

EDGE-BASED WL AND MESSAGE PASSING

Anonymous authors

Paper under double-blind review

ABSTRACT

We propose EB-1WL, an edge-based color-refinement test, and a corresponding GNN architecture, EB-GNN. Our architecture is inspired by a classic triangle counting algorithm by Chiba and Nishizeki, and explicitly uses triangles during message passing. We achieve the following results: (1) EB-1WL is significantly more expressive than 1-WL. Further, we provide a complete logical characterization of EB-1WL based on first-order logic, and matching distinguishability results based on homomorphism counting. (2) In an important distinction from previous proposals for more expressive GNN architectures, EB-1WL and EB-GNN require near-linear time and memory on practical graph learning tasks. (3) Empirically, we show that EB-GNN is a highly-efficient general-purpose architecture: It substantially outperforms simple MPNNs, and remains competitive with task-specialized GNNs while being significantly more computationally efficient.

1 INTRODUCTION

Graph learning, and in particular message-passing graph neural networks (GNNs), have emerged as a fundamental tool across the sciences (Scarselli et al., 2009; Kipf & Welling, 2017; Wu et al., 2019). A major factor in the wide applicability of GNNs is the robust theoretical framework that grounds their study in principled comparisons between architectures. Core to this framework is the characterization of the expressive power of GNNs in terms of the Weisfeiler–Leman (1WL) test, a central notion in the theoretical study of graph similarity (Morris et al., 2019; Xu et al., 2019). Owing to this connection, alternative, yet equally insightful, characterizations of GNN expressive power have been developed—such as via finite-variable fragments of counting logics (Cai et al., 1992) or homomorphism counts from classes of graphs of bounded treewidth (Dvorák, 2010; Dell et al., 2018).


While these perspectives enrich our understanding of GNN capabilities, they also make clear that 1WL-based GNNs remain limited in important ways. To highlight the necessity for more expressive tests than 1WL, an illustrative limitation of 1WL-based GNNs is their inability to detect small motifs or count triangles (Chen et al., 2020; Arvind et al., 2020; Lanzinger & Barceló, 2024). The higher-order WL hierarchy offers a systematic remedy: properties that elude 1WL are often captured at some higher level, i.e., by k WL for some $k > 1$. Yet this increase in expressive power comes at a steep computational cost. Even GNNs inspired by 2WL demand quadratic memory and cubic runtime in the number of nodes, rendering them impractical for large graphs (Maron et al., 2019). This creates a central challenge: to bridge the gap between the tractability of 1WL and the expressivity of 2WL. Importantly, progress on this front must rest on strong theoretical foundations—whether refined isomorphism tests, logical characterizations, or homomorphism-count perspectives—so that the precise expressive power of new architectures is rigorously understood.

Recent work on the so-called *neighbor-communication* 1WL (NC-1WL) test (Liu et al., 2024) moves in this direction, extending 1WL by incorporating edge information within the neighborhood of each node in the graph. NC-1WL is strictly more expressive than 1WL and remains efficient, making it an attractive refinement from a practical standpoint. However, from a theoretical perspective, its study remains incomplete: while NC-1WL has a well-defined placement within the WL hierarchy, it lacks the wealth of alternative characterizations (e.g., logical or homomorphism-count based) that have proven essential in understanding the power and limitations of 1WL and its higher-dimensional variants. Moreover, its computational complexity has not been subjected to a refined analysis, leaving open how well its efficiency scales across graph classes of varying density.

We build on the previous framework but switch the focus from nodes to edges, achieving a clear improvement in expressivity without increasing complexity. This shift yields a more powerful

054 yet elegantly simple architecture, whose theoretical properties follow naturally from its design.
 055 Guided by this idea, we introduce the *edge-based* 1WL test (EB-1WL)—a refinement of 1WL
 056 that remains computationally efficient and rigorously grounded. EB-1WL updates edge colors
 057 through incidence relations and triangle-induced interactions (see Fig. 1 and Sec. 3), an edge-centric
 058 perspective inspired by the classic triangle-counting algorithm of Chiba & Nishizeki (1985). This
 059 approach retains much of 2WL’s relational strength while avoiding its combinatorial blowup: each
 060 iteration runs in $O(\alpha m)$ time, where m is the number of edges and α is the graph’s arboricity.
 061 Since α is typically small in real-world graphs¹, EB-1WL
 062 achieves near-linear performance in practice. We further give
 063 three expressiveness results for EB1-WL:

- 064 1. EB-1WL strictly extends the expressive power of NC-1WL.
- 065 2. On the logical side, it admits a precise characterization in
 066 terms of *clique-based* finite-variable fragments with count-
 067 ing quantifiers.
- 068 3. On the homomorphism-count side, we prove a lower bound,
 069 showing that EB-1WL is at least as powerful as counting
 070 homomorphisms from chordal graphs of treewidth two.
 071



072 Figure 1: Message passing scheme
 073 of EB-1WL and EB-GNN. Colors
 074 correspond to aggregation types.

072 In combination, these results position EB-1WL as a natural
 073 generalization of the 1WL test, while still being highly efficient in practice. Its improvement in
 074 expressivity over 1WL/NC-1WL is due to a conceptual change in the architecture—not a heuristic
 075 variant, but a principled intermediate point in the WL hierarchy.

076 Building on this foundation, we introduce EB-GNNs, a message-passing architecture that exactly
 077 matches the expressive power of EB-1WL. EB-GNNs perform strongly across diverse settings: (1) on
 078 synthetic benchmarks, they capture structural patterns far beyond triangle counts and, in some cases,
 079 rival higher-order tests; (2) on molecular datasets, they outperform standard baselines and remain
 080 competitive with task-specialized GNNs while being significantly more computationally efficient;
 081 and (3) on large-scale graphs, they scale effectively and achieve state-of-the-art accuracy. Together,
 082 these results establish EB-GNNs as an expressive yet highly efficient general-purpose architecture.

083 **Further related work.** Surprisingly, edge-based GNN architectures are rare in the literature. Most
 084 directly, Zhang et al. (2020) propose an architecture based on edge convolution. Additionally, Cai
 085 et al. (2022) studied the application of standard message-passing GNNs to the line graph² for link
 086 prediction.

087 Another line of related work concerns the study of efficient higher-order GNNs (in the usual vertex-
 088 centered view). Here various approaches have been proposed to improve on higher-order k -WL
 089 performance bottlenecks especially for sparse graphs. To this end Morris et al. (2020) introduced
 090 δ - k -WL and its local variants, which demonstrate improvements on sparse graphs while maintaining
 091 high expressivity. Similar motivation is followed by Zhao et al. (2022) who introduce $(k, c)(\leq)$ -
 092 SetWL, a proposal to approximate higher-order power in a fine-grained way by also embedding local
 093 substructures. While these refinements are more scalable than plain 2WL, the vertex-based paradigm
 094 they follow makes quadratic (in the number of vertices) running time unavoidable, even in sparse
 095 graphs³. Moreover, such practical considerations remove these methods from the strong theoretical
 096 foundations of k -WL, which admits well-known and highly influential characterizations in terms of
 097 homomorphism counts (Dvorák, 2010; Dell et al., 2018) and variants of first-order logic (Cai et al.,
 098 1992). Along the same lines, other prominent approaches to obtain efficient higher-order GNNs such
 099 as PPGN (Maron et al., 2019) also cannot avoid fundamentally quadratic (or higher) computational
 100 time complexity per layer.

101 ¹In all 39 datasets considered by Eppstein et al. (2013), the arboricity is at most 201 even though their largest
 102 graph has 3.7 million nodes and 16.5 million edges. On 32/39 of their datasets, the arboricity is less than 60. We
 103 note that Eppstein et al. (2013) report the degeneracy, which is an upper bound on the arboricity.

104 ²In the *line graph* of a graph G , the edges of G become vertices that are connected according to their
 105 incidences in G .

106 ³The closest comparison to our method in terms of $(k, c)(\leq)$ -SetWL is achieved with $k = 3, c = 1$ (lower k
 107 yields expressivity at most 1-WL). On a star graph this materializes all quadratically many 2-hop paths as nodes
 in a “super-graph” on which message-passing is performed. In contrast, a star has arboricity 1 and our method is
 strictly linear in its running time.

In contrast, our proposed EB-GNNs differ substantially from these methods: rather than adapting classic WL-style message passing to edges, we leverage established algorithmic insights to introduce a fundamentally new type of message passing (cf. Figure 1) that is only possible at the edge level.

2 PRELIMINARIES

Graphs. We study undirected graphs $G = (V, E)$, where V is the set of vertices and $E \subseteq \binom{V}{2}$ is the set of edges. We set $n = |V|$ and $m = |E|$. We write $N(v) := \{w \mid \{v, w\} \in E\}$ for the *neighborhood* of $v \in V$. For technical simplicity, we assume that graphs do not have isolated nodes.

The *arboricity* α of a graph $G = (V, E)$ is the minimum number of forests that partition its edge set E (Diestel (2012)). More formally, the arboricity is the smallest integer of k such that there exist forests F_1, \dots, F_k with $F_i = (V, E_i)$ such that $\bigcup_{i=1}^k E_i = E$. The arboricity is also tightly related to other important graph parameters, such as the *degeneracy* or the density of the *densest subgraph* (Nash-Williams, 1961). It is widely known that real-world datasets such as social networks, road networks or planar graphs have very small arboricities (Eppstein et al. (2013)).

WL test. The Weisfeiler–Leman test (WL test) is a family of combinatorial algorithms for distinguishing graphs through iterative refinement of vertex- or tuple-colorings (Weisfeiler & Leman, 1968; Cai et al., 1992). The most widely used variant is the 1WL test, or *color refinement*. Given a graph $G = (V, E)$, this algorithm assigns each vertex $v \in V$ a color $\text{cr}^{(\ell)}(G, v)$ at iteration $\ell \geq 0$, defined inductively as follows: the initial color is constant, $\text{cr}^{(0)}(G, v) := 1$, and the update rule is

$$\text{cr}^{(\ell+1)}(G, v) := \left(\text{cr}^{(\ell)}(G, v), \{ \{ \text{cr}^{(\ell)}(G, u) \mid u \in N(v) \} \} \right).$$

At each iteration ℓ , the coloring induces a partition of the vertex set of G , where the partition at iteration $\ell + 1$ always refines that at iteration ℓ . Once this process stabilizes, we write $\text{cr}(G, v)$ for the final color assigned to vertex v . We define $\text{cr}(G)$ as the multiset $\{ \{ \text{cr}(G, v) \mid v \in V \} \}$, and call two graphs G and G' *distinguishable by 1WL* if $\text{cr}(G) \neq \text{cr}(G')$.

Higher-order versions of the WL test are also widely studied in the literature. Rather than focusing on individual vertices, the k WL test, for $k > 1$, considers k -tuples of vertices. Each k -tuple is assigned a color that reflects the isomorphism type of the subgraph induced by those vertices. At each refinement step, the color of a k -tuple is updated by considering all possible ways of replacing one of its vertices with another vertex in the graph. For example, when $k = 2$, each ordered pair (u, v) refines its color by looking at pairs such as (u, w) and (w, v) for all possible w . In this way, the k WL test captures not only the view of individual vertices, but also the structural relations within patterns of size k .

Neighbor-communication WL test. An extension of 1WL that incorporates information about the edges between the neighbors of a node has been proposed recently (Liu et al., 2024). The resulting test, called *NC-1WL test*, where NC stands for *neighbor communication*, proceeds similarly to 1WL, but updates the color $\text{nc}^{(\ell)}(G, v)$ of each vertex v in a graph $G = (V, E)$ according to the rule

$$\text{nc}^{(\ell+1)}(G, v) := \left(\text{nc}^{(\ell)}(G, v), \{ \{ \text{nc}^{(\ell)}(G, u) \mid u \in N(v) \}, \{ \{ \text{nc}^{(\ell)}(G, u), \text{nc}^{(\ell)}(G, w) \mid u, w \in N(v), \{u, w\} \in E \} \} \} \right).$$

In other words, besides taking into account the multiset of colors of the neighbors of a vertex v , as in 1WL, the NC-1WL test also considers the multiset of color pairs corresponding to the edges within the neighborhood of v . We let $\text{nc}(G, v)$ denote the color of node v once the coloring partition on vertices defined by the NC-1WL test becomes stable. We write $\text{nc}(G)$ for the multiset $\{ \{ \text{nc}(G, v) \mid v \in V \} \}$, and call two graphs G and G' *distinguishable by NC-1WL* if $\text{nc}(G) \neq \text{nc}(G')$.

The additional structural information collected by NC-1WL makes it strictly more powerful than 1WL in distinguishing certain classes of non-isomorphic graphs. In turn, every pair of graphs that can be distinguished by NC-1WL can also be distinguished by 2WL, but the converse does not hold. The advantage of NC-1WL over 2WL, however, is that it achieves stronger discriminative power while still operating at the *vertex level*, in the same spirit as 1WL. In particular, its computational cost is closer to that of 1WL than to the more demanding 2WL procedure.

3 EDGE-BASED WL TEST

In this section, we introduce the EB-1WL test (for edge-based 1WL), an extension of 1WL and NC-1WL that colors edges rather than vertices—analogueous to higher-order WL tests that color vertex tuples. Unlike those tests, EB-1WL colors only edge pairs, reducing space complexity from quadratic in nodes to linear in edges and making it more practical. By focusing on edges, EB-1WL gains greater expressive power than NC-1WL while maintaining reasonable computational cost.

