, \quad d_H(X) = \{\{d(1, X), \dots, d(n, X)\}\}.$$

It is not difficult to see that if $d_H(X) \neq d_H(Y)$ then X and Y can be separated by a single 1-EWL iteration. An example of such a pair is shown in the left of Figure 1. With two 1-EWL iterations, we show that can separate X and Y even if $d_H(X) = d_H(Y)$, provided that they both belong to the set of point clouds defined by

$$\mathbb{R}_{distinct}^{3 \times n} = \{X \in \mathbb{R}^{3 \times n} \mid d(i, X) \neq d(j, X) \ \forall i \neq j\}.$$

194 Such an example, taken from Pozdnyakov et al. [2020], is visualized in the middle column of Figure 1
195

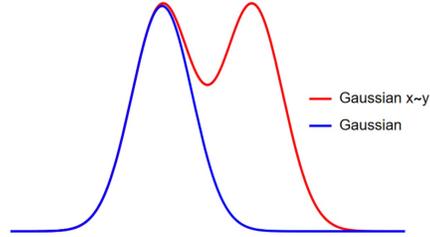
196 **Theorem 2.1.** *Two iterations of the 1-EWL test assign two point clouds $\mathcal{X}, Y \in \mathbb{R}_{distinct}^{3 \times n}$ the same*
197 *value, if and only if $X \stackrel{\mathcal{O}[3,n]}{=} Y$.*

198 In the appendix we show that the complement of $\mathbb{R}_{distinct}^{3 \times n}$ has measure zero. Thus this result
199 complements long-standing results for combinatorial graphs, stating that 1-WL can classify almost
200 all such graphs as the number of nodes tends to infinity Babai et al. [1980].

201 The right-most pair of point clouds ('Harder') in Figure 1 is taken from Pozdnyakov and Ceriotti
202 [2022]. The degree histograms of these point clouds are identical, and they are not in $\mathbb{R}_{distinct}^{3 \times n}$.
203 Pozdnyakov and Ceriotti [2022] show that this pair cannot be separated by any number of 1-EWL
204 iterations.

205 2.2 Is 1-EWL All You Need?

206 Theorem 2.1 shows that the probability of a failure of
207 the 1-EWL is zero. A natural question to ask is whether
208 more powerful tests are needed. We believe the answer to
209 this question is yes. Typical hypothesis classes used for
210 machine learning, such as neural networks, are Lipschitz
211 continuous Gama et al. [2020]. In this setting, failure to
212 separate on a measure zero set could have implications for
213 non-trivial positive measure. This phenomenon is depicted
214 in the figure in the inset. On the right, a plot of a Gaussian
215 distribution centered at $x \in \mathbb{R}$, depicting a target function is shown in blue. In red, a schematic plot
216 of how a Lipschitz continuous function that does not distinguish x from y would model the target
217 function.



218 3 2-SEWL and Vanilla 3-EWL are complete

219 We now prove that the vanilla 3-EWL test is complete.

220 **Theorem 3.1.** *For every $X, Y \in \mathbb{R}^{3 \times n}$, a single iteration of the vanilla 3-EWL test assigns X and Y*
221 *the same value if and only if $X \stackrel{\mathcal{O}[3,n]}{=} Y$.*

