Graph Homomorphism Distortion: A STRUCTURAL (PSEUDO)-METRIC FOR ATTRIBUTED GRAPHS

Anonymous authors
Paper under double-blind review

ABSTRACT

A large driver of the complexity of graph learning is the interplay between *structure* and *features*. When analyzing the expressivity of graph neural networks, however, existing approaches ignore features in favor of structure, making it nigh-impossible to assess to what extent two graphs with close features should be considered similar. We address this by developing a new (pseudo-)metric based on graph homomorphisms. Inspired by concepts from metric geometry, our *graph homomorphism distortion* measures the minimal worst-case distortion that node features of one graph are subjected to when mapping one graph to another. We demonstrate the utility of our (pseudo-)metric by showing that (i) it can be efficiently approximated while retaining distinction capabilities, (ii) it provides a granular approach to graph distinction, and (iii) it serves as a structural encoding, extending the concept of homomorphism-based descriptors to attributed graphs, thus improving the predictive capabilities of graph neural networks.

1 INTRODUCTION

A pivotal concept in graph learning involves assessing the *expressivity* of a model, intended as its ability to distinguish between graphs that are considered to be "different." Typically, "different" is taken to mean "non-isomorphic" under the Weisfeiler–Leman (WL) hierarchy of graph isomorphism tests, which remain the *de facto* benchmark for expressivity. Prior work has linked the WL hierarchy to message-passing neural networks (Morris et al., 2019; 2023; Xu et al., 2019); however, it is fundamentally limited in that it yields a *binary* quantity and does not account for an underlying measure of similarity, i.e., it cannot capture "small differences" between graphs. This is due to graph isomorphism being a brittle property, meaning that even a single edge modification can break it. Moreover, since WL is a *combinatorial* measure, it only captures graph structure while ignoring any features. The case where features are weighted by importance by using a probability measure has openned the door for more principled considerations, such as using a Gromov-Haussdorf distance as similarity (Vayer et al., 2020). As a notable extension of WL to geometric graphs, Joshi et al. (2023) develop a model of expressivity based on joint equivalence with respect to graph isomorphism (*structure*) and Euclidean isometry (*features*). While this leads to a principled measure for vertex features with an inherent Euclidean geometry, it still does not account for approximate (geometric) similarity. Despite alternatives to strict structural views, like spectral-based methods (Isufi et al., 2024), being studied, the current ecosystem is still skewed towards structure (Guo et al., 2025) and *should* move away from previous restrictive views of both (Vasileiou et al., 2026).

Graph homomorphisms: A *melange* of features and structure. Our work proposes a more general perspective that compares vertex features while incorporating the graph structure as a restriction. To this end, we make use of *graph homomorphisms*. For structural comparisons, graph homomorphism *counts* have been shown to distinguish any pair of non-isomorphic graphs (Lovász, 1967) while characterizing graph properties (Lovász, 2012). Thus, homomorphism counts have been successfully used in machine learning (Bao et al., 2025; Curticapean et al., 2017; Jin et al., 2024; Wolf et al., 2023), but they still lead to an inherently binary measure of similarity. To overcome this, we introduce the concept of *graph homomorphism distortion*, a novel notion of similarity for vertex-attributed graphs inspired by the *Hausdorff distance*, which measures the distance between two subsets of a space via per-point *distortions*.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Through the exploration of our notion of similarity, our paper makes the following key **contributions**:

- We introduce the graph homomorphism distortion and prove its theoretical properties, including its ability to distinguish between non-isomorphic graphs, induced subgraphs,and strongly symmetrical graphs.
- We devise an approximate method to compute our pseudo-metric while retaining strong discriminative capabilities.
- We prove that the graph homomorphism distortion *complements* existing expressivity measures like 1-WL by providing a more granular measure of graph similarity.
- We develop a *structural encoding* based on the graph homomorphism distortion and demonstrate its utility in classifying non-isomorphic graphs and serving as a novel *inductive bias* for graph-level tasks.

2 Background

Our work focuses on *simple*,[1] *connected*, *undirected* graphs. An undirected graph is tuple denoted $G = (V, E)$, where $v \in V$ are the vertices and $\{u, v\} \in E$ the edges.When we deal with multiple graphs at the same time, we denote the vertex set of a graph G with $V(G)$ and the edge set with $E(G)$. The number of vertices (*order*) of a graph is denoted $|G|$. Furthermore, we say that $G' = (V', E')$ is a *subgraph* of $G = (V, E)$, denoted $G' \subseteq G$, if $V' \subseteq V$ and $E' \subseteq E$. If G' contains all edges $\{u', v'\} \in E$ for $u', v' \in V'$ then G' is an *induced subgraph* of G, denoted by $G' \coloneqq G[V']$. A *rooted graph* is a tuple (G, v), where G is the graph and v is a fixed node in the graph. The classes of finite graphs with n nodes and finite rooted graphs are denoted as $\mathcal{G}_n$ and $\mathcal{G}_\bullet$ respectively.

A *(vertex)-attributed graph*[2] is a triplet $G = (V, E, f_G)$, where $f_G \colon V \to \mathcal{X}$ assigns to each vertex $v \in V$ a value $f_G(v)$ in some normed space $(\mathcal{X}, \|\cdot\|)$. Throughout this paper, we assume that all attribute functions map into the *same* co-domain. This assumption is commonly satisfied.