Let $G = (V, E)$ be an undirected graph. We iteratively associate a color $\text{eb}^{(\ell)}(G, (u, v))$ with each ordered pair (u, v) such that $\{u, v\} \in E$. That is, each edge receives two colors, one for each ordering of its endpoints. The coloring of the ordered pair (u, v) is defined inductively. Initially, every pair has the same color: $\text{eb}^{(0)}(G, (u, v)) = 1$. At iteration $\ell + 1$, the color of (u, v) is updated according to

$$\text{eb}^{(\ell+1)}(G, (u, v)) = \left(\text{eb}^{(\ell)}(G, (u, v)), \right. \quad (1)$$

$$\left. \{\{\text{eb}^{(\ell)}(G, (u, x)) \mid x \in N(u)\}\}, \right. \quad (2)$$

$$\left. \{\{\text{eb}^{(\ell)}(G, (u, y)), \text{eb}^{(\ell)}(G, (v, y))\} \mid y \in N(v) \cap N(u)\}\}, \right. \quad (3)$$

$$\left. \{\{\text{eb}^{(\ell)}(G, (v, z)) \mid z \in N(v)\}\}\right). \quad (4)$$

This update rule refines the color of an edge based on the colors of edges incident to its endpoints as can be seen in Fig. 1. Eq. (2) captures the influence of edges incident to u , Eq. (4) does the same for edges incident to v , and Eq. (3) encodes the interaction between edges that are incident to both u and v , thereby forming a triangle and incorporating the local neighborhood structure around the edge.


We let $\text{eb}(G, (u, v))$ denote the color of the ordered pair (u, v) once the coloring partition of the edges defined by the NC-1WL test becomes stable. We write $\text{eb}(G)$ for the multiset $\{\{\text{eb}(G, (u, v)), \text{eb}(G, (v, u)) \mid \{u, v\} \in E\}\}$ and call two graphs G and G' *distinguishable by EB-1WL* if they have a different number of vertices or $\text{eb}(G) \neq \text{eb}(G')$.

Computational cost of EB-1WL. Next, we study the computational cost of EB-1WL. We consider an idealized computational model where a multiset of s elements can be created in time $O(s)$ and stored in $O(1)$ space (in practice, this can be achieved with high probability using hashing). In this model, we need $O(m)$ space to store the graph and the colors of the edges.

Now we analyze the time needed to perform a single iteration. First, computing the multisets in Eq.s (2) and (4) for all nodes u and v can be done in total time $O(m)$, since for every node u (and, resp., v) we spend time proportional to its degree. The more interesting part is computing Eq. (3) efficiently for all *ordered* edges (u, v) . For this, we need to go through all common neighbors y of u and v . A naïve approach would iterate over all neighbors y of u and then check whether (y, v) exists. This becomes slow if u has a much higher degree than v , potentially taking total time $O(md)$, where d is the maximum degree of the graph. Instead, following the triangle-enumeration algorithm of Chiba & Nishizeki (1985), we iterate only over the neighbors of the lower-degree endpoint of (u, v) and check adjacency to the other endpoint. This still visits every common neighbor y of u and v exactly once, so every triangle (u, v, y) is included in the summation in Eq. (3); only the order and implementation of the enumeration change, and the permutation-invariant aggregation therefore remains unaffected. The analysis of Chiba & Nishizeki (1985) then implies a running time of $O(\alpha m)$, where α is the arboricity of the graph. Appendix A provides more details on this algorithm and its connection to EB1-WL.

Proposition 1. *An iteration of EB-1WL can be performed using $O(m)$ space and $O(\alpha m)$ time.*

Comparing with the NC-1WL architecture, one can note that although this architecture requires just $O(n)$ space to store colors, it still requires $O(m)$ for storing the graph, providing no advantage over EB-1WL. Similarly, the NC-1WL architecture requires going through all triangles of the graph, and doing this polynomially faster than in $O(\alpha m)$ time would violate conjectures from the computational complexity community (Kopelowitz et al., 2016; Vassilevska Williams & Xu, 2020).

216
217
218
219
220
221
222
223
224
225
226Figure 2: The graphs G and H from the proof of Theorem 2.227
228

4 THE DISTINGUISHING POWER OF EB-1WL

229

4.1 EXPRESSIVITY

230

We first observe that the EB-1WL test is at least as expressive as the NC-1WL test, and consequently also at least as expressive as the classical 1WL test, in distinguishing non-isomorphic graphs. Moreover, EB-1WL is strictly more expressive than NC-1WL for distinguishing graphs.

231
232
233

Theorem 2. *Every pair of graphs distinguishable by NC-1WL is also distinguishable by EB-1WL. Additionally, the graphs G and H in Fig. 2 can be distinguished by EB-1WL but not by NC-1WL. Thus, EB-1WL is strictly more expressive than NC-1WL and 1WL.*

234
235
236

Proof idea. All nodes of G and H have the same degree and lie in the same number of triangles. Hence, they are not distinguished by NC-1WL. However, they are distinguished by EB-1WL since every edge in H is part of two triangles, whereas G has edges that belong to only a single triangle. \square

240
241
242

Similar to NC-1WL, any pair of graphs distinguishable by EB-1WL is also distinguishable by 2WL, though the reverse implication does not hold. A concrete example of a pair of graphs that is distinguishable by 2WL but neither by EB-1WL nor NC-1WL is shown in Fig. 3.

243
244
245
246

4.2 LOGICAL CHARACTERIZATION

247
248

The distinguishing power of the k -WL test can be characterized via a fragment of first-order logic, namely the $(k+1)$ -variable fragment with counting quantifiers (Cai et al., 1992), with 1WL and 2WL corresponding to the two- and three-variable fragments, respectively. These logical characterizations have been an important tool in studying GNNs, as they offer helpful insights into what GNNs can accomplish. We show that EB-1WL also admits a natural logical characterization, further demonstrating the robustness and naturalness of our newly proposed framework. Specifically, EB-1WL corresponds to a novel natural logic we introduce below. This logic lies strictly between the two- and three-variable counting fragments, providing a *precise* conceptualization of its position “between” 1WL and 2WL.

249
250
251

We assume familiarity with the syntax and semantics of first-order logic (FO). We write $\phi(\bar{x})$ to indicate that the free variables of ϕ are exactly those in the tuple \bar{x} . When interpreted over graphs, FO formulas are defined over a vocabulary consisting of a single binary relation symbol E . Specifically, the formula $E(x, y)$ is interpreted over a graph $G = (V, E)$ as the set of all pairs $(u, v) \in V \times V$ such that $\{u, v\} \in E$. Given a tuple $\bar{x} = (x_1, \dots, x_n)$ of distinct variables, we define the FO formula

252
253
254

$$\text{clique}(\bar{x}) := \bigwedge_{1 \leq i < j \leq n} E(x_i, x_j),$$

255
256
257

so that its interpretation over G is the set of all n -tuples of vertices that form a clique in G .

266

We now introduce our novel *clique-based first-order logic with counting (CFOC)* via the following syntax:

267
268
269

1. If \bar{x} is a tuple of distinct variables, then $\text{clique}(\bar{x})$ is a formula in CFOC.




Figure 3: Graphs distinguishable by 2WL but not by EB-1WL.

- 270 2. If $\phi(\bar{x})$ is a formula in CFOC, then so is $\text{clique}(\bar{x}) \wedge \neg\phi(\bar{x})$.
 271
 272 3. If $\phi(\bar{x})$ and $\psi(\bar{y})$ are formulas in CFOC, then so is $\text{clique}(\bar{z}) \wedge (\phi(\bar{x}) \star \psi(\bar{y}))$ for any
 273 $\star \in \{\wedge, \vee\}$, where \bar{z} is the tuple obtained by collecting all variables that occur in \bar{x} and \bar{y} ,
 274 removing any duplicates, so that each variable appears exactly once.
 275
 276 4. If $\phi(\bar{x}, y)$ is a formula in CFOC, then $\text{clique}(\bar{x}) \wedge \exists^{\geq k} y \phi(\bar{x}, y)$ is a formula in CFOC, for
 277 any integer $k \geq 1$, where the semantics of $\exists^{\geq k} y \phi(\bar{x}, y)$ assert that there exist at least k
 278 vertices v such that $\phi(\bar{x}, v)$ holds when $y = v$.

279 Intuitively, CFOC restricts FO with counting quantifiers to formulas whose free variables form a
 280 clique in the graph. We note that this is reminiscent of, but not directly related to, the so-called
 281 *clique-guarded fragments* of first-order logic; see, e.g., Grädel (1999).

282 We denote CFOC^3 the fragment of CFOC that consists of formulas that use at most three variables.
 283 For example, the CFOC^3 formula $\psi := \exists^{\geq 1} x \exists^{\geq 1} y (\text{clique}(x, y) \wedge \exists^{\geq 3} z \text{clique}(x, y, z))$ checks if
 284 there is an edge which is a part of at least 3 triangles. A CFOC^3 *sentence* is a CFOC^3 formula without
 285 free variables. We call graphs G and H *distinguishable by CFOC^3* , if there is a CFOC^3 sentence ϕ
 286 such that $G \models \phi$ but $H \not\models \phi$. We can now establish our characterization:

287 **Theorem 3.** *The pairs of graphs that are distinguishable by EB-1WL are precisely those that are*
 288 *distinguishable by CFOC^3 .*

289 **Example 1.** The sentence ψ shown above holds in graph H from Figure 2, but not in graph G . ▲
 290

291 292 4.3 DISTINGUISHING POWER BASED ON HOMOMORPHISM COUNTS

293 A central theme in the study of the WL test is its characterization via homomorphism counts. Formally,
 294 a *homomorphism* from a graph $G = (V, E)$ to a graph $H = (V', E')$ is a mapping $h : V \rightarrow V'$ such
 295 that $\{h(u), h(v)\} \in E'$ for every edge $\{u, v\} \in E$. The connection between WL and homomorphism
 296 counts is as follows: two graphs are distinguishable by k WL if and only if they differ in the number
 297 of homomorphisms from some graph of *treewidth* at most k (Dvorák, 2010; Dell et al., 2018). For
 298 $k = 1$, this corresponds to the class of trees, and for $k = 2$ to the class of *series-parallel* graphs.
 299

300 We show that distinguishability by EB-1WL is at least as powerful as distinguishing graphs by
 301 homomorphism counts from the class of *chordal* graphs of treewidth two, which strictly lies between
 302 the classes of graphs of treewidth one and two⁴. Here, recall that a graph is chordal if it has no
 303 induced subgraph that is a cycle of length 4 or larger. This offers additional support for viewing
 304 EB-1WL as a natural and well-founded counterpart to the standard 1WL test.

305 **Theorem 4.** *If two graphs have a different number of homomorphisms from some chordal graph of*
 306 *treewidth at most 2, they are distinguishable by EB-1WL.*

307 **Example 2.** Consider a graph J formed by two triangles sharing an edge. We observe that the
 308 number of homomorphisms from J into G is strictly smaller than the number of homomorphisms
 309 from J into H , where G and H are the graphs shown in Fig. 2. Refer to Appendix B.2 for details. ▲
 310

311 312 5 EDGE-BASED GRAPH NEURAL NETWORKS

313 Now we introduce the *EB-GNN architecture*, a message-passing framework whose expressive power
 314 coincides with that of the EB-1WL test. Formally, a d -dimensional *EB-GNN* \mathcal{T} with $t > 0$ layers is
 315 specified by parameters $a_i, b_i, c_i, u_i, v_i \in \mathbb{R}^d$ and $A_i, C_i, U_i, V_i \in \mathbb{R}^{d \times d}$, for $i = 1, \dots, t$. Given a
 316 graph $G = (V, E)$, the EB-GNN \mathcal{T} assigns to each ordered edge (u, v) with $\{u, v\} \in E$ and each
 317 layer $0 \leq i \leq t$ a feature vector $f^{(i)}(u, v) \in \mathbb{R}^d$. At the input layer, we set⁵
 318

$$319 f^{(0)}(u, v) = (1 \quad 0 \quad \dots \quad 0)^T \in \mathbb{R}^d. \quad (5)$$

320
 321
 322 ⁴Note that, by definition, trees are always chordal.

323 ⁵When nodes or edges have additional input features, one can take them into account while defining
 $f^{(0)}(u, v)$, see Section 6.

For $1 \leq i \leq t$, the update rules are given by

$$\alpha^{(i)}(u) = \sum_{x \in N(u)} \text{ReLU}\left(A_i \cdot f^{(i-1)}(u, x) + a_i\right), \quad (6)$$

$$\beta^{(i)}(u, v) = \sum_{y \in N(u) \cap N(v)} \text{ReLU}\left(B_i \cdot \begin{pmatrix} f^{(i-1)}(u, y) \\ f^{(i-1)}(v, y) \end{pmatrix} + b_i\right), \quad (7)$$

$$\gamma^{(i)}(v) = \sum_{z \in N(v)} \text{ReLU}\left(C_i \cdot f^{(i-1)}(v, z) + c_i\right), \quad (8)$$

$$g^{(i)}(u, v) = f^{(i-1)}(u, v) + \alpha^{(i)}(u) + \beta^{(i)}(u, v) + \gamma^{(i)}(v), \quad (9)$$

$$f^{(i)}(u, v) = g^{(i)}(u, v) + \text{FFN}_{U_i, u_i, V_i, v_i}(g^{(i)}(u, v)), \quad (10)$$

where $\text{ReLU}(x) = \max\{0, x\}$ is applied coordinate-wise in the equations above, and $\text{FFN}_{U, u, V, v}(x) = V \cdot \text{ReLU}(Ux + u) + v$. The overall output of \mathcal{T} on G is defined as

$$\mathcal{T}(G) = \sum_{\{u, v\} \in E} f^{(t)}(u, v). \quad (11)$$

Graphs G and H are *distinguishable* by EB-GNNs if there exists an EB-GNN \mathcal{T} with $\mathcal{T}(G) \neq \mathcal{T}(H)$.

It is immediate that the distinguishing power of EB-GNNs cannot exceed that of EB-1WL. Notably, we can show that this upper bound is tight.