Proof. First, it is clear that if $X \stackrel{\mathcal{O}[3,n]}{=} Y$ then $\mathbf{C}_G(X) = \mathbf{C}_G(Y)$ since the vanilla 3-EWL test is invariant by construction. The challenge is proving the other direction. To this end, let us assume that $\mathbf{C}_G(X) = \mathbf{C}_G(Y)$, and assume without loss of generality that $r := \text{rank}(X) \geq \text{rank}(Y)$. Note that X has rank $r \leq 3$, and so it must contain some three points whose rank is also r . By applying a permutation to \bar{X} we can assume without loss of generality that these three points are the first three points. The initial coloring $\mathbf{C}_0(1, 2, 3)(X)$ of this triplet is their Gram matrix $(\langle x_i, x_j \rangle)_{1 \leq i, j \leq 3}$, which has the same rank r as the space spanned by the three points. Next, since $\mathbf{C}_G(X) = \mathbf{C}_G(Y)$ are the same, there exists a triplet of points i, j, k such that $\mathbf{C}_{(1)}(1, 2, 3)(X) = \mathbf{C}_{(1)}(i, j, k)(Y)$ which implies that the initial colorings are also the same. By applying a permutation to Y we can assume without loss of generality that $i = 1, j = 2, k = 3$. Next, since the Gram matrix of x_1, x_2, x_3 and y_1, y_2, y_3 are identical, there is an orthogonal transformation that takes x_i to y_i for $i = 1, 2, 3$, and by applying this transformation to all points in X we can assume without loss of generality that $x_i = y_i$ for $i = 1, 2, 3$. It remains to show that the rest of the points of X and Y are equal, up to permutation. To see this, first note that X and Y have the same rank since

$$r = \text{rank}(X) \geq \text{rank}(Y) \geq \text{rank}(y_1, y_2, y_3) = \text{rank}(x_1, x_2, x_3) = r.$$

Thus the space spanned by $x_1 = y_1, x_2 = y_2, x_3 = y_3$ contains all points in X and Y . Next, we can deduce from the aggregation rule defining $\mathbf{C}_1(1, 2, 3)(X)$ in (2), that

$$\{\{\langle x_j, x_1 \rangle, \langle x_j, x_2 \rangle, \langle x_j, x_3 \rangle \mid j \in [n]\}\} = \{\{\langle y_j, y_1 \rangle, \langle y_j, y_2 \rangle, \langle y_j, y_3 \rangle \mid j \in [n]\}\}.$$

222 Since all points in X and Y belong to the span of $x_1 = y_1, x_2 = y_2, x_3 = y_3$, X and Y are the same
223 up to permutation of the last $n - 3$ coordinates. This concludes the proof of the theorem. \square

224 We next outline the completeness proof of the more efficient 2-SEWL.

225 **Theorem 3.2.** For every $X, Y \in \mathbb{R}^{3 \times n}$, a single iteration of the 2-SEWL test assigns X and Y the
226 same value if and only if $X \stackrel{\mathcal{SO}[3, n]}{=} Y$.

227 *Proof idea.* The completeness of Vanilla-3-EWL was based on the fact that its initial coloring captures
228 the Gram matrix of triplets of vectors that span the space spanned by X , and on the availability of
229 projections onto this basis in the aggregation step defined in (2). Our proof for 2-EWL completeness
230 relies on the fact that a pair of non-degenerate vectors x_i, x_j induces a basis $x_i, x_j, x_i \times x_j$ of \mathbb{R}^3 .
231 The Gram matrix of this basis can be recovered from the Gram matrix of the first two points x_i, x_j ,
232 and the projection onto this basis can be obtained from the extra geometric information we added in
233 (18). A full proof is given in the appendix. \square

234 To conclude this section, we note that the above theorem can be readily used to also show that the
235 2-EWL test is also complete with respect to $\mathcal{O}[3, n]$. For details see Appendix A.

236 4 WL-equivalent GNNs with continuous features

237 In the previous section we discussed the generic completeness of 1-EWL and the completeness of
238 2-SEWL and vanilla 3-EWL. The **Embed** functions in these tests are hash functions, which can be
239 redefined independently for each pair of point clouds X, Y . In this section, our goal is to explain how
240 to construct GNNs with equivalent separation power to that of these tests, while choosing continuous,
241 piecewise differentiable **Embed** functions that are injective. While this question is well studied for
242 combinatorial graphs with discrete features Xu et al. [2018], Morris et al. [2018], Maron et al. [2019],
243 Aamand et al. [2022], here we focus on addressing it for Euclidean graphs with continuous features.

244 4.1 Multiset injective functions

245 Let us first review some known results on injective multiset functions. Recall that a function defined on
246 multisets with n elements coming from some alphabet $\Omega \subseteq \mathbb{R}^D$ can be identified with a permutation
247 invariant function defined on Ω^n . A multiset function is injective if and only if its corresponding
248 function on Ω^n is separating with respect to the action of the permutation group (see Definition 1.1).