For two graphs $G = (V, E)$ and $H = (V', E')$, a *homomorphism* is a function $\varphi \colon V \to V'$ such that if $\{u, v\} \in E$, then $\{\varphi(u), \varphi(v)\} \in E'$. A graph G is called a *core* if for every proper subgraph G' there is no homomorphism from G to G'. We denote the set of all homomorphisms from G to H by $\mathrm{Hom}(G, H)$ and its cardinality by $\hom(G, H)$. The function φ is an *isomorphism* if it additionally satisfies that its inverse φ^{-1} exists and is a homomorphism. We write $G \simeq H$ to indicate that two graphs are isomorphic. Additionally, two rooted graphs (G, v) and (H, w) are isomorphic if there exists an isomorphism φ such that $\varphi(v) = w$. A *graph encoding* is a function $\Gamma \colon \mathcal{G} \to \mathcal{X}$ that acts on a set of graphs $\mathcal{G}$ and embeds them into some normed space $\mathcal{X}$. The function Γ is an *invariant graph encoding* if it is *permutation-invariant*, i.e., for $G, H \in \mathcal{G}$ with $G \simeq H$, we have $\Gamma(G) = \Gamma(H)$. Additionally, if for $G \not\simeq H$ we have $\Gamma(G) \neq \Gamma(H)$, we call a graph encoding *complete*.

3 The Graph Homomorphism Distortion

The main thread running through this work is on measuring (dis)-similarity between graphs. The only constraint we establish is that the graph theoretic notion of equivalence, isomorphism, shall be respected, i.e., our distance assigns 0 to two isomorphic graphs. Our source of inspiration is the *Hausdorff metric*. For two subsets of a metric space $A, B \subseteq \mathcal{X}$ with metric $\mathrm{d}_\mathcal{X}$, it measures the distance between A and B as

$$\mathrm{d_H}(A, B) \coloneqq \max\Big(\sup_{a\in A}\inf_{b\in B} \mathrm{d}_\mathcal{X}(a, b),\ \sup_{b\in B}\inf_{a\in A} \mathrm{d}_\mathcal{X}(a, b)\Big). \tag{1}$$

Focusing on one of the directions of the interior terms, we interpret it as a *distortion*, where

$$\overbrace{\sup_{a\in A}\inf_{b\in B}}^{\text{pairing}} \underbrace{\mathrm{d}_\mathcal{X}(a, b)}_{\text{cost}}. \tag{2}$$

This defines a non-symmetric relationship between A and B, made symmetric by taking the maximum in each direction.

[1] A graph is *simple* if it does not have any self-loops.

[2] For simplicity, we refer to it as an *attributed graph*.

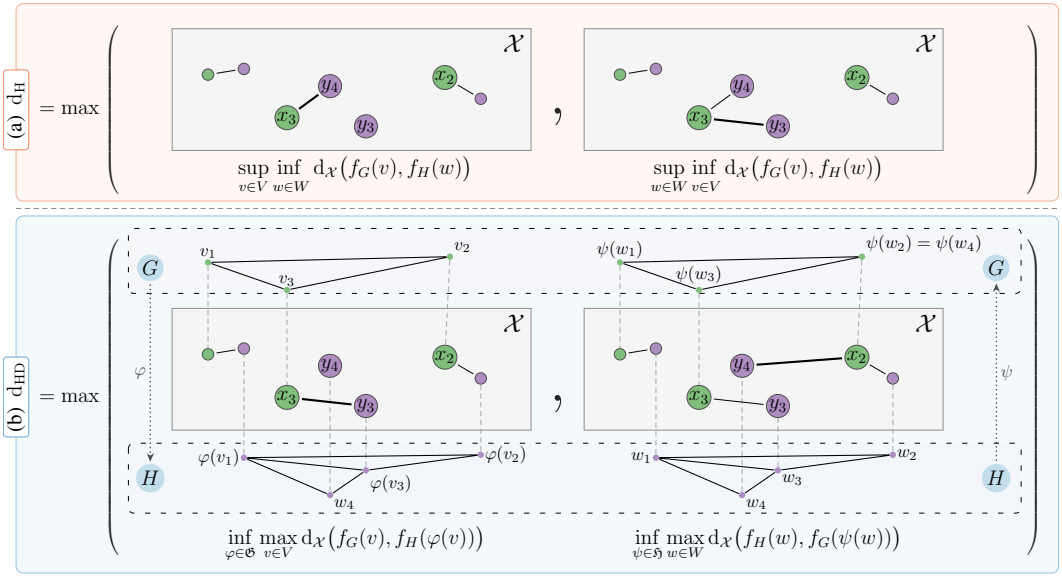


Figure 1: In this example $\mathcal{X} = \mathbb{R}^2$, $G = (V, E, f_G)$, $H = (W, E', f_H)$ and $\in \mathrm{Im} f_G \subset \mathbb{R}^2$ and $\in \mathrm{Im} f_H \subset \mathbb{R}^2$. Row (a) shows $\mathrm{d_H}$, the *Hausdorff metric*, evaluated on all pairs, resulting in $\mathrm{d_H}(G, H) = \max\left(\mathrm{d}_{\mathcal{X}}(x_3, y_4), \mathrm{d}_{\mathcal{X}}(x_3, y_3)\right)$. Row (b) shows $\mathrm{d_{HD}}$, the **homomorphism distortion**, where the comparison of a pair is contingent on to the *existence* of a homomorphism. Concretely, in the $G \to H$ direction, the pair (x_3, y_4) is not accounted for and instead, it is replaced by the pair (x_3, y_3). Similarly, in the $H \to G$ direction, the resulting pair is (x_2, y_4). In the end, $\mathrm{d_H}(G, H) = \mathrm{d}_{\mathcal{X}}(x_3, y_4) \neq \mathrm{d}_{\mathcal{X}}(x_2, y_4) = \mathrm{d_{HD}}(G, H)$.