Theorem 5. *Pairs of graphs distinguishable by EB-1WL are also distinguishable by EB-GNNs.*

In our proof, EB-GNNs require dimension $O(m + t)$, where m is the number of edges and t the number of triangles. Whether this can be reduced to $O(\log n)$, as shown for simulating 1WL by MPNNs (Aamand et al., 2022), remains open. Moreover, by replacing ReLU with any analytic non-polynomial activation and concatenation in Eq. (7) with a Hadamard product, one can prove Thm. 5 already holds for $d = 1$, following techniques from Amir et al.; Bravo et al.; Hordan et al.

Algorithmic implementation. In our implementation we perform a preprocessing step in which we enumerate all triangles in time $O(\alpha m)$ using the algorithm of Chiba & Nishizeki (1985). Since each triangle (u, v, y) corresponds to an edge (u, v) and a node $y \in N(u) \cap N(v)$, we can obtain the sets $N(u) \cap N(v)$ for all edges (u, v) in time $O(\alpha m)$. After that, each iteration of EB-GNN takes time $O(m + t)$, where t is the number of triangles—implying a better running time per iteration than in Proposition (1) because $t = O(m\alpha)$. However, due to this preprocessing step now we use memory $O(m + t)$ (instead of just $O(m)$).

6 EXPERIMENTAL EVALUATION

Now we empirically evaluate EB-GNN. The primary goal of our experiments is to show that EB-GNN provides a fast and expressive, general-purpose GNN architecture, and therefore we evaluate EB-GNN across tasks from diverse domains. We compare EB-GNN against Message Passing Neural Networks (MPNNs), another widely used general-purpose architecture, as well as against state-of-the-art models specifically optimized for each corresponding task. Our implementation is available at <https://anonymous.4open.science/r/EB-GNN-CB41/>.

We focus on graph-level predictions and predictions on existing edges. While graph-level prediction tasks are a well-established use case for predictive GNNs (Morris et al., 2019; Paolino et al., 2024; Southern et al., 2025), prediction tasks on existing edges have received less attention. This latter task is relevant in chemistry: rather than relying on costly molecular simulations, GNNs can directly predict quantum mechanical properties, e.g., the bond length between atoms (Li et al., 2024).

Experiment setup. Next, we describe the datasets and baseline models used for comparison. We evaluate EB-GNN on two synthetic datasets designed to measure its practically realized expressivity. Additionally, we assess performance on three real-world datasets: two composed of small molecular

Table 1: Empirical results on expressivity datasets. GIN + C_3 is an MPNN that uses triangle subgraph counts as additional node features. MPNN and 2WL results on BREC are from Wang & Zhang (2024). For details on runtime constants see Table 4.

Model	Runtime	CSL	BREC			
		Accuracy (\uparrow)	# Distinguishable Graph Pairs (\uparrow)			
			Basic	Reg.	Ext.	CFI
MPNN	$\mathcal{O}(n)$	10%	0	0	0	0
MPNN + C_3	$\mathcal{O}(\alpha m)$	20%	0	0	0	0
2WL	$\mathcal{O}(n^3)$	–	60	50	100	60
I^2 -GNN	$\mathcal{O}(nd^2)$	–	60	100	100	21
DRFWL	$\mathcal{O}(nd^4)$	–	60	50	99	0
4- ℓ -GIN	$\mathcal{O}(nd^3)$	60%	60	100	95	2
NC-GNN	$\mathcal{O}(\alpha m)$	20%	52	48	0	0
EB-GNN (ours)	$\mathcal{O}(\alpha m)$	20%	59	48	60	0

Table 2: Results on MalNet-Tiny. Top three models as 1st, 2nd, 3rd. All results except EB-GNN are from Southern et al. (2025).

Method	MalNet-Tiny Accuracy (\uparrow)
MPNN	91.10 \pm 0.98
HyMN	92.84 \pm 0.52
GPS (Perf.)	92.14 \pm 0.24
GPS (BigBird)	91.02 \pm 0.48
GPS (Transf.)	90.85 \pm 0.68
NC-GNN	92.50 \pm 0.56
EB-GNN	93.22 \pm 0.41

graphs and one consisting of large cybersecurity graphs. Consistent with previous observations that real-world graphs have low arboricity, we find that these datasets exhibit small arboricity: ≤ 3 for molecular graphs and ≤ 15 for large cybersecurity graphs with each dataset having < 4 mean arboricity. For each dataset, we compare EB-GNN against both general-purpose models and state-of-the-art task-specific baselines. Details on model and experiment setup are in App. C

Synthetic datasets for measuring expressivity. We evaluate the empirical expressivity of EB-GNN on two synthetic datasets: CSL (Murphy et al., 2019; Dwivedi et al., 2023) and BREC (Wang & Zhang, 2024). The CSL dataset consists of 150 graphs with 41 nodes each, grouped into 10 distinct isomorphism classes. These classes, defined by skip connections between Hamiltonian cycles, are all indistinguishable by 1WL. The graph-level task is to classify each graph according to its isomorphism class. The BREC dataset comprises pairs of graphs that are indistinguishable by 1WL but distinguishable by 3WL. For each pair, the graph-level task is to compute embeddings that correctly differentiate the two graphs.

On the synthetic datasets, we compare EB-GNN against 2WL to assess how much of 2WL’s expressivity EB-GNN can replicate while maintaining asymptotically faster runtimes. In addition, we include a comparison with the MPNN GIN (Xu et al., 2019) augmented with triangle subgraph counts as node features (Bouritsas et al., 2022), referred to as MPNN + C_3 . This comparison helps isolate how much of EB-GNN’s expressivity gain stems from its ability to count triangles, as captured by the β aggregation in Eq. (7). We also compare against NC-GNN to determine whether our theoretical increase in expressivity (Thm. 2) can also be measured empirically. Furthermore, we compare against all models used as baselines in other experiments for which we could find publicly available results.

Molecular edge-level and graph-level tasks. We further evaluate EB-GNN on a range of molecular edge-level and graph-level prediction tasks. Li et al. (2024) introduce the QMD dataset, which contains 65 000 molecular graphs annotated with various quantum mechanical properties. We assess EB-GNN on all four edge-level regression tasks from QMD. We evaluate four diverse graph-level regression tasks from QMD, selected to represent varied objectives (most other graph-level tasks in QMD focus on predicting HOMO/LUMO gaps). We also evaluate EB-GNN on 12 graph-level regression tasks from the widely used QM9 dataset (Wu et al., 2018), following common practice (Morris et al., 2019; Paolino et al., 2024). QM9 comprises 130 000 molecular graphs representing molecules of up to 9 atoms. On average, graphs in both QMD and QM9 contain fewer than 20 nodes.

For QMD tasks, we compare EB-GNN against D-MPNN (Yang et al., 2019; Dai et al., 2016), a directed MPNN that has demonstrated strong performance on chemical prediction tasks (Vermeire et al., 2022; Heid & Green, 2022) and is integrated into the widely adopted ChemProp framework (Heid et al., 2024). Similar to EB-GNN, D-MPNN performs message passing on edges rather than nodes. Furthermore, we also compare against NC-GNN (Liu et al., 2024) which we train using the same hyperparameter tuning procedure as EB-GNN. For QM9, we compare against a standard MPNN and several expressive GNN architectures: 1-2-3 GNN (Morris et al., 2019), DTNN (Schütt et al., 2017; Wu et al., 2018), NestedGNN (Zhang & Li, 2021), I2-GNN (Huang et al., 2023), DRFWL (Zhou et al., 2023), and the recent state-of-the-art 5- ℓ -GIN baseline (Paolino et al., 2024).

Table 3: MAE (\downarrow) on QMD. Best model per task marked **blue**. D-MPNN results from Li et al. (2024).

Model	Edge-level Tasks (MAE \downarrow)				Graph-level Tasks (MAE \downarrow)			
	Bond Index (unitless) $\times 10^{-3}$	Bond Length (\AA) $\times 10^{-3}$	Bonding Electrons (e) $\times 10^{-2}$	Natural Ionicity (unitless) $\times 10^{-4}$	IP $\times 10^{-3}$	EA $\times 10^{-3}$	Dipole Moment (debye) $\times 10^{-1}$	Traceless Quad. Mom. (debye \AA) $\times 10^0$
D-MPNN	6.65	4.48	1.46	9.00	4.29	4.06	4.59	1.62
NC-GNN	4.82 ± 0.04	3.33 ± 0.03	1.28 ± 0.01	5.39 ± 0.25	5.3 ± 0.03	4.35 ± 0.03	4.45 ± 0.02	1.57 ± 0.01
EB-GNN	4.42 ± 0.01	3.17 ± 0.011	1.19 ± 0.018	4.64 ± 0.01	5.49 ± 0.26	4.46 ± 0.12	4.33 ± 0.05	1.55 ± 0.01

Table 4: Normalized test MAE (\downarrow) on QM9 dataset. Top three models as **1st**, **2nd**, **3rd**. Table based on Paolino et al. (2024). For runtime, n is the number of nodes, m the number of edges, c and s are maximum size of subgraph sizes, d the maximum degree, and α the arboricity.

Target (MAE \downarrow)	Model							
Runtime	MPNN $\mathcal{O}(n)$	1-2-3 GNN $\mathcal{O}(n^3)$	DTNN $\mathcal{O}(m)$	NestedGNN $\mathcal{O}(ncd)$	I2-GNN $\mathcal{O}(nsd^2)$	DRFWL $\mathcal{O}(nd^4)$	5- ℓ GIN $\mathcal{O}(nd^5)$	EB-GNN $\mathcal{O}(\alpha)$
μ ($\times 10^{-1}$)	4.93	4.76	2.44	4.28	4.28	3.46	3.50 ± 0.11	3.15 ± 0.02
α ($\times 10^{-1}$)	7.8	2.7	9.5	2.90	2.30	2.22	2.17 ± 0.25	2.08 ± 0.03
ϵ_{homo} ($\times 10^{-3}$)	3.21	3.37	3.88	2.65	2.61	2.26	2.05 ± 0.05	2.18 ± 0.01
ϵ_{lumo} ($\times 10^{-3}$)	3.55	3.51	5.12	2.97	2.67	2.25	2.16 ± 0.04	2.17 ± 0.02
$\Delta(\epsilon)$ ($\times 10^{-3}$)	4.9	4.8	11.2	3.8	3.8	3.24	3.21 ± 0.14	3.02 ± 0.01
R^2	34.1	22.9	17.0	20.5	18.64	15.04	13.21 ± 0.19	13.83 ± 0.12
ZVPE ($\times 10^{-4}$)	12.4	1.9	17.2	2.	1.4	1.7	1.27 ± 0.03	1.26 ± 0.01
U_0	2.32	0.0427	2.43	0.295	0.211	0.156	0.0418 ± 0.0520	0.063 ± 0.002
U	2.08	0.111	2.43	0.361	0.206	0.153	0.023 ± 0.023	0.078 ± 0.008
H	2.23	0.0419	2.43	0.305	0.269	0.145	0.0352 ± 0.0304	0.0564 ± 0.0075
G	1.94	0.0469	2.43	0.489	0.261	0.156	0.0118 ± 0.0015	0.076 ± 0.006
C_v	0.27	0.0944	2.43	0.174	0.0730	0.0901	0.0702 ± 0.0024	0.092 ± 0.001

Cybersecurity. To evaluate our model on a different domain with larger graphs, we conduct experiments on MalNet-Tiny (Freitas et al., 2021). MalNet-Tiny consists of 5000 graphs with an average of over 1500 nodes, sampled from the MalNet dataset. The graph-level task is to classify whether a function call is benign or belongs to one of four malicious classes (AdWare, Trojan, Addisplay, Downloader). For comparison, we include a standard MPNN and the recently proposed subgraph GNN HyMN (Southern et al., 2025). We further compare against the graph transformer GPS (Rampásek et al., 2022) using different attention mechanisms: Performer (Choromanski et al., 2021), Big Bird (Zaheer et al., 2020), and the standard Transformer (Vaswani et al., 2017). Finally, we also compare against NC-GNN (Liu et al., 2024) which we train in the same fashion as EB-GNN.

Results. Recall that comparisons are made against both general-purpose MPNNs and state-of-the-art models for each dataset. Our experiments demonstrate that EB-GNN consistently outperforms standard MPNNs and remains competitive with state-of-the-art approaches.

Results on synthetic data. Table 1 presents the results of our expressivity experiments on the synthetic datasets. On CSL, EB-GNN achieves an accuracy of 20%, outperforming a vanilla MPNN (10% accuracy) and matching an MPNN augmented with triangle counts (MPNN + C_3 , 20% accuracy). On BREC, the gains of EB-GNN cannot be attributed to triangle counting, as evidenced by the MPNN + C_3 results. Remarkably, EB-GNN achieves performance on Basic and Regular graphs that is nearly identical to 2WL, though it fails to distinguish any CFI pairs. These results indicate that EB-GNN’s empirically realized expressivity extends well beyond triangle counting and, in many cases, approaches 2WL expressivity while maintaining significantly faster runtimes. We outperform NC-GNN in two out of five categories and tie it in the remaining two. Notably, on Extension graphs EB-GNN solves 60 instances while NC-GNN solves 0, highlighting our empirical increase in expressivity.