249 In Corso et al. [2020], Wagstaff et al. [2022] it was shown that for any separating, permutation
250 invariant mappings from \mathbb{R}^n to \mathbb{R}^K , the embedding dimension K will be at least n . Two famous
251 examples of continuous functions that achieve this bound are

$$\Psi_{\text{sort}}(x_1, \dots, x_n) = \text{sort}(x_1, \dots, x_n) \quad \text{and} \quad \Psi_{\text{pow}}(x_1, \dots, x_n) = \left(\sum_{i=1}^n x_i^t \right)_{t=1}^n. \quad (5)$$

252 When the multiset elements are in \mathbb{R}^D , the picture is similar: if there exists a continuous, permutation
253 invariant and separating mapping from $\mathbb{R}^{D \times n}$ to \mathbb{R}^K , then necessarily $K \geq n \cdot D$ Joshi et al. [2022].
254 In Dym and Gortler [2023] it is shown that continuous separating invariants for $D > 1$, with near-
255 optimal dimension, can be derived from the $D = 1$ separating invariants $\Psi = \Psi_{\text{pow}}$ or $\Psi = \Psi_{\text{sort}}$,
256 by considering random invariants of the form

$$\mathbf{Embed}_\theta(x_1, \dots, x_n) = \langle b_j, \Psi(a_j^T x_1 \dots, a_j^T x_n) \rangle, \quad j = 1, \dots, K. \quad (6)$$

257 where each a_j and b_j are d and n dimensional random vectors, and we denote $\theta =$
258 $(a_1, \dots, a_K, b_1, \dots, b_K) \in \mathbb{R}^{K(D+n)}$. When $K = 2nD + 1$, for almost any choice of θ , the
259 function \mathbf{Embed}_θ will be separating on $\mathbb{R}^{D \times n}$. Thus the embedding dimension in this construction is
260 optimal up to a multiplicative constant of two.

261 An important property of this results of Dym and Gortler [2023] for our discussion, is that the
 262 embedding dimension K can be reduced if the domain of interest is a non-linear subset of $\mathbb{R}^{D \times n}$
 263 of low dimension. For example, if the domain of interest is a finite union of lines in $\mathbb{R}^{D \times n}$, then
 264 the *intrinsic dimension* of the domain is 1, and so we will only need an embedding dimension of
 265 $K = 2 \cdot 1 + 1 = 3$. Thus, the required embedding dimension depends on the intrinsic dimension of
 266 the domain rather than on its *ambient dimension*, which in our case is $n \cdot D$.

267 To formulate these results precisely we will need to introduce some real algebraic geometry terminol-
 268 ogy (see Basu et al. [2006] for more details): A *semi-algebraic subset* of a real finite-dimensional
 269 vector space is a finite union of subsets that are defined by polynomial equality and inequality
 270 constraints. For example, polygons, hyperplanes, spheres, and finite unions of these sets, are all
 271 semi-algebraic sets. A semi-algebraic set is always a finite union of manifolds, and its dimension is
 272 the maximal dimension of the manifolds in this union. Using these notions, we can now state the
 273 ‘intrinsic version’ of the results in Dym and Gortler [2023]:

274 **Theorem 4.1** (Dym and Gortler [2023]). *Let \mathcal{X} be an S_n -invariant semi-algebraic subset of $\mathbb{R}^{D \times n}$ of*
 275 *dimension $D_{\mathcal{X}}$. Denote $K = 2D_{\mathcal{X}} + 1$. Then for Lebesgue almost every $\theta \in \mathbb{R}^{K(D+n)}$ the mapping*
 276 *$\mathbf{Embed}_{\theta} : \mathcal{X} \rightarrow \mathbb{R}^K$ is S_n invariant and separating.*