Homomorphism-restricted distances. An attributed graph encodes two kinds of information: features and structure. The Hausdorff metric accounts only for features, missing the opportunity to exploit structural information. Our proposal, in contrast, looks at attributed graphs and incorporate the structural information by restricting how we can map one-to-another and *vice versa*. One natural choice of map is a *homomorphism*, which encodes how a graph can be embedded into another while respecting immediate neighborhoods. Our *graph homomorphism distortion* aims to encode a notion of distance depending on the vertex attributes of graphs, thus providing a trade-off between expressivity and computation. It consists in a *restriction* of the Hausdorff metric, in terms of which pairs (a, b) for $a \in A$ and $b \in B$ are evaluated, with respect to the existence of a homomorphism. We change Eq. (2) as follows:

Definition 1 (Attribute distortion)**.** Let $G = (V, E, f_G)$ and $H = (V', E', f_H)$ be two attributed graphs with $\mathfrak{G} = \mathrm{Hom}(G, H)$ the set of all homomorphisms $G \to H$, $f_G \colon V(G) \to \mathcal{X}$ and $f_H \colon V(H) \to \mathcal{X}$. Their *attribute distortion* is

$$\mathrm{dis}(G, H) \coloneqq \inf_{\varphi\in\mathfrak{G}} \max_{v\in V(G)} \mathrm{d}_{\mathcal{X}}\left(f_G(v), f_H(\varphi(v))\right), \tag{3}$$

where we set $\inf \emptyset \coloneqq \infty$ and $\mathrm{d}_{\mathcal{X}}(x, y) = \|x - y\|$

Figure 1 illustrates an example. To obtain the symmetric structure of our original muse, the Hausdorff distance, we symmetrize Definition 1 as follows.

Definition 2 (Graph homomorphism distortion)**.** Given two attributed graphs $G = (V, E, f_G)$ and $H = (V', E', f_H)$, we define their *graph homomorphism distortion* $\mathrm{d_{HD}}$ as

$$\mathrm{d_{HD}}(G, H) \coloneqq \max\left(\mathrm{dis}(G, H), \mathrm{dis}(H, G)\right). \tag{4}$$

In the remainder of the manuscript, we refer to this as *homomorphism distortion*. When we talk about an attributed graph G we implicitly refer to the tuple (V, E, f_G), and we use $\mathfrak{G}$ to denote the set of homomorphisms to another graph H. While Definition 2 appears restrictive at first glance, as it requires the homomorphisms to be known explicitly, we will later show how to evaluate it based on families of graphs for which homomorphisms can be efficiently calculated.

3.1 General properties of the Homomorphism Distortion

We first show that the homomorphism distortion is a pseudo-metric (cf. Appendix A) on vertex-attributed graphs. We defer all proofs to Appendix B for clarity.

Lemma 1 (Homomorphism distortion is a pseudo-metric)**.** *The* $\mathrm{d}_{\mathrm{HD}}(\cdot,\cdot)$ *is a pseudo-metric, i.e., the following properties hold for all attributed graphs* G, G', G''*:*

1. $\mathrm{d}_{\mathrm{HD}}(G,G') \geq 0$ *and* $\mathrm{d}_{\mathrm{HD}}(G,G) = 0$,
2. $\mathrm{d}_{\mathrm{HD}}(G,G') = \mathrm{d}_{\mathrm{HD}}(G',G)$,
3. $\mathrm{d}_{\mathrm{HD}}(G,G') \leq \mathrm{d}_{\mathrm{HD}}(G,G'') + \mathrm{d}_{\mathrm{HD}}(G'',G')$.

The homomorphism distortion is also 1-*Lipschitz continuous*, i.e., the closer the attribute functions, the smaller the distortion. This property will prove useful for improving the computability of d_{HD}.

Proposition 1 (Lipschitz continuity)**.** *Given two attribute functions* $f, g\colon V \to \mathcal{X}$ *on the same graph* G*, we have* $\mathrm{d}_{\mathrm{HD}}((V,E,f),(V,E,g)) \leq \|f-g\|_\infty$ *where* $\|\cdot\|_\infty$ *denotes the supremum norm induced by the norm of* $\mathcal{X}$*.*

Restricting the space of vertex attributes. We are interested in attributed graphs whose attribute function is to some extent related to the graph structure. As such, we restrict our focus to functions respecting some desirable properties. Concretely, for an attributed graph $G = (V, E, f_G)$, we expect that if $u, v \in V$ are *automorphic*, then $f_G(u) = f_G(v)$, i.e., f_G respects *global* structure. Similarly, for another graph $H = (W, E', f_H)$, where $u \in V$ and $w \in W$, if there exists an isomorphism from G to H such that $\varphi(u) = w$, then $f_G(u) = f_H(w)$, i.e., f_G and f_H respect *local* structure.[3] To formalize the previous restriction, we need to use the viewpoint of rooted graphs. Recall that the space of all rooted graphs is $\mathcal{G}_\bullet$ and their equivalence, isomorphism, is $\simeq_\bullet$. We can enforce both previous restrictions by looking at the quotient space, denoted $\mathcal{G}_\bullet/\simeq_\bullet$. Consider a function Ψ on the quotient space $\Psi\colon \mathcal{G}_\bullet/\simeq_\bullet \to \mathcal{X}$ and let $(G,v),(H,w) \in \mathcal{G}_\bullet$ be two rooted graphs. It follows immediately that if $G \simeq H$ and $G \neq H$, for any pair of nodes $v \in V(G)$ and $w \in V(H)$ and any isomorphism $\varphi(v) = w$, the rooted graphs will collapse to the same equivalence class. If $G = H$, the claim still holds since an automorphism is also an isomorphism. The attribute restriction we propose has a name: a *rooted-graph invariant*. This is a function on rooted graphs that evaluates to the same value in case of isomorphism, i.e., $(G,v) \simeq_\bullet (H,w) \implies f_G(v) = f_H(w)$. With this in mind, we introduce one final notational trick in the spirit of rigor and to show how this machinery applies to general graphs: For a graph G, let $\iota_G\colon V(G) \to \mathcal{G}_\bullet/\simeq_\bullet$ be the quotient map that takes each possible rooted graph in G as (G,v) for $v \in V(G)$ and maps it to its corresponding equivalence class. We define the attribute function on a graph induced by Ψ as