486 *Molecular edge-level and graph-level tasks.* Table 3 presents the edge-level and graph-level results
 487 on QMD, compared against the baseline D-MPNN (Li et al., 2024). EB-GNN outperforms D-MPNN
 488 across all edge-level tasks, reducing the mean absolute error by 20% to 50% depending on the
 489 task. On graph-level tasks, EB-GNN and D-MPNN perform comparably, with EB-GNN slightly
 490 outperforming D-MPNN on two tasks and slightly underperforming on the other two. EB-GNN
 491 outperforms NC-GNN in six out of eight datasets and only loses to NC-GNN for datasets where
 492 D-MPNN is the best performing model. Table 4 summarizes the results on QM9, where we compare
 493 against several expressive GNNs. The current state-of-the-art on this dataset is 5-ℓGNN (Paolino
 494 et al., 2024), which significantly improved over previous architectures. EB-GNN ranks as the best or
 495 second-best model on 11 out of 12 tasks, achieving comparable or better performance than 5-ℓGNN
 496 on 7 tasks (with at most 6% lower accuracy) and outperforming it on 4 tasks. Moreover, EB-GNN is
 497 substantially more computationally efficient than 5-ℓGNN: while 5-ℓGNN has asymptotic runtime
 498 $\mathcal{O}(nd^5)$, EB-GNN only requires $\mathcal{O}(am)$ time. In our experiments, EB-GNN is approximately 4
 499 times faster (see App. C).

500 *Results on graph-level tasks for malware detection.* Table 2 shows the results on MalNet-Tiny.
 501 EB-GNN outperforms all other models, including a standard MPNN, various graph transformers
 502 (GPS), and the subgraph GNN HyMN, demonstrating that EB-GNN generalizes well to other domains
 503 and scales effectively to larger graphs.

504 7 CONCLUSION, LIMITATIONS AND FUTURE WORK

506 We propose an edge-based message passing algorithm that combines high expressivity with near-
 507 linear runtime on sparse graphs. Our analysis fully characterizes its expressivity in terms of first-order
 508 logic and establishes a meaningful lower bound via homomorphism counting. Empirically, our
 509 architecture outperforms standard MPNNs while remaining competitive with more expressive models,
 510 all at a substantially lower runtime. We observe that EB-GNN is a general-purpose architecture:
 511 while it achieves strong empirical results, it does not rely on specialized techniques used on top of
 512 basic architectures known to improve GNN performance, such as using subgraph counts (Bouritsas
 513 et al., 2022), homomorphism counts (Barceló et al., 2021; Welke et al., 2023; Jin et al., 2024) or
 514 positional encodings (You et al., 2019; Ying et al., 2021; Ma et al., 2023; Bao et al., 2025).

515 *Limitations.* Although EB-1WL and EB-GNN achieve near-linear running times for sparse graphs,
 516 their efficiency depends on arboricity and triangle enumeration, which can be expensive on very dense
 517 graphs, potentially limiting scalability in such settings. Furthermore, EB-GNN produces embeddings
 518 only for edges preventing direct application to node-level tasks and standard link-prediction methods
 519 that rely on node embeddings. We leave extending EB-GNN to these tasks as future work.

520 *Future work.* An interesting open problem left by our work is whether the converse of Thm. 4 holds,
 521 i.e., whether graphs distinguishable by EB-1WL are exactly those distinguishable by homomorphism
 522 counts from some chordal graph of treewidth 2. Another line of future work concerns the tradeoff
 523 between expressiveness and generalization: recent results show that greater expressiveness need not
 524 harm generalization if matched to task demands and training data (Maskey et al., 2025), and can even
 525 help when graphs are well separated by large margins (Li et al., 2025). EB-GNNs strike a principled
 526 balance — more expressive than NC-1WL, yet far cheaper than 2WL — while showing strong results
 527 across benchmarks. Future work should broaden empirical evaluation to fully assess this balance of
 528 expressiveness, scalability, and generalization.

529
530
531
532
533
534
535
536
537
538
539

540 **Reproducibility statement.** Our theory is based on our EB-1WL algorithm which is described
 541 in Section 3. We analyze its properties in Section 4 and propose EB-GNN in Section 5 which is
 542 at least as expressive as EB-1WL. Our proofs are in Appendix B. We describe our experiments in
 543 Section 6 with more details in Appendix C. Our implementation as well as instructions to recreate our
 544 experiments are available at <https://anonymous.4open.science/r/EB-GNN-CB41/>.

546 REFERENCES

- 547
 548 Anders Aamand, Justin Y. Chen, Piotr Indyk, Shyam Narayanan, Ronitt Rubinfeld, Nicholas Schiefer,
 549 Sandeep Silwal, and Tal Wagner. Exponentially improving the complexity of simulating the
 550 weisfeiler-lehman test with graph neural networks. In *NeurIPS*, 2022.
- 551 Tal Amir, Steven J. Gortler, Ilai Avni, Ravina Ravina, and Nadav Dym. Neural injective functions for
 552 multisets, measures and graphs via a finite witness theorem. In *NeurIPS*, 2023.
- 553
 554 Vikraman Arvind, Frank Fuhlbrück, Johannes Köbler, and Oleg Verbitsky. On weisfeiler-lehman
 555 invariance: Subgraph counts and related graph properties. *Journal of Computer and System*
 556 *Sciences*, 113:42–59, 2020.
- 557 Linus Bao, Emily Jin, Michael Bronstein, İsmail İlkan Ceylan, and Matthias Lanzinger. Homomor-
 558 phism counts as structural encodings for graph learning. In *ICLR*, 2025.
- 559
 560 Pablo Barceló, Floris Geerts, Juan L. Reutter, and Maksimilian Ryschkov. Graph neural networks
 561 with local graph parameters. In *NeurIPS*, pp. 25280–25293, 2021.
- 562
 563 Giorgos Bouritsas, Fabrizio Frasca, Stefanos Zafeiriou, and Michael M Bronstein. Improving graph
 564 neural network expressivity via subgraph isomorphism counting. *IEEE Transactions on Pattern*
 565 *Analysis and Machine Intelligence*, 45(1):657–668, 2022.
- 566
 567 César Bravo, Alexander Kozachinskiy, and Cristobal Rojas. On dimensionality of feature vectors in
 mpnns. In *ICML*, 2024.
- 568
 569 Jin-yi Cai, Martin Fürer, and Neil Immerman. An optimal lower bound on the number of variables
 570 for graph identification. *Comb.*, 12(4):389–410, 1992.
- 571
 572 Lei Cai, Jundong Li, Jie Wang, and Shuiwang Ji. Line graph neural networks for link prediction.
IEEE Trans. Pattern Anal. Mach. Intell., 44(9):5103–5113, 2022.
- 573
 574 Zhengdao Chen, Lei Chen, Soledad Villar, and Joan Bruna. Can graph neural networks count
 575 substructures? In *NeurIPS*, 2020.
- 576
 577 Norishige Chiba and Takao Nishizeki. Arboricity and subgraph listing algorithms. *SIAM Journal on*
computing, 14(1):210–223, 1985.
- 578
 579 Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas
 580 Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, David Belanger, Lucy
 581 Colwell, and Adrian Weller. Rethinking attention with performers. In *ICLR*, 2021.
- 582
 583 Hanjun Dai, Bo Dai, and Le Song. Discriminative embeddings of latent variable models for structured
 data. In *ICML*, 2016.
- 584
 585 Holger Dell, Martin Grohe, and Gaurav Rattan. Lovász meets weisfeiler and leman. In *ICALP*,
 586 volume 107, pp. 40:1–40:14, 2018.
- 587
 588 Reinhard Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate texts in mathematics*. Springer,
 2012. ISBN 978-3-642-14278-9.
- 589
 590 Zdenek Dvorák. On recognizing graphs by numbers of homomorphisms. *J. Graph Theory*, 64(4):
 591 330–342, 2010.
- 592
 593 Vijay Prakash Dwivedi, Chaitanya K. Joshi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and
 Xavier Bresson. Benchmarking graph neural networks. *Journal of Machine Learning Research*, 24
 (43):1–48, 2023.

- 594 David Eppstein, Maarten Löffler, and Darren Strash. Listing all maximal cliques in large sparse
595 real-world graphs. *ACM J. Exp. Algorithmics*, 18, 2013.
- 596
- 597 Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In
598 *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- 599
- 600 Scott Freitas, Yuxiao Dong, Joshua Neil, and Duen Horng Chau. A large-scale database for graph
601 representation learning. In *NeurIPS Datasets and Benchmarks Track*, 2021.
- 602
- 603 Erich Grädel. Decision procedures for guarded logics. In *CADE*, pp. 31–51, 1999.
- 604
- 605 Esther Heid and William H. Green. Machine learning of reaction properties via learned representations
606 of the condensed graph of reaction. In *Journal of Chemical Information and Modeling*, 2022.
- 607
- 608 Esther Heid, Kevin P. Greenman, Yunsie Chung, Shih-Cheng Li, David E. Graff, Florence H. Vermeire,
609 Haoyang Wu, William H. Green, and Charles J. McGill. Chemprop: A machine learning package
610 for chemical property prediction. In *Journal of Chemical Information and Modeling*, 2024.
- 611
- 612 Snir Hordan, Tal Amir, and Nadav Dym. Weisfeiler leman for Euclidean equivariant machine learning.
613 In *ICML*, pp. 18749–18784, 2024.
- 614
- 615 Yinan Huang, Xingang Peng, Jianzhu Ma, and Muhan Zhang. Boosting the cycle counting power of
616 graph neural networks with i^2 -gnns. In *ICLR*, 2023.
- 617
- 618 Emily Jin, Michael M. Bronstein, İsmail İlkan Ceylan, and Matthias Lanzinger. Homomorphism
619 counts for graph neural networks: All about that basis. In *ICML*, 2024.
- 620
- 621 Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks.
622 In *ICLR*, 2017.
- 623
- 624 Tsvi Kopelowitz, Seth Pettie, and Ely Porat. Higher lower bounds from the 3sum conjecture. In
625 *SODA*, pp. 1272–1287. SIAM, 2016.
- 626
- 627 Matthias Lanzinger and Pablo Barceló. On the power of the weisfeiler-leman test for graph motif
628 parameters. In *ICLR*, 2024.
- 629
- 630 Shih-Cheng Li, Haoyang Wu, Angiras Menon, Kevin A. Spiekermann, Yi-Pei Li, and William H.
631 Green. When do quantum mechanical descriptors help graph neural networks to predict chemical
632 properties? *Journal of the American Chemical Society*, 146(33):23103–23120, 2024.
- 633
- 634 Shouheng Li, Floris Geerts, Dongwoo Kim, and Qing Wang. Towards bridging generalization and
635 expressivity of graph neural networks. In *ICLR*, 2025.
- 636
- 637 Meng Liu, Haiyang Yu, and Shuiwang Ji. Empowering gnns via edge-aware weisfeiler-leman
638 algorithm. *Trans. Mach. Learn. Res.*, 2024, 2024.
- 639
- 640 Liheng Ma, Chen Lin, Derek Lim, Adriana Romero-Soriano, Puneet K. Dokania, Mark Coates, Philip
641 Torr, and Ser-Nam Lim. Graph inductive biases in transformers without message passing. In *ICML*,
642 2023.
- 643
- 644 Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably powerful graph
645 networks. In *NeurIPS*, pp. 2153–2164, 2019.
- 646
- 647 Sohir Maskey, Raffaele Paolino, Fabian Jögl, Gitta Kutyniok, and Johannes F. Lutzeyer. Graph repre-
648 sentational learning: When does more expressivity hurt generalization? *CoRR*, abs/2505.11298,
649 2025.
- 650
- 651 Christopher Morris, Martin Ritzert, Matthias Fey, William L. Hamilton, Jan Eric Lenssen, Gaurav
652 Rattan, and Martin Grohe. Weisfeiler and Leman go neural: Higher-order graph neural networks.
653 In *AAAI*, 2019.
- 654
- 655 Christopher Morris, Gaurav Rattan, and Petra Mutzel. Weisfeiler and leman go sparse: Towards
656 scalable higher-order graph embeddings. In *NeurIPS*, 2020.

- 648 Ryan L Murphy, Balasubramaniam Srinivasan, Vinayak Rao, and Bruno Ribeiro. Relational Pooling
649 for Graph Representations. In *ICML*, 2019.
- 650
- 651 C St JA Nash-Williams. Edge-disjoint spanning trees of finite graphs. *Journal of the London*
652 *Mathematical Society*, 1(1):445–450, 1961.
- 653
- 654 Raffaele Paolino, Sohir Maskey, Pascal Welke, and Gitta Kutyniok. Weisfeiler and leman go loopy:
655 A new hierarchy for graph representational learning. In *NeurIPS*, 2024.
- 656
- 657 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
658 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward
659 Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner,
660 Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep
661 learning library. In *NeurIPS*, 2019.
- 662
- 663 Ladislav Rampásek, Mikhail Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Do-
664 minique Beaini. Recipe for a General, Powerful, Scalable Graph Transformer. In *NeurIPS*,
665 2022.
- 666
- 667 Donald J Rose. Triangulated graphs and the elimination process. *Journal of Mathematical Analysis*
668 *and Applications*, 32(3):597–609, 1970.
- 669
- 670 Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The
671 graph neural network model. *IEEE Trans. Neural Networks*, 20(1):61–80, 2009.
- 672
- 673 Kristof T. Schütt, Farhad Arbabzadah, Stefan Chmiela, Klaus R. Müller, and Alexandre Tkatchenko.
674 Quantum-chemical insights from deep tensor neural networks. In *Nature Communications*, 2017.
- 675
- 676 Joshua Southern, Yam Eitan, Guy Bar-Shalom, Michael Bronstein, Haggai Maron, and Fabrizio
677 Frasca. Balancing efficiency and expressiveness: Subgraph gnn with walk-based centrality. In
678 *ICLR*, 2025.
- 679
- 680 Virginia Vassilevska Williams and Yinzhan Xu. Monochromatic triangles, triangle listing and APSP.
681 In *FOCS*, pp. 786–797. IEEE, 2020.
- 682
- 683 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz
684 Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- 685
- 686 Florence H. Vermeire, Yunsie Chung, and William H. Green. Predicting solubility limits of organic
687 solutes for a wide range of solvents and temperatures. In *Journal of the American Chemical*
688 *Society*, 2022.
- 689
- 690 Yanbo Wang and Muhan Zhang. An empirical study of realized gnn expressiveness. In *ICML*, 2024.
- 691
- 692 Boris Weisfeiler and Andrei A. Leman. The reduction of a graph to canonical form and the algebra
693 which appears therein. *NTI, Series 2*, 9:12–16, 1968. In Russian. English translation in *Transl. of*
694 *Math. Monographs*, American Mathematical Society, 2001.
- 695
- 696 Pascal Welke, Maximilian Thiessen, Fabian Jögl, and Thomas Gärtner. Expectation-complete graph
697 representations with homomorphisms. In *ICML*, 2023.
- 698
- 699 Zhenqin Wu, Bharath Ramsundar, Evan N. Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S.
700 Pappu, Karl Leswing, and Vijay Pande. MoleculeNet: a benchmark for molecular machine learning.
701 In *Chemical Science*, 2018.
- 702
- 703 Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. A
704 comprehensive survey on graph neural networks. *arXiv preprint arXiv:1901.00596*, 2019. URL
705 <https://arxiv.org/abs/1901.00596>.
- 706
- 707 Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural
708 networks? In *ICLR*, 2019.