277 4.2 Multiset injective functions for GNNs

278 We now return to discuss GNNs and explain the importance of the distinction between the intrinsic
 279 and ambient dimensions in our context. Suppose we are given n initial features $(h_1^{(0)}, \dots, h_n^{(0)})$ in
 280 \mathbb{R}^d , and for simplicity let us assume they are recursively refined via the simple aggregation rule:

$$h_i^{(t+1)} = \mathbf{Embed}^{(t)} \left(\{ \{ h_j^{(t)} \}_{j=1, j \neq i}^n \} \right). \quad (7)$$

281 Let us assume that each $\mathbf{Embed}^{(t)}$ is injective on the space of all multisets with $n - 1$ elements in
 282 the ambient space of $h_j^{(t)}$. Then the injectivity of $\mathbf{Embed}^{(1)}$ implies that $h_i^{(1)}$ is of dimension at least
 283 $(n - 1) \cdot d$. The requirement that $\mathbf{Embed}^{(2)}$ is injective on a mult-set of $n - 1$ features in $\mathbb{R}^{(n-1) \cdot d}$
 284 implies that $h_i^{(2)}$ will be of dimension at least $(n - 1)^2 \cdot d$. Continuing recursively with this argument
 285 we obtain an estimate of $\sim (n - 1)^T d$ for the dimensions of each $h_i^{(T)}$ after T iterations of (7).
 286

287 Fortunately the analysis presented above is
 288 overly pessimistic, because it focused only on
 289 the *ambient dimension*. Let us denote the matrix
 290 containing all n features at time t by $H^{(t)}$.
 291 Then $H^{(t)} = F_t(H^{(0)})$, where F_t is the con-
 292 catenation of all $\mathbf{Embed}^{(t')}$ functions from all
 293 previous time-steps. Thus $H^{(t)}$ resides in the set
 294 $F_t(\mathbb{R}^{d \times n})$. Here we again rely on results from
 295 algebraic geometry: if F_t is a composition of
 296 piecewise linear and polynomial mappings, then
 297 it is a semi-algebraic mapping, which means
 298 that $F_t(H^{(0)})$ will be a semi-algebraic set of di-
 299 mension $\dim(\mathbb{R}^{n \times d}) = n \cdot d$. This point will be
 300 explained in more detail in the proof of Theo-
 301 rem 4.2. By Theorem 4.1 we can then use \mathbf{Embed}_{θ}
 302 as a multiset injective function on \mathcal{X}_t with a fixed
 embedding dimension of $2n \cdot d + 1$ which does not depend on T . This is visualized in Figure 2.

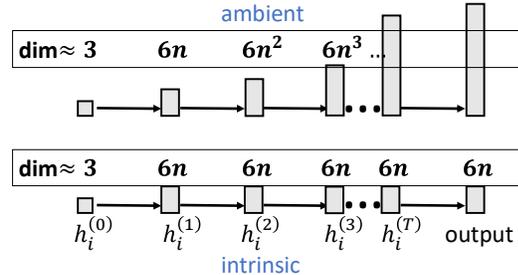


Figure 2: The exponential growth in the dimension that would result from only considering the ambient feature dimension can be avoided by exploiting the constant intrinsic dimension.

303 **2-SEWLnet** Based on the discussion above, we can devise architectures that simulate the various
 304 tests discussed in this paper and have reasonable feature dimensions throughout the construction, In
 305 particular, we can simulate T iterations of the 2-SEWL test by replacing all $\mathbf{Embed}^{(t)}$ functions¹ with
 306 $\mathbf{Embed}_{\theta}^{(t)}$, where in our implementation we choose $\Psi = \Psi_{\text{sort}}$ in (6). The embedding dimension for
 307 all t is taken to be $6n + 1$, since the input is in $\mathbb{R}^{3 \times n}$. We denote the obtained parametric function
 308 by F_{ϕ} . Based on a formalization of the discussion above, we prove in the appendix that F_{ϕ} has the
 309 separation power of the complete 2-SEWL test, and therefore F_{ϕ} is separating.