$$f_G \coloneqq \Psi \circ \iota_G, \quad f_G(v) = \Psi\left([(G,v)]\right). \tag{5}$$

Moving forward, we only refer to Ψ to discuss restrictions on the classes of functions we consider in attributed graphs. Also, we will assume that, whenever we refer to an attributed graph with a function f_G, said function is always induced by some Ψ. This implies that, if the graph attributes are defined algorithmically, then they should be obtained through the *same* algorithm on all graphs $G \in \mathcal{G}_n$. For example, for two attribute functions f_G, f_H on graphs G, H it cannot happen that f_G is the degree of a vertex and f_H is constant. We now turn to illustrate the discriminative power of the homomorphism distortion in terms of isomorphisms.

3.2 Discriminative power: Non-isomorphic distinction and its flaws

We circle back to our initial inspiration, *what is the difference between the homomorphism distortion and the Hausdorff metric on the feature space*? Consider a fixed Ψ and the normed space $(\bigcup_{G\in\mathcal{G}} V(G), \mathrm{d}_\Psi)$ with $\mathrm{d}_\Psi(v \in G, w \in H) \coloneqq \mathrm{d}_{\mathcal{X}}(f_G(v), f_H(w)) = \|f_G(v) - f_H(w)\| = \|\Psi\left([(G,v)]\right) - \Psi\left([(H,w)]\right)\|$, then the setting described in Eq. (1) is the Hausdorff metric on the normed spaces, i.e.,

$$\mathrm{d}_{\mathrm{H}}(V(G),V(H)) = \max\left(\sup_{v\in V(G)}\inf_{w\in V(H)} \mathrm{d}_\Psi(v,w), \sup_{w\in V(H)}\inf_{v\in V(G)} \mathrm{d}_\Psi(v,w)\right). \tag{6}$$

[3] The first criterion means that f_G is structurally dependent, whereas the second one means that f_G and f_H are *permutation invariant*.

The following relationship between representation and the homomorphism distortion holds for any scale in terms of the Hausdorff distance.

Theorem 1. *The* d_{HD} *is* at least as discriminative *as* Ψ*, i.e.,* $\mathrm{d}_{\mathrm{H}}(V(G), V(H)) \leq \mathrm{d}_{\mathrm{HD}}(G, H)$.

Importantly, $\mathrm{d}_{\mathrm{HD}}(G, H) = 0$ implies that the Hausdorff distance is 0. The homomorphism distortion is also permutation invariant because of the composition of homomorphisms.

Proposition 2 (Permutation invariance)**.** *The* d_{HD} *is permutation-invariant, i.e., for two attributed graphs* G, H *with permutations* $\rho\colon V(G) \to V(G)$ *and* $\rho'\colon V(H) \to V(H)$*, we have* $\mathrm{d}_{\mathrm{HD}}(\rho(V(G)), \rho'(V(H))) = \mathrm{d}_{\mathrm{HD}}(G, H)$.

Thus, the choice of Ψ provides us with a *knob* that changes its power to distinguish between graphs. This implies that for certain graph encodings, even some simple graphs cannot be distinguished. We illustrate in the following corollaries both ends of the spectrum of expressivity that Ψ provides.

Corollary 1. *The following statements for* Ψ *hold:*

1. *If* Ψ *is constant, then* $\mathrm{d}_{\mathrm{HD}}(G, H) = 0 \iff \mathrm{core}(G) \simeq \mathrm{core}(H)$.
2. *If* Ψ *is injective, then* $\mathrm{d}_{\mathrm{HD}}(G, H) = 0 \iff G \simeq H$.
3. *If* Ψ *is injective, then* d_{HD} *is a metric on the quotient.*

Between cases 1. and 2. there is a range of choices of Ψ which will allow certain graphs to be distinguished and other not. For example, if two graph G, H are *co-spectral*, i.e., their adjacency matrices have the same multiset of eigenvalues, then a d_{HD} with Ψ based on the spectrum will not be able to distinguish them. On the other hand, certain kinds of hard graphs such as the Rook's graph on a 4×4 grid and the Shrikhande graphs do not share a core, resulting in an infinite distance. The final scenario, where Ψ is injective and each isomorphism class is assigned a "canonical" labeling, is the *graph canonization problem* (Babai & Luks, 1983). Solving the *graph isomorphism problem* consists in verifying that the output of two "canonical" graph representations is equal.

Beyond strict isomorphism. Distinction criteria based on isomorphism tests are fundamentally limited when comparing (attributed) graphs: They fail to account for subtle differences that might be defining characteristics in how a label is assigned to a graph. For example, a label could depend on the *core* of a graph and not necessarily on its whole connectivity. This phenomenon is also present in common graph datasets (Ivanov et al., 2019). We identify cases where highly similar graphs cannot be distinguished by the homomorphism distortion. In these instances, the high similarity is given by induced subgraphs. We begin by analyzing the directional distortion between a graph G and an induced subgraph G' of G.