- 702 Kevin Yang, Kyle Swanson, Wengong Jin, Connor Coley, Philipp Eiden, Hua Gao, Angel Guzman-
703 Perez, Timothy Hopper, Brian Kelley, Miriam Mathea, Andrew Palmer, Volker Settels, Tommi
704 Jaakkola, Klavs Jensen, and Regina Barzilay. Analyzing learned molecular representations for
705 property prediction. *Journal of Chemical Information and Modeling*, 59(8):3370–3388, 2019.
- 706 Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and
707 Tie-Yan Liu. Do transformers really perform bad for graph representation?, 2021.
- 708
- 709 Jiaxuan You, Rex Ying, and Jure Leskovec. Position-aware graph neural networks. In *ICML*, 2019.
- 710
- 711 Manzil Zaheer, Guru Guruganesh, Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon,
712 Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. Big bird: Transformers for
713 longer sequences. In *NeurIPS*, 2020.
- 714 Muhan Zhang and Pan Li. Nested graph neural networks. In *NeurIPS*, 2021.
- 715
- 716 Xikun Zhang, Chang Xu, Xinmei Tian, and Dacheng Tao. Graph edge convolutional neural networks
717 for skeleton-based action recognition. *IEEE Trans. Neural Networks Learn. Syst.*, 31(8):3047–3060,
718 2020. doi: 10.1109/TNNLS.2019.2935173. URL [https://doi.org/10.1109/TNNLS.
719 2019.2935173](https://doi.org/10.1109/TNNLS.2019.2935173).
- 720 Lingxiao Zhao, Neil Shah, and Leman Akoglu. A practical, progressively-expressive GNN. In
721 *NeurIPS*, 2022. URL [http://papers.nips.cc/paper_files/paper/2022/hash/
722 dc89a0709f213fd0ac4b1172719b2c38-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/dc89a0709f213fd0ac4b1172719b2c38-Abstract-Conference.html).
- 723
- 724 Junru Zhou, Jiarui Feng, Xiyuan Wang, and Muhan Zhang. Distance-restricted folklore weisfeiler-
725 leman GNNs with provable cycle counting power. In *NeurIPS*, 2023.
- 726
- 727
- 728
- 729
- 730
- 731
- 732
- 733
- 734
- 735
- 736
- 737
- 738
- 739
- 740
- 741
- 742
- 743
- 744
- 745
- 746
- 747
- 748
- 749
- 750
- 751
- 752
- 753
- 754
- 755

756 **Use of LLMs in paper writing.** We used LLMs to polish sentences and improve phrasing.
757

758 A RELATIONSHIP WITH THE CHIBA–NISHIZEKI ALGORITHM

759 We provide a thorough explanation of the algorithm by Chiba & Nishizeki (1985), state some of its
760 properties which are important for our architecture and describe how the algorithm inspired our GNN
761 architecture.
762
763

764 **Description of the Chiba–Nishizeki algorithm.** The algorithm by Chiba & Nishizeki (1985)
765 obtains as input an undirected, unweighted graph $G = (V, E)$ and it returns a list of all triangles in G .
766 Here, we provide a slightly modified version of the algorithm, which is slightly easier to describe but
767 has the same properties.
768

769 Concretely, the algorithm works as follows. It iterates over all edges $(u, v) \in E$ and for each edge
770 (u, v) does the following: Suppose w.l.o.g. that u is the lower-degree endpoint of the endpoint of
771 the edge (u, v) , i.e., assume that $|N(u)| \leq |N(v)|$. Now the algorithm iterates over all neighbors
772 $w \in N(u)$ and checks if $w \in N(v)$; if this is the case then it returns that (u, v, w) is a triangle.

773 Note that the algorithm returns all triangles: Clearly, for any triangle (u, v, w) its edge (u, v) will be
774 considered in the outer loop and we will have that $w \in N(u)$ (thus w is considered in the inner loop)
775 and $w \in N(v)$ (satisfying the if-condition). Thus, the triangle (u, v, w) will be reported.
776

777 Further, the algorithm’s running time is $O(\alpha m)$: Note that for each edge (u, v) we spend time
778 $O(\min\{|N(u)|, |N(v)|\})$ since we only spend time proportional to the neighborhood size of the
779 lower-degree endpoint of (u, v) . Here, we use that the check whether $w \in N(v)$ can be done in time
780 $O(1)$ using hash maps.⁶ Thus, its total running time is given by $\sum_{(u,v) \in E} O(\min\{|N(u)|, |N(v)|\})$
781 and Chiba & Nishizeki (1985, Lemma 2) showed that this quantity is bounded by $O(\alpha m)$.

782 **Properties of the Chiba–Nishizeki algorithm.** The algorithm has several important properties:
783

- 784 1. The algorithm lists all triangles in time $O(\alpha m)$. This implies that graphs with arboricity α
785 contain at most $O(\alpha m)$ triangles. This allows us to bound the number of messages we need
786 to send in our architecture due to triangles by $O(\alpha m)$.
- 787 2. The running time of the Chiba–Nishizeki algorithm is optimal under standard assumptions
788 from the complexity theory community (Kopelowitz et al., 2016; Vassilevska Williams &
789 Xu, 2020). Thus, the preprocessing time of our algorithm cannot be improved if all triangles
790 need to be enumerated.
- 791 3. It is well-known in the algorithms community that the class of graphs with arboricity
792 $O(1)$ contains natural graph families, such as planar graphs, minor-closed families, and
793 preferential attachment graphs, among others. Thus, for all graphs from these families, all
794 triangles can be enumerated in time $O(n)$ and thus in time linear in the size of the input
795 graph. See, e.g., (Chiba & Nishizeki, 1985, Lemma 1) for the planar case.
- 796 4. It holds that $\alpha = O(\sqrt{m})$ (Chiba & Nishizeki, 1985, Lemma 1). As a consequence, the
797 arboricity can be at most a $O(\sqrt{n})$ factor larger than the average degree in the graph. Indeed,
798 this is the case for the simple example consisting of a path with $n - \sqrt{n}$ vertices, which is
799 connected to a clique with \sqrt{n} vertices; this graph has $m = O(n)$ edges and thus average
800 degree $O(1)$ but its arboricity is $O(\sqrt{n})$ due to the clique.
- 801 5. The arboricity is highly related to other graph parameters, such as the *degeneracy* or the
802 *densest subgraph* (Nash-Williams, 1961) and differs from them by at most a factor of 2.
- 803 6. In practice, it is well-known that practical networks have very small arboricities (Eppstein
804 et al., 2013). Indeed, in all 39 datasets considered by Eppstein et al. (2013), the arboricity is
805 at most 201 even though their largest graph has 3.7 million nodes and 16.5 million edges.
806 On 32/39 of their datasets, the arboricity is less than 60. We note that Eppstein et al. (2013)
807 report the degeneracy, which is an upper bound on the arboricity.
808

809 ⁶This is where our version of the algorithm differs from the original. The paper by Chiba & Nishizeki (1985)
does not use hash maps and instead uses a slightly more complicated marking procedure.

Relationship to our architecture. Next, we briefly describe how the Chiba–Nishizeki algorithm inspired our architecture. Recall that the algorithm iterates over all (u, v) ; then it iterates over all $w \in N(u)$ and checks whether $w \in N(v)$ to see if it should report a triangle. In other words, it only reports triangles for vertices w such that $w \in N(u) \cap N(v)$.

When looking at this procedure from the perspective of distributed computing, one can view this as follows: Each edge (u, v) aggregates the neighborhoods $N(u)$ and $N(v)$ and then computes their intersection $N(u) \cap N(v)$ to obtain in which triangles it appears in.

Indeed, this distributed point of view is the motivation for our Equations (6), (7) and (8): Equation (6) aggregates the embeddings of neighbors $N(u)$ of u , Equation (8) aggregates the embeddings of neighbors $N(v)$ of v , and Equation (7) aggregates the embeddings of all edges (u, w) and (v, w) such that $w \in N(u) \cap N(v)$ forms a triangle with u and v .

B OMITTED PROOFS

In this section, we provide missing proofs from the main text.

B.1 PROOF OF PROPOSITION 1

This follows immediately from the discussion in the two paragraphs preceding the statement of Proposition 1.

B.2 PROOF OF THEOREM 2

Part 1: We first show the claim that every pairs of graphs distinguishably by NC-1WL is also distinguished by EB-1WL.

Fix a graph $G = (V, E)$ (we omit G in the notation for colors from now on). In the proof, we use the following terminology. An “ordered edge” is an ordered pair of nodes, connected by an edge (EB-1WL assigns colors to ordered edges). Now, an “ordered triangle” is an ordered triple of nodes where all nodes are connected by an edge.

To simplify the presentation, we will also use the following notation. For example, if u is a node of G , then $(u, *)$ is the set of ordered pairs of nodes connected by an edge, where the first node in the pair is u . Likewise, $(*, u)$ will be a similar set but for pairs where the second node is u .

More generally, if we have a k -tuple where some coordinates are nodes of G , and some coordinates are $*$ (meaning “undefined”), this tuple denotes the set of all ways to replace $*$ ’s by nodes of G such that all pairs of nodes in the tuple are connected by an edge. For example, $(*, *, *)$ means the set of all ordered triangles, and $(u, *, *)$ denotes the set of all ordered triangles that have u as the first coordinate.

Next, if c is a coloring of nodes, and t is a tuple of nodes, we will write $c(t)$ for the tuple of colors of nodes in t , for instance, if $t = (u, v, w)$, then $c(t) = (c(u), c(v), c(w))$. Likewise, if c is a coloring of pairs of nodes, and t is a tuple, then we will write $c(t)$ for the tuple of colors of all pairs of nodes in t . We will use this when t is at most a triple of nodes, where then it is defined as:

$$c(u, v, w) = (c(u, v), c(u, w), c(v, w)).$$

Moreover, we extend this notation to tuples with $*$ ’s. Namely, if T is a tuple of with stars (a set of tuples of nodes without stars), then $c(T) = \{c(t) \mid t \in T\}$.

In this notation, the updates of NC-1WL and EB-1WL can be defined as follows. The color $nc^{(\ell+1)}(u)$ is defined by multisets $nc^{(\ell+1)}(u, *)$ and $nc^{(\ell+1)}(u, *, *)$. In turn, the color $eb^{(\ell+1)}(u, v)$ is defined by $eb^{(\ell)}(u, *)$, $eb^{(\ell)}(v, *)$, and $eb^{(\ell)}(u, v, *)$.

To establish the claim, it is enough to show that $eb(*, *)$ uniquely determines $nc^{(\ell)}(*)$ for every $\ell \geq 0$.

Lemma 6. For every $\ell \geq 0$, and for every ordered edge (u, v) , we have that $\text{eb}(u, v)$ uniquely determines $\text{nc}^{(\ell)}(u)$ and $\text{nc}^{(\ell)}(v)$.

Proof. The proof is by induction on ℓ . For $\ell = 0$, all nodes have the same $\text{nc}^{(0)}$ -color, so there is nothing to prove. Assume now that the statement is proved for ℓ , we establish it for $\ell + 1$. That is, we have to show that $\text{eb}(u, v)$ uniquely determines $\text{nc}^{(\ell+1)}(u)$ and $\text{nc}^{(\ell+1)}(v)$. We only show that it determines $\text{nc}^{(\ell+1)}(u)$. It is enough because by induction in (1–4), it can be shown that $\text{eb}(u, v)$ uniquely determines $\text{eb}(v, u)$.

By definition, $\text{eb}^{(\ell)}(u, v)$ uniquely determines $\text{eb}^{(\ell-1)}(u, *)$ and $\text{eb}^{(\ell-1)}(u, v, *)$. If we make one more step, from $\text{eb}^{(\ell-1)}(u, *)$ we can determine $\text{eb}^{(\ell-2)}(u, *, *)$. Indeed, we use the fact that we can determine $\text{eb}^{(\ell-2)}(u, w, *)$ from $\text{eb}^{(\ell-1)}(u, w)$ for $(u, w) \in (u, *)$.

For a stable EB-1WL coloring eb , we thus get that $\text{eb}(u, v)$ uniquely determines $\text{eb}(u, *)$ and $\text{eb}(u, *, *)$. They, in turn, by the induction hypothesis, uniquely determined $\text{nc}^\ell(u, *)$ and $\text{nc}^\ell(u, *, *)$. Thus, from this information, we get $\text{nc}^{(\ell+1)}(u)$, as required. \square

We now finish the proof of the claim. Assume that we want to know how many times a node color c appears in $\text{nc}^{(\ell)}(*)$. We can assume that $\ell \geq 1$ (we can determine the multiset for $\ell = 0$ from the multiset for $\ell = 1$), then the color c uniquely determines the degree d of a node (which is not 0 by the assumption of the absence of isolated nodes). We go through all ordered edges (u, v) and count how many times we have $\text{nc}^{(\ell)}(u) = c$. We count every node d times, so we have to divide this number by d .