310 **Theorem 4.2.** Let F_ϕ denote the parametric function simulating the 2-SEWL test. Then for Lebesgue
 311 almost every ϕ the function $F_\phi : \mathbb{R}^{3 \times n} \rightarrow \mathbb{R}^{6n+1}$ is separating with respect to the action of $\mathcal{SO}[3, n]$.

312 To conclude this subsection, we note that while sort-based permutation invariants are used as aggrega-
 313 tors in GNNs Zhang et al. [2020, 2018], Blondel et al. [2020], the polynomial-based aggregators Ψ_{pow}
 314 are not as common. To a certain extent, one can use the approach in Xu et al. [2018], Maron et al.
 315 [2019], replace the polynomials in Ψ_{pow} by MLPs, and justify this by the universal approximation
 316 power of MLPs. A limitation of this approach is that it only guarantees separation at the limit.

317 5 Synthetic Experiments

318 In this section we implement 2-SEWLnet, described in Section 4, and empirically evaluate its
 319 separation power, and the separation power of alternative $\mathcal{SO}[3, n]$ invariant point cloud architectures.
 320 We trained the architectures on permuted and rotated variations of highly-challenging point-cloud
 321 pairs, and measured separation by the test classification accuracy. We considered three pairs of point
 322 clouds (Hard1-Hard3) from Pozdnyakov et al. [2020]. These pairs were designed to be challenging
 323 for distance-based invariant methods. However, our analysis reveals that they are in fact separable
 324 by two iterations of the 1-EWL test. We then consider a pair of point clouds from Pozdnyakov and
 325 Ceriotti [2022] which was proven to be indistinguishable by the 1-EWL tests. The results of this
 326 experiment are given in Table 1. Further details on the experimental setup appear in Appendix B.

Separation	complete	\cong 1-EWL	unknown	unknown	unknown
Point Clouds	2-SEWLnet	EGNN	MACE	TFN	GVPGNN
Hard1	100 %	100 %	100 %	100 %	100 %
Hard2	100 %	100 %	100 %	100 %	50 %
Hard3	100 %	100 %	100 %	100 %	95.0 \pm 15.0 %
Harder	100 %	50 %	100 %	100 %	53.7 \pm 13.1 %

Table 1: Separation accuracy on challenging 3D point clouds. Hard examples correspond to point clouds which cannot be distinguished by a single 1-EWL iteration but can be distinguished by two iterations, according to Theorem 2.1. The Harder example is a point cloud not distinguishable by 1-EWL Pozdnyakov and Ceriotti [2022]. GNN implementations and code pipeline based on Joshi et al. [2022].

327 As expected, we find that 2-SEWLnet, which has complete separation power, succeeded in perfectly
 328 separating all examples. We also found that EGNN Victor Garcia Satorras [2021], which is essentially
 329 an implementation of 1-EWL, does not separate the **Harder** example, but *does* separate the **Hard**
 330 example after two iterations, as predicted by Theorem 2.1. We also considered three additional
 331 invariant point cloud models whose separation power is not as well understood. We find that MACE
 332 Batatia et al. [2022] and TFN Thomas et al. [2018] achieve perfect separation, (when applying them
 333 with at least 3-order correlations and three-order $\mathcal{SO}(3)$ representations). The third GVPGNN Jing
 334 et al. [2021] architecture attains mixed results. We note that we cannot necessarily deduce from our
 335 empirical results that MACE and TFN are complete. While it is true that TFN is complete when
 336 considering arbitrarily high order representations Dym and Maron [2020], it is not clear whether
 337 order three representation suffices for complete separation. We conjecture that this is not the case.
 338 However, finding counterexamples is a challenging problem we leave for future work.

339 **Future Work** In this work, we presented several invariant tests for point clouds that are provably
 340 complete, and have presented and implemented 2-SEWL-net which simulates the complete 2-SEWL
 341 test. Currently, this is a basic implementation that only serves to corroborate our theoretical results. A
 342 practically useful implementation requires addressing several challenges, including dealing with point
 343 clouds of different sizes, the non-trivial $\sim n^4$ complexity of computing even the relatively efficient
 344 2-SEWL-net, and finding learning tasks where complete separation leads to gains in performance.
 345 We are actively researching these directions and hope this paper will inspire others to do the same.