Proposition 3. *Let* $G = (V, E, f_G)$ *be an attributed graph and* $G' = (V', E', f_{G'})$ *be an attributed subgraph of* G*. The following statements hold if* $f_{G'}$ *is a restriction of* f_G *to* V'*:*

1. $\mathrm{dis}(G, G') = 0$,
2. *If additionally for any* $n \in \mathbb{N}$, $G = K_n$ *or* $G = C_\ell$ *for odd* $\ell > 2 \in \mathbb{N}$*, then* $\mathrm{dis}(G, G') > 0$.

The distortion *from* an induced subgraph to its supergraph is uninformative, i.e., it is always 0. However, the statement above is asymmetrical in that the distortion *to* an induced subgraph can be nonzero. Note that our overall graph homomorphism distortion d_{HD} maintains its symmetry, because d_{HD} is defined as the maximum over both directional infima; cf. Eq. (4). This concludes our analysis of the basic properties of the homomorphism distortion. In light of recent results from Jin et al. (2024) and Bao et al. (2025), we investigate the properties and feasibility of an *encoding* consisting of homomorphism distortions to a reference set.

3.3 Homomorphism Distortion Encoding

Our main contribution to machine learning is to zoom in on the potential of the homomorphism distortion as a way to *characterize* graphs, or analogously to obtain a graph encoding. A robust characterization acts as a fine-grained *signature*, enabling effective comparison between graphs. While permutation invariance is the theoretical ideal condition for distinguishing isomorphic graphs, practical machine learning tasks often require broader capabilities. In this context, encodings have emerged as a standard empirical approach. Let $\mathcal{F}$ be a family of graphs. The homomorphism distortion encoding of an attributed graph G is $\Gamma(G)_{\mathcal{F}} = [\mathrm{d}_{\mathrm{HD}}(F, G)]_{F \in \mathcal{F}}$, where $\Gamma(G)_{\mathcal{F}} \in \mathbb{R}^{|\mathcal{F}|}$ is

a *signature* of homomorphism distortions from each graph $F \in \mathcal{F}$ to G. We now demonstrate that, to achieve permutation-invariance of our encoding, it is *sufficient* to have $\mathcal{F} = \mathcal{G}_n$.

Lemma 2. *Let $G \in \mathcal{G}_n$ be an attributed graph. If $\mathcal{F} = \mathcal{G}_n$, then $\Gamma(G)_{\mathcal{F}}$ is (permutation) invariant.*

Notice that Lemma 2 also follows from Proposition 2. With this, we can extend the previous result to show that the homomorphism distortion encoding is complete.

Proposition 4. *Let $G \in \mathcal{G}_n$ be an attributed graph. If $\mathcal{F} = \mathcal{G}_n$, then $\Gamma(G)_{\mathcal{F}}$ is complete up to the distinguishing power of f_G*

All the appealing properties of the homomorphism distortion and its promise as an encoding have a drawback: Enumerating homomorphisms, i.e., computing the maps, and computing each of the encoding for a large $\mathcal{F}$ is intractable. We remedy this by introducing alternative formulations that make our method computationally feasible.

3.4 Computational Feasibility

In practice, computing the number of homomorphisms between two graphs G, H is an NP-complete problem (Böker et al., 2019). However, for specific choices or classes of G and H, this problem is considerably easier. Díaz et al. (2002) show that counting homomorphisms becomes *computable* in the case of $\mathrm{Hom}(F, G)$, where F has bounded treewidth k. Algorithms that leverage tree decomposition exist (Bodlaender & Koster, 2007), thus we rely on dynamic programming solutions to make the computations (Bodlaender et al., 2013). Nevertheless, we need to ensure that we control the treewidth k. We observe that the following relationship hints at a possible strategy to reduce treewidth.

Proposition 5. *Let G, H be two attributed graphs and G', H' subgraphs of each one. Then*

$$\mathrm{d}_{\mathrm{HD}}(G, H) \geq \mathrm{d}_{\mathrm{HD}}(G', H') - (\mathrm{d}_{\mathrm{HD}}(G, G') + \mathrm{d}_{\mathrm{HD}}(H, H'))\,.$$

If we choose a method for obtaining G', H' that results in a smaller treewidth, we have an effective proxy for the distance. The utility of Proposition 5 will become apparent when coupled with Proposition 3 in the experimental section. Here, picking G', H' to be induced subgraphs results in a bound of the form $\mathrm{d}_{\mathrm{HD}}(G, H) \geq \mathrm{d}_{\mathrm{HD}}(G', H')$, where we never "overestimate" the difference in graphs and only underestimate it.[4] We can also make use of this bound using other reduction methods, like isospectral reductions (Kempton & Tolbert, 2024), where G', H' become co-spectral and therefore, for an appropriate Ψ, also yield $\mathrm{d}_{\mathrm{HD}}(G, G') = \mathrm{d}_{\mathrm{HD}}(H, H') = 0$.

We now turn our attention to the discriminative power of the homomorphism distortion with respect to Weisfeiler–Leman test by framing it as a (pseudo)-metric to allow for comparison.

4 The WL Distance and Its Shortcomings

Given that the Weisfeiler–Leman test for graph isomorphism remains the *de facto* benchmark for graph expressivity analysis, we first discuss it and its limitations. We reframe the 1-WL test in two steps. The first part is a procedure that produces an encoding of two graphs based on their topology. Secondly, it compares their encodings to measure their distance. To begin with, we algebraically formalize what it means for two graphs to be WL-equivalent.