Part 2: We move on to the second part of the theorem: that there are graphs that are distinguished by EB-1WL but not by NC-1WL. The graphs G and H are given in Fig. 4. Both have 16 nodes and are based on the 16-gons. The graph G has also all chords for vertices at distance 2 and at distance 4 in the 16-gon. In turn, the graph H has all chords for vertices at distance 3 and distance 4 in the 16-gon.




Figure 4: Example

All nodes in both graphs have degree 6. Moreover, all nodes in both graphs appear in the same number of triangles, namely, 6. Indeed, every triangle in both graphs is given by an equation $a + b = c$, for some chords, connecting vertices at distances a, b and c , respectively. In the first graph, there are two types of triangles, $1 + 1 = 2$ and $2 + 2 = 4$. Each type, when we rotate it over the 16-gon, covers each vertex exactly 3 times, giving 6 triangles for each node. In the second graph, there are also two types of triangles, $1 + 3 = 4$ and $3 + 1 = 4$ (differing just by their orientation), so the same count gives 6 triangles for every node.

This implies that the NC-1WL gives the same color to all nodes of both graphs in the first iteration, and thus, cannot distinguish these two graphs.

On the other hand, these two graphs can be distinguished by the EB-1WL. Indeed, in H , every edge appears in 2 triangles (every 1-edge or every 3-edge we can complement either from the left or from the right, and every 4-edge we can complete in two ways inside). But, say, for the 4-edge in the first graph, there is just one $2 + 2$ triangle that has it.

B.3 PROOF OF THEOREM 3

We start by showing the following.

Lemma 7. *For each possible color c that the EB-1WL test obtains after t rounds, there is a formula $\phi_c(x, y)$ of quantifier depth t in CFOC³ such that $\text{eb}^t(G, (u, v)) = c$ if and only if $G \models \phi_c(u, v)$, for each graph $G = (V, E)$ and edge $\{u, v\} \in E$.*

Proof. We prove the result by induction on the number of rounds $t \geq 0$. By definition of EB-1WL, before any refinement every edge $\{u, v\}$ receives the same initial color $c = 1$. Accordingly, we can define $\phi_c(x, y) := E(x, y)$.

Suppose we already have, for every color d occurring after t rounds of EB-1WL, a formula $\phi_d(x, y)$ that defines it. Consider now a new color c obtained after $t + 1$ rounds. By the definition of EB-1WL, the color of an edge $\{u, v\}$ at step $t + 1$ is completely determined by the following information from step t :

1. the color d previously assigned to $\{u, v\}$,
2. the multiset \mathcal{E} of colors of edges of the form $\{u, w\}$,
3. the multiset \mathcal{G} of colors of edges of the form $\{v, w\}$, and
4. the multiset \mathcal{F} of pairs of colors (f_1, f_2) corresponding to edges $\{u, w\}$ and $\{v, w\}$ incident with a common neighbor w .

It follows that the new color c can be described in CFOC³ by the formula:

$$\begin{aligned} \phi_c(x, y) := & \phi_d(x, y) \wedge \\ & \bigwedge_{e \in \mathcal{E}} \exists^= n_e z (E(x, z) \wedge \phi_e(x, z)) \wedge \exists^= \sum_{e \in \mathcal{E}} n_e z E(x, z) \\ & \bigwedge_{g \in \mathcal{G}} \exists^= m_g z (E(z, y) \wedge \phi_g(z, y)) \wedge \exists^= \sum_{g \in \mathcal{G}} m_g z E(y, z) \\ & \bigwedge_{(f_1, f_2) \in \mathcal{F}} \exists^= p_{(f_1, f_2)} z (E(x, z) \wedge E(y, z) \wedge \phi_{f_1}(x, z) \wedge \phi_{f_2}(y, z)) \wedge \\ & \exists^= \sum_{(f_1, f_2) \in \mathcal{F}} p_{(f_1, f_2)} z (E(x, z) \wedge E(y, z)). \end{aligned}$$

Here, n_e is the multiplicity of e in \mathcal{E} , m_g the multiplicity of g in \mathcal{G} , and $p_{(f_1, f_2)}$ the multiplicity of (f_1, f_2) in \mathcal{F} . As usual, we write $\exists^= n x \psi$ as shorthand for $\exists^{\geq n} x \psi \wedge \neg \exists^{\geq n+1} x \psi$. Notice that $\phi_c(x, y)$ has quantifier rank $t + 1$. \square

Let $G = (V, E)$ be a graph and \bar{v} a tuple of elements in V . For $t \geq 0$, we write $\text{tp}^t(G, \bar{v})$ for the t -type of (G, \bar{v}) , which is the set of formulas $\phi(\bar{x})$ with quantifier depth t in CFOC³ such that $G \models \phi(\bar{v})$. Also, $\text{eb}^t(G)$ is the multiset formed by all elements of the form $\text{eb}^t(G, (u, v))$, where $\{u, v\} \in E$.

Lemma 8. *Consider graphs $G = (V, E)$ and $G' = (V', E')$ and edges $\{u, v\} \in E$ and $\{u', v'\} \in E'$. Then for each $t \geq 0$, the following hold:*

1. If $\text{eb}^t(G) = \text{eb}^t(G')$ then $\text{tp}^t(G) = \text{tp}^t(G')$.

- 972 2. If $\text{eb}^t(G, (u, v)) = \text{eb}^t(G', (u', v'))$ and $\text{eb}^t(G) = \text{eb}^t(G')$, then $\text{tp}^t(G, u) = \text{tp}^t(G', u')$,
 973 $\text{tp}^t(G, v) = \text{tp}^t(G', v')$, and $\text{tp}^t(G, (u, v)) = \text{tp}^t(G', (u', v'))$.
 974

975 *Proof.* We prove this by induction on $t \geq 0$. The base case $t = 0$ holds trivially. In fact, there
 976 are no sentences in CFOC^3 of quantifier depth 0, and hence $\text{tp}^0(G) = \text{tp}^0(G')$ holds vacuously.
 977 Moreover, any formula $\phi(x)$ of quantifier depth 0 with only one free variable x in CFOC^3 is a Boolean
 978 combination of formulas of the form $E(x, x)$. Since both $G \models \neg E(u, u)$ and $G' \models \neg E(u', u')$, it
 979 is the case that $G \models \phi(u) \Leftrightarrow G' \models \phi(u')$. The same holds for v . Finally, any formula $\phi(x, y)$ of
 980 quantifier depth 0 with only two free variables x and y in CFOC^3 is a Boolean combination of formulas
 981 of the form $E(x, y)$, $E(x, x)$, and $E(y, y)$. Since both $G \models \neg E(u, u)$ and $G' \models \neg E(u', u')$, both
 982 $G \models \neg E(v, v)$ and $G' \models \neg E(v', v')$, and both $G \models E(u, v)$ and $G' \models \neg E(u', v')$, it is the case
 983 that $G \models \phi(u, v) \Leftrightarrow G' \models \phi(u', v')$.

984 Let us consider now the inductive case $t + 1$, for $t \geq 0$. We prove (1) and (2) separately.
 985

- 986 • We first prove (1). Assume that $\text{eb}^{t+1}(G) = \text{eb}^{t+1}(G')$. Consider a sentence ϕ in CFOC^3
 987 of quantifier depth $t + 1$. First, note that the case $t = 0$ is trivial: in this situation, ϕ must be
 988 a Boolean combination of formulas of the form $\exists^{\geq n} x \psi(x)$, where $\psi(x)$ is itself a Boolean
 989 combination of formulas of the form $E(x, x)$. Consequently, either every graph G satisfies
 990 ϕ or none does.

991 Now assume $t > 0$, and let ϕ be a Boolean combination of CFOC^3 formulas of the form
 992 $\exists^{\geq n} x \phi(x)$, where $\phi(x)$ has quantifier depth t . Our assumption $\text{eb}^{t+1}(G) = \text{eb}^{t+1}(G')$
 993 implies

$$994 \{\{\text{eb}^t(G, (u, v)) \mid \{u, v\} \in E\}\} = \{\{\text{eb}^t(G', (u', v')) \mid \{u', v'\} \in E'\}\},$$

995 and therefore, by the induction hypothesis,
 996

$$997 \{\{\text{tp}^t(G, (u, v)) \mid \{u, v\} \in E\}\} = \{\{\text{tp}^t(G', (u', v')) \mid \{u', v'\} \in E'\}\}.$$

998 Moreover, $\text{tp}^t(G, (u, v))$ determines both $\text{tp}^t(G, u)$ and $\text{tp}^t(G, v)$.
 999

1000 To establish that $G \models \phi \iff G' \models \phi$, it suffices to show

$$1001 \{\{\text{tp}^t(G, u) \mid u \in V\}\} = \{\{\text{tp}^t(G', u') \mid u' \in V'\}\}.$$

1002 Take an arbitrary $\tau = \text{tp}^t(G, u)$ for some $u \in V$. Since $t > 0$, every vertex $v \in V$ with
 1003 $\text{tp}^t(G, v) = \tau$ must have the same degree d as u (because they satisfy the same formulas of
 1004 the form $\exists^{\geq n} y E(x, y)$). Define
 1005

$$1006 S(G, \tau) = \{\{\text{tp}^t(G, (u, v)) \mid \{u, v\} \in E \text{ and } \text{tp}^t(G, u) = \tau\}\}.$$

1007 Then the number of vertices $v \in V$ and $v' \in V'$ such that $\text{tp}^t(G, v) = \tau = \text{tp}^t(G', v')$ is
 1008

$$1009 |S(G, \tau)|/d \quad \text{and} \quad |S(G', \tau)|/d,$$

1010 respectively. Because $S(G, \tau) = S(G', \tau)$ by the above, the two counts coincide, which
 1011 completes the proof.
 1012

- 1013 • We now prove (2). Assume that $\text{eb}^{t+1}(G, (u, v)) = \text{eb}^{t+1}(G', (u', v'))$ and $\text{eb}^{t+1}(G) =$
 1014 $\text{eb}^{t+1}(G')$. We only show that $\text{tp}^{t+1}(G, (u, v)) = \text{tp}^{t+1}(G', (u', v'))$, as the remaining
 1015 cases are conceptually analogous. Consider a formula $\phi(x, y)$ in CFOC^3 of quantifier depth
 1016 $t + 1$. Then $\phi(x, y)$ must be of the form $E(x, y) \wedge \alpha(x, y)$, where $\alpha(x, y)$ is a CFOC^3
 1017 formula of quantifier depth $t + 1$. Since both $G \models E(u, v)$ and $G' \models E(u', v')$, we only
 1018 have to show that $G \models \alpha(u, v) \Leftrightarrow G' \models \alpha(u', v')$. By definition, $\alpha(x, y)$ is a Boolean
 1019 combination of: (a) CFOC^3 formulas $\beta(x, y)$ of quantifier depth $t + 1$ with two free variables,
 1020 x and y , (b) CFOC^3 formulas $\gamma(x)$ of quantifier depth $t + 1$ with only one free variable, x ,
 1021 (c) CFOC^3 formulas $\delta(y)$ of quantifier depth $t + 1$ with only one free variable, y , and (d)
 1022 CFOC^3 sentences η of quantifier depth $t + 1$.

1023 Consider first a formula as above of the form $\beta(x, y)$. By construction, $\beta(x, y) = E(x, y) \wedge$
 1024 $\exists^{\geq n} z (E(x, z) \wedge E(y, z) \wedge \beta_1(x, y, z))$, where $\beta_1(x, y, z)$ is a CFOC^3 formula of quantifier
 1025 depth t . But since the logic only allows formulas to mention three variables, β_1 must be
 a Boolean combination of CFOC^3 formulas of quantifier depth t with at most two free

variables in the set $\{x, y, z\}$. By assumption, $\text{eb}^{t+1}(G, (u, v)) = \text{eb}^{t+1}(G', (u', v'))$, and hence both $\text{eb}^t(G, (u, v)) = \text{eb}^t(G', (u', v'))$ and

$$\begin{aligned} & \{(\text{eb}^t(G, (u, w)), \text{eb}^t(G, (v, w))) \mid \{u, w\}, \{v, w\} \in E\} = \\ & \{(\text{eb}^t(G', (u', w')), \text{eb}^t(G', (v', w'))) \mid \{u', w'\}, \{v', w'\} \in E'\}. \end{aligned}$$

Also by assumption, $\text{eb}^{t+1}(G) = \text{eb}^{t+1}(G')$, which implies that $\text{eb}^t(G) = \text{eb}^t(G')$. By induction hypothesis on (1) and (2), we conclude that $\text{tp}^t(G) = \text{tp}^t(G')$, $\text{tp}^t(G, (u, v)) = \text{tp}^t(G', (u', v'))$, $\text{tp}^t(G, u) = \text{tp}^t(G', u')$, $\text{tp}^t(G, v) = \text{tp}^t(G', v')$, and:

$$\begin{aligned} & \{(\text{tp}^t(G, (u, w)), \text{tp}^t(G, (v, w)), \text{tp}^t(G, w)) \mid \{u, w\}, \{v, w\} \in E\} = \\ & \{((\text{tp}^t(G', (u', w')), \text{tp}^t(G', (v', w')), \text{tp}^t(G', w')) \mid \{u', w'\}, \{v', w'\} \in E'\}. \end{aligned}$$

This implies that

$$\begin{aligned} & |\{w \in N(u) \cap N(v) \mid G \models \beta_1(u, v, w)\}| = \\ & |\{w' \in N(u') \cap N(v') \mid G' \models \beta_1(u', v', w')\}|, \end{aligned}$$

which means that $G \models \beta(u, v) \Leftrightarrow G' \models \beta(u', v')$.