¹A minor technicality is that the **Embed** functions are actually defined on vector-multiset pairs. This issue is discussed in the proof of the theorem.

References

- 346
347 Patrick Reiser, Marlen Neubert, Andr'e Eberhard, Luca Torresi, Chen Zhou, Chen Shao, Houssam
348 Metni, Clint van Hoesel, Henrik Schopmans, Timo Sommer, and Pascal Friederich. Graph neural
349 networks for materials science and chemistry. *Communications Materials*, 3(1):93, 2022. doi:
350 10.1038/s43246-022-00315-6.
- 351 Zixuan Yi, Iadh Ounis, and Craig Macdonald. Graph contrastive learning with positional represen-
352 tation for recommendation. In Jaap Kamps, Lorraine Goeuriot, Fabio Crestani, Maria Maistro,
353 Hideo Joho, Brian Davis, Cathal Gurrin, Udo Kruschwitz, and Annalina Caputo, editors, *Ad-
354 vances in Information Retrieval*, pages 288–303, Cham, 2023. Springer Nature Switzerland. ISBN
355 978-3-031-28238-6.
- 356 Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural
357 message passing for quantum chemistry. *CoRR*, 2017.
- 358 Limei Wang, Yi Liu, Yuchao Lin, Haoran Liu, and Shuiwang Ji. Comenet: Towards complete and
359 efficient message passing for 3d molecular graphs, 2022.
- 360 Marc Finzi, Max Welling, and Andrew Gordon Wilson. A practical method for constructing equivari-
361 ant multilayer perceptrons for arbitrary matrix groups. *arXiv preprint arXiv:2104.09459*, 2021.
- 362 Xu Ma, Yuqian Zhou, Huan Wang, Can Qin, Bin Sun, Chang Liu, and Yun Fu. Image as set of points,
363 2023.
- 364 Soledad Villar, David W Hogg, Kate Storey-Fisher, Weichi Yao, and Ben Blum-Smith. Scalars
365 are universal: Equivariant machine learning, structured like classical physics. In M. Ran-
366 zato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in
367 Neural Information Processing Systems*, volume 34, pages 28848–28863. Curran Associates,
368 Inc., 2021. URL [https://proceedings.neurips.cc/paper_files/paper/2021/file/
369 f1b0775946bc0329b35b823b86eeb5f5-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/f1b0775946bc0329b35b823b86eeb5f5-Paper.pdf).
- 370 Chaitanya K. Joshi, Cristian Bodnar, Simon V. Mathis, Taco Cohen, and Pietro Liò. On the expressive
371 power of geometric graph neural networks. *NeurIPS Workshop on Symmetry and Geometry in
372 Neural Representations*, 2022.
- 373 Zhengdao Chen, Soledad Villar, Lei Chen, and Joan Bruna. On the equivalence between graph
374 isomorphism testing and function approximation with gnns. *Advances in neural information
375 processing systems*, 32, 2019.
- 376 Nadav Dym and Haggai Maron. On the universality of rotation equivariant point cloud networks.
377 *ArXiv*, abs/2010.02449, 2020.
- 378 Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley.
379 Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds.
380 *arXiv preprint arXiv:1802.08219*, 2018.
- 381 Derek Lim, Joshua Robinson, Lingxiao Zhao, Tess Smidt, Suvrit Sra, Haggai Maron, and Stefanie
382 Jegelka. Sign and basis invariant networks for spectral graph representation learning. *arXiv
383 preprint arXiv:2202.13013*, 2022.
- 384 Sergey N. Pozdnyakov and Michele Ceriotti. Incompleteness of graph neural networks for points
385 clouds in three dimensions, 2022.
- 386 Vikraman Arvind and Gaurav Rattan. The complexity of geometric graph isomorphism. In *Electron.
387 Colloquium Comput. Complex.*, volume 21, page 70, 2014.
- 388 Nadav Dym and Shahar Ziv Kovalsky. Linearly converging quasi branch and bound algorithms for
389 global rigid registration. In *Proceedings of the IEEE/CVF International Conference on Computer
390 Vision*, pages 1628–1636, 2019.
- 391 V. Kurlin. Computable complete invariants for finite clouds of unlabeled points under euclidean
392 isometry. *ArXiv*, abs/2207.08502, 2022.