Definition 3 (Stable 1-WL coloring encoding)**.** A function $c\colon V \to \mathbb{N}$ is a stable 1-WL coloring encoding if it is the result of applying the 1-WL algorithm until convergence, i.e., there is no change in colors from one iteration to another. We denote the set of m colors that this procedure yields by $\mathcal{C} = \{1, \ldots, m\} \subset \mathbb{N}$ and call the multiset $h(G) = \{\!\{c(v) | v \in G\}\!\}$ the *color histogram* of a graph G.

In the following, we refer to stable 1-WL colorings as *WL colorings*. Operating on our previous definitions of vertex-attribution, during this section, we fix Ψ to be a WL coloring and let the resulting attribute functions, i.e., f_G, reflect said coloring such that for a graph G, $\mathrm{Im}\, f_G = h(G)$. It then follows that the homomorphism distortion between graphs with a WL coloring can distinguish *at least as many* non-isomorphic graphs as the WL test.

[4] A consequence is that the distance, as well as the encoding, will remain *invariant* with respect to isomorphism.

Proposition 6 (Generalization of WL distinguishability). *Let G and H be two non-isomorphic graphs whose attribute functions $f_G = c$ and $f_H = c'$ correspond to their respective WL colorings. If WL can distinguish them, then $\mathrm{d_{HD}}(G, H) > 0$.*

In the following, we show that if we define a *(pseudo)-metric* associated to the two tests, a similar relation holds between them. To start with, we consider the equality comparison of two color histograms, on which WL equivalence is based.

Definition 4 (WL (pseudo)-metric). Let G, H be two graphs with Ψ a WL coloring and $h(G), h(H)$ the corresponding color histograms. This leads to a discrete (pseudo)-metric

$$\mathrm{d_{WL}}(G, H) = \begin{cases} 0 & \text{if } h(G) = h(H), \\ 1 & \text{otherwise.} \end{cases} \tag{7}$$

This distance is too coarse for common machine-learning applications. We thus extend it to use the WL encoding under a metric that compares multisets. Following Definition 4 and Theorem 1 we now establish an *ordering* of these metrics, which we interpret as their discriminative power.

Theorem 2 (Discriminative power of coloring). *Let G and H be two graphs and Ψ be a WL coloring. Then $\mathrm{d_{WL}}(G, H) \leq \mathrm{d_H}(V(G), V(H)) \leq \mathrm{d_{HD}}(G, H)$.*

Thus, the $\mathrm{d_{WL}}$ (pseudo)-metric is the *least* discriminative due to the information loss inherent in binary comparisons. While our extension in terms of the Hausdorff metric, $\mathrm{d_H}$, provides a richer signal in the form of a *continuous* measurement, $\mathrm{d_{HD}}$ achieves the highest discriminative power: It simultaneously accounts for both the structural attributes and the homomorphisms between G and H. This ordering of discriminative power implies a relationship between the induced topological spaces. Specifically, we obtain an inclusion of the associated topologies:

Proposition 7 (Granularity of metric topologies). *Let $\mathcal{T}_{\mathrm{HD}}$ and $\mathcal{T}_{\mathrm{WL}}$ be the metric topologies induced by $\mathrm{d_{HD}}$ and $\mathrm{d_{WL}}$, respectively. Then, on a set of attributed graphs $\mathcal{G}$ with Ψ as the WL coloring, we have $\mathcal{T}_{\mathrm{WL}} \subset \mathcal{T}_{\mathrm{HD}}$.*

Intuitively, this inclusion holds because a "stronger" pseudo-metric like $\mathrm{d_{HD}}$, which assigns larger distances, generates smaller, more selective open balls. The $\mathcal{T}_{\mathrm{HD}}$ topology is thus *finer*, meaning it distinguishes graph structures more precisely than $\mathcal{T}_{\mathrm{WL}}$. In terms of convergence, this implies that the condition for $\mathrm{d_{HD}}$ is stricter: Any sequence of graphs converging in the HD topology must necessarily converge in the WL topology, but not vice versa.

5 Experiments

Our theoretical analysis highlights the utility of the homomorphism distortion both as a *notion of difference* and as a *graph encoding*. In this section, we empirically validate these properties.[5] Specifically, we aim to answer the following research questions:

- Q1 Can the homomorphism distortion distinguish non-isomorphic highly symmetric graphs ?
- Q2 How do the features in the attributed graphs affect the homomorphism distortion?
- Q3 Does the homomorphism distortion encoding serve as a useful inductive bias?
- Q4 Does approximation by treewidth decomposition work in practice?

Functions Ψ over the graphs. The theoretical results require the attribute function Ψ to be a rooted-graph invariant. Thus, we use known *positional and structural encodings* (Kanatsoulis et al., 2025, PSEs), specifically, (i) Random Walk positional encodings (Dwivedi et al., 2022, RW), which use probabilities of random walks to characterize local structural roles, (ii) Shortest Path encodings (SP), which use sorted vectors of shortest-path distances to other vertices, and (iii) Effective Resistance encodings (ER), which assign a sorted vector of resistances from each vertex to each other.

Computing distortion between large graphs. By Proposition 5, we can approximate the true distance by constructing subgraphs that have a smaller treewidth and are thus computable. We restrict ourselves to $\mathrm{tw}(G) \in \{1, 2\}$ due to computability reasons and ensure all graphs satisfy in either case.

[5] Please refer to Appendix D for hardware details and the full list of hyperparameters.