Consider second a formula as above of the form $\gamma(x)$. By construction, $\gamma(x) = \exists^{\geq n} z (E(x, z) \wedge \gamma_1(x, z))$, where $\gamma_1(x, z)$ is a CFOC³ formula of quantifier depth t . By assumption, $\text{eb}^{t+1}(G, (u, v)) = \text{eb}^{t+1}(G', (u', v'))$, and hence

$$\{(\text{eb}^t(G, (u, w)) \mid \{u, w\} \in E\} = \{(\text{eb}^t(G', (u', w')) \mid \{u', w'\} \in E'\}.$$

By induction hypothesis, this implies that:

$$\{(\text{tp}^t(G, (u, w)) \mid \{u, w\} \in E\} = \{(\text{tp}^t(G', (u', w')) \mid \{u', w'\} \in E'\}.$$

By focusing on those types that contain $\gamma_1(x, z)$, we obtain:

$$|\{w \in N(u) \mid G \models \gamma_1(u, w)\}| = |\{w' \in N(u') \mid G' \models \gamma_1(u', w')\}|,$$

which means that $G \models \gamma(u) \Leftrightarrow G' \models \gamma(u')$.

The formulas of the form $\delta(y)$ are handled analogously to the previous case.

Finally, consider a sentence as above of the form η . Since $\text{eb}^{t+1}(G) = \text{eb}^{t+1}(G')$, we have by part (1) that $G \models \eta \Leftrightarrow G' \models \eta$.

This finishes the proof of the lemma. \square

We now prove Theorem 3. Assume first that $\text{eb}(G) = \text{eb}(G')$, for graphs G and G' . In particular, then, $\text{eb}^t(G) = \text{eb}^t(G')$ for every $t \geq 0$. Take an arbitrary sentence ϕ of CFOC³, and assume that its quantifier depth is t . From Lemma 8, we conclude that $\text{tp}^t(G) = \text{tp}^t(G')$, and hence $G \models \phi \Leftrightarrow G' \models \phi$.

Assume, on the contrary, that $G \models \phi \Leftrightarrow G' \models \phi'$, for every CFOC³ sentence ϕ . Hence, $\text{tp}^t(G) = \text{tp}^t(G')$ for every $t \geq 0$. To prove $\text{eb}(G) = \text{eb}(G')$ it suffices to show $\text{eb}^t(G) = \text{eb}^t(G')$ for each $t \geq 0$. For a node $v \in V$, define $\text{eb}^t(v) = \{(\text{eb}^t(v, w) \mid \{v, w\} \in E)\}$. Notice, then, that $\text{eb}^t(G)$ is completely determined by the multiset $\Gamma(G) = \{(\text{eb}^t(v) \mid v \in V)\}$. For each element $\tau \in \Gamma(G)$, we write \mathcal{C}_τ for the multiset of colors computed by EB-IWL after t steps that belong to τ .

Define the following formula from CFOC³:

$$\phi_G := \left(\bigwedge_{\tau \in \Gamma(G)} \exists^{\ell_\tau} x \phi_\tau(x) \right) \wedge \exists^{\sum_{\tau \in \Gamma(G)} \ell_\tau} x (x = x),$$

where ℓ_τ is the multiplicity of τ in $\Gamma(G)$ and $\phi_\tau(x)$ is defined as follows:

$$\phi_\tau(x) = \left(\bigwedge_{c \in \mathcal{C}_\tau} \exists^{\ell_c} y (E(x, y) \wedge \phi_c(x, y)) \right) \wedge \exists^{\sum_{c \in \mathcal{C}_\tau} \ell_c} y E(x, y),$$

1080 where q_τ is the cardinality of τ in $\Gamma(G)$ and $\phi_c(x, y)$ is the CFOC³ formula from Lemma 7. Notice
 1081 that ϕ_G has quantifier depth bounded by $t + 2$.

1082 It is easy to see that $G' \models \phi_G \Leftrightarrow \Gamma(G) = \Gamma(G')$. Since $G \models \phi_G$ and $\text{tp}^{t+2}(G) = \text{tp}^{t+2}(G')$, we
 1083 conclude that $G' \models \phi_G$, and hence $\Gamma(G) = \Gamma(G')$. Since $\Gamma(G)$ determines $\text{eb}^t(G)$, we conclude
 1084 that $\text{eb}^t(G) = \text{eb}^t(G')$.
 1085

1086 B.4 PROOF OF THEOREM 4

1087 It is well-known that chordal graphs admit a *perfect elimination order*: an ordering of its nodes such
 1088 that every node v and its neighbors that go before v in the order form a clique (Rose, 1970). Graphs
 1089 of tree-width 2 cannot have cliques larger than a triangle; this means that for any chordal graph H of
 1090 tree-width 2 there exists an ordering v_1, \dots, v_m of its nodes such that for any $k \in \{1, \dots, m\}$, one
 1091 of the following holds:
 1092

- 1093 • (a) v_k is not connected to any node out of v_1, \dots, v_{k-1} ;
- 1094 • (b) v_k is connected to exactly one node out of v_1, \dots, v_{k-1} ;
- 1095 • (c) v_k is connected to exactly two nodes $v_i, v_j \in \{v_1, \dots, v_{k-1}\}$, and v_i and v_j are also
 1096 connected.
 1097

1098 Let eb be the stable EB-1WL coloring of G . Let us show that the multiset $\text{eb}(G)$ of eb-labels
 1099 of ordered edges uniquely determines the number of homomorphisms from H to G , for any tree-
 1100 width 2 chordal graph H . This means that if two graphs G_1 and G_2 have a different number of
 1101 homomorphisms from some graph H like that, they will be distinguished by the EB-1WL on the
 1102 stage where colorings of both graphs stabilize.
 1103

1104 For a node u , define:

$$1105 \text{eb}(u) = \{\{\text{eb}(u, w) \mid w \in N(u)\}\}.$$

1106 Note that $\text{eb}(u, v)$, as it is stable, uniquely determines induced labels of u and v through (2) and
 1107 (4). Moreover, $\text{eb}(G)$ uniquely determines the multiset $\text{eb}(V) = \{\{\text{eb}(u) \mid u \in V\}\}$. Namely, we
 1108 go through all $\text{eb}(u, v)$, compute $\text{eb}(u)$ from it, which in turn determines the degree of u . We then
 1109 divide the number of occurrences of $\text{eb}(u)$ by the degree.
 1110

1111 Consider any labeling of edges of H by labels from $\text{eb}(G)$, and of its nodes by labels from $\text{eb}(V)$.
 1112 We show that for any such labeling, the number of homomorphisms from H to G that “preserve”
 1113 this labeling is determined just by the multiset $\text{eb}(G)$. The total number of homomorphisms is
 1114 hence also determined by $\text{eb}(G)$, since we can go through all possible labelings of H and sum up
 1115 homomorphisms for all of them.

1116 Here, “preserve” formally means that (a) a node v_j with a label ℓ goes into a node in G that has this
 1117 eb-label (b) if v_i, v_j is an edge of H (where $i < j$ are indices of these nodes in the perfect elimination
 1118 order), and if v_i and v_j go to some nodes u_i, u_j in G , respectively, then the label of the edge v_i, v_j in
 1119 H has to be equal to $\text{eb}(u_i, u_j)$.

1120 We show that we can compute the number of homomorphisms that preserve a given labeling of H ,
 1121 just knowing $\text{eb}(G)$, by first computing how many ways we can define the image of v_1 , then the
 1122 image of v_2 , then of v_3 , and so on.

1123 As for v_1 , it is assigned a label ℓ in the labeling; we know the multiset of eb-labels of the nodes of G ,
 1124 which determines the number of ways we can define the image of v_1 (this is the number of nodes of
 1125 G that have label ℓ).
 1126

1127 Now, assume that we have defined images of v_1, \dots, v_{k-1} . Now, it’s v_k ’s turn. Firstly, it is possible
 1128 that v_k is not connected to any node among v_1, \dots, v_{k-1} . Then we can freely map v_k to any node of
 1129 G that has the same label ℓ as assigned to v_k in the coloring. We just have to multiply the current
 1130 number of homomorphisms by the number of nodes in G with this label ℓ .

1131 If v_k is connected to a single node v_i , $i < k$, and v_i is already mapped to some node u_i , then we have
 1132 to map v_k to some node u_k that is connected to u_i and such that the $\text{eb}(u_i, u_k)$ coincides with the
 1133 color of the edge v_i, v_k in H (this color also determines the label of u_k which has to be consistent with
 the label that v_k has in the labeling of H , otherwise the number of homomorphisms is just 0). The

number of ways to choose such u_k is thus determined by $\{\{\text{eb}(u_i, w) \mid w \in N(u_i)\}\} = \text{eb}(u_i)$, which in turn equals the label of v_i in the labeling of H . We multiply the current number of homomorphisms by the number of occurrences of the label of the edge v_i, v_k in the label of v_i .

The same argument holds for the third case when v_k is connected to previous nodes v_i, v_j that are connected by an edge. The number of ways to choose the image of v_k is the number of occurrences of the pair (ℓ_{ik}, ℓ_{jk}) in (3) in the label of the edge ℓ_{ij} , where ℓ_{ik}, ℓ_{jk} , and ℓ_{ij} are labels of edges v_i, v_k, v_j, v_k , and v_i, v_j in H , respectively.

B.5 PROOF OF THEOREM 5

The theorem is deduced from the following lemma.

Lemma 9. *Let $f_1, \dots, f_n \in \mathbb{R}^d$ be n distinct vectors. Then there exists a matrix $A \in \mathbb{R}^{n \times d}$ and a vector $b \in \mathbb{R}^n$ such that the vectors*

$$g_1 = \text{ReLU}(Af_1 + b), \dots, g_n = \text{ReLU}(Af_n + b)$$

are linearly independent.

Indeed, consider any two graphs G and H , distinguishable by the EB-1WL test. We construct an EB GNN \mathcal{T} that distinguishes G and H .

Consider the disjoint union of G and H . Let m be the number of ordered edges of this union, and t be the number of ordered triangles. The dimension of \mathcal{T} will be

$$d = 1 + 2m + t.$$

We show that there exists a choice of parameter matrices such that, for any i , a) EB-1WL labels after i iterations are in a one-to-one correspondence with feature vectors $f^{(i)}(u, v)$; b) all coordinates of $f^{(i)}$, except the first one, are 0s.

For $i = 0$, this holds because all edges initially have the same EB-1WL label, and because of (5).

Assume that it holds after $i - 1$ iterations. We use Lemma 9 in (6–8) to map distinct feature vectors $f^{(i-1)}(u, v)$ (there are at most m of them) or their pairs as in (7) (there are at most t such pairs) to linearly independent vectors. This ensures that we obtain different sums for different multisets of terms in (6–8). We can use 3 blocks of m, m and t disjoint coordinates that are “free” in vectors $f^{(i-1)}(u, v)$ so that the sum in (9) uniquely determines the whole 4-tuple in the definition of the updated EB-1WL feature.

We can then define a feed-forward network to injectively map vectors $g^{(i)}(u, v)$ to vectors $f^{(i)}(u, v)$ that have all coordinates, except the first one, equal to 0. There exists a vector $w = (w_1, \dots, w_d) \in \mathbb{R}^d$ such that $\langle w, g^{(i)}(u, v) \rangle \neq \langle w, g^{(i)}(u', v') \rangle$ whenever $g^{(i)}(u, v) \neq g^{(i)}(u', v')$ (for each of the finitely many pairs of distinct $g^{(i)}$ -vectors, the set of w for which we have equality is a hyperplane in \mathbb{R}^d). Hence, it is enough to realize the following linear transformation by a FFN:

$$g^{(i)}(u, v) \mapsto \begin{pmatrix} w_1 & w_2 & \dots & w_d \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{pmatrix} g^{(i)}(u, v) = Wg^{(i)}(u, v).$$

This can be achieved by the following FFN:

$$x \mapsto W \text{ReLU}(x + u_i) - Wu_i,$$

where $u_i \in \mathbb{R}^d$ is an arbitrary vector such that $u_i \geq g^{(i)}(u, v)$ coordinate-wise for every ordered edge (u, v) . Indeed, then we obtain:

$$x \mapsto W \text{ReLU}(x + u_i) - Wu_i = W(x + u_i) - Wu_i = Wx$$

for any x of the form $x = g^{(i)}(u, v)$, as required.

Now, assume that we are at an iteration when G and H have different multisets of EB-1WL labels. We now need to do one more iteration that map distinct feature vectors into linearly independent

vectors so that G and H will have different final representations in (11). We set all matrices and bias vectors in (6–8) to 0 so that $g^{(i)}(u, v) = f^{(i-1)}(u, v)$. We then use Lemma 9 again to construct a matrix U and a vector u such that the following transformation:

$$f^{(i-1)}(u, v) \mapsto \text{ReLU}(U f^{(i-1)}(u, v) + v) \quad (12)$$

maps distinct vectors into linearly independent ones. There are at most m distinct feature vectors, which means we can use coordinates that are 0 in $g^{(i)}(u, v)$ (by the invariant b)) for the image of (12) so it does not interfere with the first coordinate of $g^{(i)}(u, v)$ in (10). By setting $V = Id, v = 0$, we obtain a FFN that realizes this transformation.