- 393 Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen, Gorkirt
394 Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks.
395 *arXiv preprint arXiv:1810.02244*, 2018.
- 396 Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably powerful graph
397 networks, 2019. URL <https://arxiv.org/abs/1905.11136>.
- 398 Emiel Hooeboom. Max Welling Victor Garcia Satorras. E(n) equivariant graph neural networks.
399 *Proceedings of the 38-th International Conference on Machine Learning*, PMLR(139), 2021.
- 400 Kristof Schütt, Pieter-Jan Kindermans, Huziel Enoc Saucedo Felix, Stefan Chmiela, Alexandre
401 Tkatchenko, and Klaus-Robert Müller. Schnet: A continuous-filter convolutional neural network
402 for modeling quantum interactions. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus,
403 S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*,
404 volume 30. Curran Associates, Inc., 2017. URL [https://proceedings.neurips.cc/paper_](https://proceedings.neurips.cc/paper_files/paper/2017/file/303ed4c69846ab36c2904d3ba8573050-Paper.pdf)
405 [files/paper/2017/file/303ed4c69846ab36c2904d3ba8573050-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/303ed4c69846ab36c2904d3ba8573050-Paper.pdf).
- 406 Yael Kapon, Abhijit Saha, Tal Duanis-Assaf, Thijs Stuyver, Amir Ziv, Tzuriel Metzger, Shira Yochelis,
407 Sason Shaik, Ron Naaman, Meital Reches, et al. Evidence for new enantiospecific interaction
408 force in chiral biomolecules. *Chem*, 7(10):2787–2799, 2021.
- 409 Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural
410 networks?, 2018.
- 411 Nadav Dym and Steven J. Gortler. Low dimensional invariant embeddings for universal geometric
412 learning, 2023.
- 413 Ben Finkelshtein, Chaim Baskin, Haggai Maron, and Nadav Dym. A simple and universal rotation
414 equivariant point-cloud network. *arXiv preprint arXiv:2203.01216*, 2022.
- 415 Johannes Gasteiger, Florian Becker, and Stephan Günnemann. Gemnet: Universal directional graph
416 neural networks for molecules, 2021. URL <https://arxiv.org/abs/2106.08903>.
- 417 Georg Bökman, Fredrik Kahl, and Axel Flinth. Zz-net: A universal rotation equivariant architecture
418 for 2d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
419 *Recognition*, pages 10976–10985, 2022.
- 420 Daniel Widdowson and Vitaliy Kurlin. Recognizing rigid patterns of unlabeled point clouds by
421 complete and continuous isometry invariants with no false negatives and no false positives. *arXiv*
422 *preprint arXiv:2303.15385*, 2023.
- 423 Peter Brass and Christian Knauer. Testing the congruence of d-dimensional point sets. In *Proceedings*
424 *of the sixteenth annual symposium on Computational geometry*, pages 310–314, 2000.
- 425 Daniel Widdowson and Vitaliy Kurlin. Resolving the data ambiguity for periodic crystals. In
426 *Advances in Neural Information Processing Systems*, 2022.
- 427 Omri Puny, Matan Atzmon, Edward J Smith, Ishan Misra, Aditya Grover, Heli Ben-Hamu, and
428 Yaron Lipman. Frame averaging for invariant and equivariant network design. In *International*
429 *Conference on Learning Representations*, 2021.
- 430 Sergey N Pozdnyakov, Michael J Willatt, Albert P Bartók, Christoph Ortner, Gábor Csányi, and
431 Michele Ceriotti. Incompleteness of atomic structure representations. *Physical Review Letters*,
432 125(16):166001, 2020.
- 433 Boris Weisfeiler and A. A. Leman. The reduction of a graph to canonical form and the algebra which
434 appears therein. *Nauchno-Technicheskaya Informatsia*, 2:12–16, 1968.
- 435 Jin-Yi Cai, Martin Fürer, and Neil Immerman. An optimal lower bound on the number of variables
436 for graph identification. *Combinatorica*, 12(4):389–410, 1992. doi: 10.1007/BF01305232.
- 437 Martin Grohe. *Descriptive complexity, canonisation, and definable graph structure theory*, volume 47.
438 Cambridge University Press, 2017.