For $\mathrm{tw}(G) = 1$ we take the *minimum spanning tree* of G, i.e., the smallest tree that ensures G is connected, and for $\mathrm{tw}(G) = 2$, we start with the vertex set V and progressively add edges from E that do not increase the treewidth past the threshold. If all of the remaining edges increase the treewidth, we terminate the process. The computational complexity of calculating $\mathrm{d_{HD}}$ is $\mathcal{O}(m^{k+1} \cdot k^2 \log k \cdot n)$, where $k := \mathrm{tw}(G)$ and $m := |H|$, for a target graph H. See Appendix D for a breakdown.

5.1 Non-learning Scenarios

We measure the discriminative power, the impact of features in attributed graphs and the effectiveness of treewidth decomposition (Q1, Q2, Q4) by employing the expressivity benchmark dataset `BREC` (Wang & Zhang, 2024). The task on this dataset is to distinguishing between highly similar non-isomorphic graphs. Following the methodology of Ballester & Rieck (2025), we establish a threshold under which graphs are isomorphic as `float32` machine precision, i.e., $\epsilon = 1.19 \times 10^{-7}$. Table 1 shows the results in distinguishing each class of graphs.

Table 1: The % of distinguished graphs in `BREC`.

Method	B	R	E	C	4	D
1-WL	0.0	0.0	0.0	0.0	0.0	0.0
3-WL	**100.0**	37.0	**100.0**	60.0	0.0	0.0
PH	**100.0**	**97.0**	**100.0**	6.0	**100.0**	5.0
Ψ			$\mathrm{tw}(G) = 1$			
SP	26.7	13.0	41.0	12.0	0.0	0.0
RW	**100.0**	55.0	**100.0**	13.0	0.0	**15.0**
ER	**100.0**	50.0	**100.0**	**94.0**	0.0	0.0
Ψ			$\mathrm{tw}(G) \leq 2$			
SP	58.3	14.0	51.0	12.0	0.0	0.0
RW	**100.0**	57.0	**100.0**	13.0	0.0	**15.0**
ER	**100.0**	50.0	**100.0**	**94.0**	0.0	0.0

The $\mathrm{d_{HD}}$ distinguishes non-isomorphic graphs. Our method performs particularly well on CFI (C) and distance-regular (D) graphs compared to non-GNN method 3-WL (Wang & Zhang, 2024) and results from Ballester & Rieck (2025), obtained using persistent homology (PH). Naturally, 1-WL cannot distinguish any graphs. Notice that PH outperforms all other methods in the case of regular and strongly-regular (R) and 4-distance regular (4). These results leads to a strong positive answer to Q1 ✓, with ER, followed by RW as the most effective encodings. This phenomenon provides insight into Q2, indicating the role that the choice of Ψ plays in the performance of the homomorphism distortion.

Decompositions are useful. For graphs with a treewidth larger than 3, enumerating homomorphisms is *intractable*. Thus, we have to rely on approximations. The quality of the approximation and the speedup it provides are inevitably related. We observe that in our experiments the trade-off is minimal. For $\mathrm{tw}(G) = 2$, computations take 3 hours; for $\mathrm{tw}(G) = 1$, computations only take 6 minutes, which is $\times 28$ faster, with a big and slight degradation on SP and RW, respectively. Overall, both decompositions provides reliable differentiation in addition to making the computations tractable with provides positive evidence for Q4 ✓.

5.2 Learning Scenarios

We address the remaining research questions (Q2, Q3, Q4) by using the homomorphism distortion encoding in a set of machine learning tasks. We focus on graph level tasks, since we propose a graph-level encoding, We add the signature information resulting from $\Gamma(G)$ as graph-level information. We embed this signature in a high-dimensional space using a function $\omega \colon \mathbb{R}^M \to \mathbb{R}^d$, which we parametrize with a 2-layer MLP.

Table 2: The MAE (↓) for graph regression on `ZINC-12k` (without edge features) using $\mathcal{F} = \{C_i \mid 3 \leq i \leq 8\}$ for comparability. Hom, Sub and Spasm are from Jin et al. (2024) and MoSe from Bao et al. (2025). Super-indices † and $*$ denote the addition of homomorphism and subgraph counts as node features, respectively. Standard deviations are given in the last digit, e.g., $0.380(9) = 0.380 \pm 0.009$.

		Count-based					*Ours*		
	Base	$\mathrm{Hom}^{\dagger}$	Sub^{*}	$\mathrm{Spasm}^{\circledast}$	Hom+F	MoSe	$\Gamma(G)$	$\Gamma(G)^{\dagger}$	$\Gamma(G)^{*}$
GAT	0.380(9)	0.229(8)	0.201(4)	0.147(4)	–	–	0.317(12)	0.183(5)	0.168(3)
GCN	0.282(7)	0.235(5)	0.198(3)	0.165(4)	0.207(8)	–	0.231(16)	0.189(6)	0.164(3)
GIN	0.246(19)	0.197(15)	0.143(2)	0.143(4)	0.174(5)	0.158(4)	0.214(16)	0.174(6)	0.131(3)

Evaluation. We evaluate molecular datasets `ZINC-12K`, with the same setup as Jin et al. (2024), which follows Dwivedi et al. (2023), `ogbg-molhiv` following Hu et al. (2020) and `NCI1` following Morris et al. (2020). We choose these due to optimal treewidth properties of the graphs in them (cf. Table S.1). For the remaining ones, we approximate the distances using the method in the previous subsection. We choose three GNN models (GAT, GCN, and GIN) and 5 random seed initializations, reporting the mean and standard deviation ($\mu \pm \sigma$). We choose the best hyperparameters based on the validation performance and report the scores on the held-out test set. Our comparison partners are HC (Wolf et al., 2023), Num (Nguyen & Maehara, 2020) and $\mathrm{Hom} + F$ (Welke et al., 2023).