Proof of Lemma 9. Since f_1, \dots, f_n are distinct, there exists $w \in \mathbb{R}^d$ such that $\langle f_1, w \rangle, \dots, \langle f_n, w \rangle$ are distinct (because the set of w such that $\langle f_i, w \rangle = \langle f_j, w \rangle$ for some $i \neq j$ is a union of finitely many hyperplanes, not covering the whole \mathbb{R}^d). Without loss of generality,

$$\langle f_1, w \rangle < \langle f_2, w \rangle < \dots < \langle f_n, w \rangle.$$

For $i = 1, \dots, n$, let γ_i be some number between $\langle f_{i-1}, w \rangle$ and $\langle f_i, w \rangle$ (for $i = 1$, this is some number smaller than $\langle f_1, w \rangle$). Define

$$A = \begin{pmatrix} w \\ \vdots \\ w \end{pmatrix}, \quad b = \begin{pmatrix} -\gamma_1 \\ -\gamma_2 \\ \vdots \\ -\gamma_n \end{pmatrix}.$$

Observe that in the vector

$$A f_i + b = \begin{pmatrix} \langle f_i, w \rangle - \gamma_1 \\ \langle f_i, w \rangle - \gamma_2 \\ \vdots \\ \langle f_i, w \rangle - \gamma_n \end{pmatrix}$$

the first i coordinates are strictly positive, and the other coordinates are strictly negative. Hence, the vector $g_i = \text{ReLU}(A f_i + b)$ is a vector where the first i coordinates are strictly positive, and the rest are 0s. Therefore, g_1, \dots, g_n are linearly independent. \square

This finishes the proof of the theorem.

C MORE DETAILS ON EXPERIMENTS

We provide additional information on our experimental procedure and more detailed results.

Model. We have defined graph as purely existing of nodes and edges $G = (V, E)$. However, real-world datasets often use node features and edge features to encode additional information in the graph for all. For every directed edge $(u, v) \in E$, we incorporate the node features of X_u, X_v of u, v and the edge features $W_{(u,v)}$ of (u, v) into our model by initializing the edge embedding $\mathcal{T}^{(0)}(G, (u, v))$ with them. We use three different embedding encoders ENC to map both the node features and the edge features to the embedding dimension of our GNN

$$\mathcal{T}^{(0)}(G, (u, v)) = \text{ENC}_{\text{left}}(X_u) + \text{ENC}_{\text{right}}(X_v) + \text{ENC}_{\text{edge}}(W_{(u,v)}).$$

After obtaining the initial edge embedding, we perform multiple iterations of edge based message passing. After final iteration t , we pool edge embeddings into the shape required by the task. For graph-level tasks, we experiment with three methods. We compute graph-level embeddings by either summing

$$\mathcal{T}_{\text{SUM}}(G) = \sum_{(u,v):\{u,v\} \in E} \mathcal{T}^t(G, (u, v)),$$

1242 computing the mean, or computing the mean scaled by the number of nodes (mimicking sum pooling
1243 in MPNNs)

$$1244 \mathcal{T}_{\text{MEAN}}(G) = \frac{1}{|E|} \mathcal{T}_{\text{SUM}}(G), \quad \mathcal{T}_{\text{NODESUM}}(G) = \frac{|V|}{|E|} \mathcal{T}_{\text{SUM}}(G).$$

1246 For edge-level tasks, note that we compute embeddings of directed edges whereas tasks we worked on
1247 were on *undirected* regression edges. we experiment with two methods of performing undirected edge
1248 predictions. First, we simply sum the embedding of the two directed edges and perform a prediction
1249 on this combined embedding. Second, we make a prediction for both directions and combine these
1250 predictions by computing the mean.
1251

1252 All our models were implemented in PyTorch Geometric (Paszke et al., 2019; Fey & Lenssen, 2019).
1253 All our models were trained on servers with one NVIDIA GeForce RTX 3080 GPU (10 GB VRAM)
1254 and 64 GB of RAM. When training a model we evaluate its performance after every epoch on the
1255 validation and test set. This performance is measured in the metric that is most commonly used on
1256 that dataset. After training, we report the validation and test performance in the epoch with the best
1257 validation performance. For real-life datasets, we perform hyperparameter tuning where we pick
1258 the hyperparameter combination based on the best validation performance. We train a model with
1259 this hyperparameter combination multiple times on different seeds reporting the mean and standard
1260 deviation of the test metric.

1261 In general, all our models are trained with a Cosine learning rate scheduler and a learning rate of
1262 0.001 (except on BREC where we use the procedure provided by Wang & Zhang (2024)). The
1263 final prediction, is made by a two-layer MLP. On all real-life datasets EB-GNN uses both skip
1264 connections and feed-forward layers, but not on the synthetic datasets CSL and BREC. Below, we
1265 discuss the different setup we used for each dataset. For more details, please consider our code at
1266 <https://anonymous.4open.science/r/EB-GNN-CB41/>.

1267 **CSL.** We train all models for 1000 epochs with a Cosine learning rate scheduler. All GNNs have 5
1268 layers and an embedding dimension of 64.

1269
1270 **BREC.** We train and evaluate our model with the procedure provided by Wang & Zhang (2024).
1271 Our EB-GNN model has 10 layers, an embedding dimension of 16 and uses some pooling. We have
1272 observed that using the output of Eq. 9 instead of Eq. 10 as edge embeddings leads to better results
1273 on the extension graphs. We believe that this is due to numerical issues caused by the large number
1274 of layers and small embedding dimension (which is necessary to fit the model into GPU memory).
1275 Our MPNN models also have 10 layers but uses an embedding dimension of 64.

1276
1277 **QMD.** We train separate models for each of the 8 different tasks. We tune the hyperparameters of
1278 EB-GNN for each task based on the grid in Table 5. We train for 500 epochs with a batch size of
1279 1024 and evaluate on 10 different seeds.

1280
1281 **QM9.** We train separate models for each of the 11 different tasks. Initially, we repeated the same
1282 procedure as for QMD. However, training with a significantly larger batch size than models in literature
1283 might harm our performance. Thus, after the initial hyperparameter sweep (Table 5) we used the
1284 best hyperparameters and additionally tuned the batch size together with another pooling operation
1285 (Tab. 6). We evaluate the best hyperparameter combination on 10 different seeds.

1286 Similar to Zhou et al. (2023) and Paolino et al. (2024), we perform a speed evaluation on QM9. For
1287 this, we train our model on the training set with batch size 64 and measure the time it takes to train for
1288 a single epoch. Additionally, we also track the pre-processing time on the entire dataset. The results
1289 can be seen in Tab. 7. There are two issues with this type of evaluation. First, we (as well as previous
1290 work) compare runtime of models trained on different hardware. This is best noticed by the fact that
1291 our expressive EB-GNN is faster than the MPNN in Tab. 7. In our case, this is less of a problem
1292 because compared to 5 ℓ -GIN (the other best model on QM9), our GPU is significantly weaker (RTX
1293 3080 vs RTX 3090 Ti/RTX A6000). Second, we believe that previous works included the initial
1294 processing time for the dataset in pre-processing. This includes time spent to generate the initial
1295 graphs which is needed for all models but can vary across hardware. We remedy this by reporting
both the time spent for our pre-processing (70 seconds) as well as the time spent on preparing the
initial graphs (30 seconds).

MalNet-Tiny. As the graphs in MalNet-Tiny are very large, we did not perform any hyperparameter tuning. Our model uses mean pooling, has 5 message passing layers and an embedding dimension of 64. It is trained for 500 epochs with a batch size of 16 and evaluated on 5 different seeds.

Table 5: Hyperparameter grid used on QMD and QM9.

Hyperparameter	Values
Embedding dimension	128, 256
Dropout rate	0, 0.2, 0.5
Number of message passing layers	3, 4, 5
Pooling (graph-level tasks)	$\mathcal{T}_{\text{SUM}}, \mathcal{T}_{\text{MEAN}}$
Pooling (edge-level tasks)	sum directed embeddings \rightarrow make undirected prediction, make directed predictions \rightarrow sum into undirected prediction

Table 6: Smaller hyperparameter grid used on QM9 after hyperparameter tuning with Tab. 5.

Hyperparameter	Values
Embedding dimension	Best from Tab. 5
Dropout rate	Best from Tab. 5
Number of message passing layers	Best from Tab. 5
Batch size	64, 1024
Pooling	$\mathcal{T}_{\text{NODESUM}}$, Best from Tab. 5

Table 7: Empirical time complexity for QM9 dataset; results from Zhou et al. (2023) and Paolino et al. (2024).

Model	Preprocessing [sec]	Training [sec/epoch]
MPNN	64	45.3
NestedGNN	2 354	107.8
I2GNN	5 287	209.9
2-DRFWL	430	141.9
5- ℓ GIN	444	130.6
EB-GNN (ours)	100	31

D RELATIONSHIP TO δ - k -LWL

Morris et al. (2020) introduced δ - k -WL and its local variant δ - k -LWL as a variant of the k -WL test that is more efficient on sparse graphs, matching our motivation for the introduction of EB-1WL. The relationship between the two formalisms, especially for the case $k = 2$, is interesting. The formalisms look superficially similar, while being conceptually quite far apart. We therefore elaborate on these differences explicitly here.

In δ - k -LWL, colors $c^{(\ell)}$ are assigned to k -tuples of vertices. Initially, $c(\mathbf{v})$ is colored by the isomorphism type of $G[\mathbf{v}]$ (the induced subgraph of G by vertices in \mathbf{v}). We refer to position i of \mathbf{v} as v_i . Moreover, let $\theta_i(\mathbf{v}, x)$ denote the tuple obtained by replacing the i component of \mathbf{v} with x . The colors are then updated iteratively according to

$$c^{(\ell+1)}(\mathbf{v}) = (c^{(\ell)}(\mathbf{v}, \tag{13}$$


$$\{\{c^{(\ell)}(\theta_1(\mathbf{v}, x)) \mid x \in N(v_1)\}\}, \tag{14}$$

$$\dots \tag{15}$$

$$\{\{c^{(\ell)}(\theta_k(\mathbf{v}, x)) \mid x \in N(v_k)\}\}) \tag{16}$$

On the surface the core update part, gathering the multisets over neighbors of each component, seems similar to parts (2) and (4) in our update function for eb. However, the mathematically literate reader will notice some key differences. In δ - l -LWL the update considers are vectors where the i -th component is replaced by neighbors of the i -th component, i.e., we replace v_i with elements in $N(v_i)$. In EB-1WL we, in a sense, do the opposite. We take the edges incident to each v_i , that is we form the tuples (v_i, x) for $x \in N(v_i)$. Instead of replacing v_i by its neighbors, we look at the tuples formed by v_i with its neighbors.

We illustrate this with a simple example graph



The δ -2-WL update for $c((a, b))$ would update based on the pair of multisets $\{\{C(a, x) \mid x \in N(b)\}\}$ and $\{C(x, c) \mid x \in N(a)\}$, that is both times $\{\{C(a, c)\}\}$. That is, we replace each component with all its neighbors, i.e., a classic k -WL style update but limited to local neighborhoods.

The seemingly similar part of our update (that is only (2) and (4), ignoring (3)) would update based on $\{C(a, x) \mid x \in N(a)\}$ (2) and $\{C(b, x) \mid x \in N(b)\}$. So $\{C(a, b)\}$ and $\{C(b, c)\}$, i.e., the update is based on the edges incident to edge (a, b) (and the edge itself).

Another noteworthy difference is in the computational complexity of a layer update. An update for δ -2-WL requires $O(n^2d)$ time, where d is the degree of the graph. As each n^2 pairs has to access $k \cdot d$ colors of neighbors. Intuitively, this bound is reasonably tight as neither part – the number of tuples, or the access of $k \cdot d$ tuples in the update – can be avoided. In contrast, our bound of $O(\alpha m)$ promises significantly better performance, especially in sparse graphs as they have $m \ll n^2$ and $\alpha \ll d$.

E EB-1WL VS 1WL WITH TRIANGLE INFORMATION

The important role of triangles in EB-1WL naturally motivates a comparison with MPGNNs that have triangle counts injected.

Here we provide an example demonstrating that EB-1WL is strictly more expressive than 1WL with triangle counts added on node and edge level. Figure 5 shows two graphs from the BREC dataset that are indistinguishable by 1WL with such triangle counts but distinguishable by EB-1WL. Figure 6 gives example code to verify the indistinguishability by 1WL. Moreover, the graphs exhibit a different number of homomorphisms from the 4-cycle with a chord and by Theorem 4 is therefore distinguished by EB-1WL.




Figure 5: The first graph pair in the Extension graphs of the BREC dataset Wang & Zhang (2024).

```

1427 1 from BRECDataset import BRECDataset
1428 2 import networkx as nx
1429 3
1430 4 def edge_idx_to_nx(ei):
1431 5     g = nx.Graph()
1432 6     for a,b in zip(ei[0], ei[1]):
1433 7         g.add_edge(int(a), int(b))
1434 8     return g
1435 9
1436 10 d = BRECDataset("BREC_data_all")
1437 11
1438 12 g0 = edge_idx_to_nx(d[160*32*2 : (160+1)*32*2][0].edge_index)
1439 13 g1 = edge_idx_to_nx(d[160*32*2 : (160+1)*32*2][1].edge_index)
1440 14
1441 15 def edge_tri(g, multiset=False):
1442 16     vt = nx.triangles(g)
1443 17     et = {m: vt[m[0]]+vt[m[1]] for m in g.edges()}
1444 18     if multiset:
1445 19         return list(sorted(et.values()))
1446 20     else:
1447 21         return et
1448 22
1449 23 nx.set_node_attributes(g0, nx.triangles(g0), 'c3')
1450 24 nx.set_node_attributes(g1, nx.triangles(g1), 'c3')
1451 25
1452 26 nx.set_edge_attributes(g0, edge_tri(g0), 'ec3')
1453 27 nx.set_edge_attributes(g1, edge_tri(g1), 'ec3')
1454 28
1455 29 h1 = nx.weisfeiler_lehman_graph_hash(g0, node_attr='c3', edge_attr='ec3')
1456 30 h2 = nx.weisfeiler_lehman_graph_hash(g1, node_attr='c3', edge_attr='ec3')
1457 31
1458 32 print(h1, h2, h1 == h2)

```

Figure 6: Example code to verify the indistinguishability of G_0 and G_1 by 1WL with triangle counts on edges and vertices.