- 439 Hanspeter Kraft and Claudio Procesi. Classical invariant theory, a primer. *Lecture Notes. Preliminary*
440 *version*, 1996.
- 441 László Babai, Paul Erdős, and Stanley M Selkow. Random graph isomorphism. *SIAM Journal on*
442 *Computing*, 9(3):628–635, 1980.
- 443 Fernando Gama, Joan Bruna, and Alejandro Ribeiro. Stability properties of graph neural networks.
444 *IEEE Transactions on Signal Processing*, 68:5680–5695, 2020. doi: 10.1109/tsp.2020.3026980.
445 URL <https://doi.org/10.1109%2Ftsp.2020.3026980>.
- 446 Anders Aamand, Justin Chen, Piotr Indyk, Shyam Narayanan, Ronitt Rubinfeld, Nicholas Schiefer,
447 Sandeep Silwal, and Tal Wagner. Exponentially improving the complexity of simulating the
448 weisfeiler-lehman test with graph neural networks. *Advances in Neural Information Processing*
449 *Systems*, 35:27333–27346, 2022.
- 450 Gabriele Corso, Luca Cavalleri, Dominique Beaini, Pietro Liò, and Petar Veličković. Principal
451 neighbourhood aggregation for graph nets. *Advances in Neural Information Processing Systems*,
452 33:13260–13271, 2020.
- 453 Edward Wagstaff, Fabian B Fuchs, Martin Engelcke, Michael A Osborne, and Ingmar Posner.
454 Universal approximation of functions on sets. *Journal of Machine Learning Research*, 23(151):
455 1–56, 2022.
- 456 Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in real algebraic geometry*,
457 volume 10. Springer, 2006.
- 458 Yan Zhang, Jonathon Hare, and Adam Prügel-Bennett. Fspool: Learning set representations with
459 featurewise sort pooling, 2020.
- 460 Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen. An end-to-end deep learning
461 architecture for graph classification. In *Proceedings of the AAAI conference on artificial intelligence*,
462 volume 32, 2018.
- 463 Mathieu Blondel, Olivier Teboul, Quentin Berthet, and Josip Djolonga. Fast differentiable sorting
464 and ranking. In *International Conference on Machine Learning*, pages 950–959. PMLR, 2020.
- 465 Ilyes Batatia, David P Kovacs, Gregor Simm, Christoph Ortner, and Gabor Csanyi. Mace:
466 Higher order equivariant message passing neural networks for fast and accurate force fields.
467 In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in*
468 *Neural Information Processing Systems*, volume 35, pages 11423–11436. Curran Associates,
469 Inc., 2022. URL [https://proceedings.neurips.cc/paper_files/paper/2022/file/](https://proceedings.neurips.cc/paper_files/paper/2022/file/4a36c3c51af11ed9f34615b81edb5bbc-Paper-Conference.pdf)
470 [4a36c3c51af11ed9f34615b81edb5bbc-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/4a36c3c51af11ed9f34615b81edb5bbc-Paper-Conference.pdf).
- 471 Bowen Jing, Stephan Eismann, Patricia Suriana, Raphael John Lamarre Townshend, and Ron Dror.
472 Learning from protein structure with geometric vector perceptrons. In *International Conference on*
473 *Learning Representations*, 2021. URL <https://openreview.net/forum?id=1YLJDvSx6J4>.
- 474 J.R. Munkres. *Topology*. Featured Titles for Topology. Prentice Hall, Incorporated, 2000. ISBN
475 9780131816299. URL <https://books.google.co.il/books?id=XjoZAQAAIAAJ>.