The $\mathrm{d_{HD}}$ is an effective inductive bias and it complements homomorphism counts. Table 2 presents the results on `ZINC-12K` against results the state-of-the-art methods using homomorphism-based inductive biases. Using the homomorphism distortion encoding yields in better performance than the baseline and it is competitive with plain homomorphism counts across *all* models. Moreover, complementing our method with either Hom or Sub counts yields the best performance in GIN. Our interpretation is that count and homomorphism distortion information, while competitive among each other, show to be *complementary* as an inductive bias. Table 3 presents the results on additional molecular datasets where the homomorphism distortion with C_9 outperforms all plain count-based methods. We interpret these results as positive confirmation of Q3 ✓; additionally they provide further evidence for Q4 ✓.

Table 3: Performance of homomorphism-based encodings with GIN.

Method	`NCI1` ACC (↑)	`ogbg-molhiv` AUROC (↑)
Base	80.73 ± 1.87	77.00 ± 1.49
Count-based		
HC	77.30 ± 4.00	74.70 ± 1.00
Num	76.67 ± 1.16	–
$\mathrm{Hom} + F$	–	77.70 ± 1.50
Ours		
$\Gamma(G)_{\mathrm{RW}}$	81.95 ± 1.30	77.35 ± 1.59
$\Gamma(G)_{\mathrm{ER}}$	81.68 ± 1.40	77.10 ± 1.11
$\Gamma(G)_{\mathrm{SP}}$	82.46 ± 1.26	77.85 ± 1.06

6 Discussion

Our paper questions the current expressivity paradigm in graph representation learning and highlights some of its shortcomings, such as the reliance on purely structural comparisons: As a first step towards addressing these limitations, we introduce the graph homomorphism distortion, a novel measure of dissimilarity for attributed graphs. This (pseudo-)metric captures a continuous notion which fully characterizes graphs under certain assumptions, inducing a metric on the space of isomorphism equivalence classes. Moreover, we show that it complements existing expressivity measures such as 1-WL in terms of their topological granularity and discriminative power. By reframing the homomorphism distortion as a *structural encoding*, we obtain an encoding that (i) fully characterizes graphs in *theory*, (ii) is computable, and (iii) improves predictive performance, alone or as an addition to homomorphism counts, when used as an inductive bias in *practice*.

Limitations and future work. Our formulation suffers from some limitations, which we aim to address in future work. Most notably, our measure currently requires general assumptions on the vertex-attributed graphs to guarantee completeness, and obtaining more specific results would be desirable. Moreover, for practical reasons, we refrain from using more structural or positional encodings as vertex attribute functions; we believe that more suitable functions exist or may be learned. A thread along this paper is developing a computational method for analyzing graph expressivity, akin to WL, we leave this for future work. We also postpone additional theoretical inquiries concerning which graph properties (e.g., the presence/absence of cycles) are captured by different proxy graph families $\mathcal{F}$, and instead restrict our experiments to a limited exploration of such classes. We believe that a *substantial shift* in graph representation learning is required to move beyond restricted expressivity measurements. While the community is already (partially) aware of these flaws (Bechler-Speicher et al., 2025; Coupette et al., 2025), additional steps need to be taken to fix them. Our work proposes one particular direction in the search for better expressivity measurements.

AI USE STATEMENT

In this work, we used generative AI tools for obtaining prototype code for submitting tasks to our HPC infrastructure. We also used generative AI tools to create TikZ figure source code and formatting tables. We have not used AI tools for any other tasks with required or recommended disclosure. We have checked and refined all AI-generated outputs and take full responsibility for the contents of this work.

REPRODUCIBILITY STATEMENT

In Section 3.4 we present a method for tackling computational feasibility. In Section 5.1 we present the method for obtaining lower treewidth graphs, denoted **Computing distortion between large graphs**. In Section 5.1 we present results on `BREC` and along that the choices of threshold, ϵ, and the chosen treewwidths. In Section 5.2 we reference the method to encode the distortion vector and what it consists of. In Section B and C we provide proofs for all of our propositions, lemmas and theorems. In Section D we reference: hardware details used for the computations in **Hardware setup**, hyperparameters we sweep through in **Hyperparameters**, and a pseudo-code sketch of the computation of the homomorphism distortion in **Algorithm computations**.

REFERENCES

Arvind, V., Fuhlbrück, F., Köbler, J., and Verbitsky, O. On Weisfeiler-Leman invariance: Subgraph counts and related graph properties. *Journal of Computer and System Sciences*, 113:42–59, 2020.

Babai, L. and Luks, E. M. Canonical labeling of graphs. In *Proceedings of the 15th Annual ACM Symposium on Theory of Computing*, pp. 171–183, 1983.

Ballester, R. and Rieck, B. On the Expressivity of Persistent Homology in Graph Learning. In Wolf, G. and Krishnaswamy, S. (eds.), *Proceedings of the Third Learning on Graphs Conference*, volume 269 of *Proceedings of Machine Learning Research*, pp. 42:1–42:31. PMLR, 2025.

Bao, L., Jin, E., Bronstein, M. M., Ceylan, I. I., and Lanzinger, M. Homomorphism counts as structural encodings for graph learning. In *International Conference on Learning Representations*, 2025.

Bechler-Speicher, M., Finkelshtein, B., Frasca, F., Müller, L., Tönshoff, J., Siraudin, A., Zaverkin, V., Bronstein, M. M., Niepert, M., Perozzi, B., Galkin, M., and Morris, C. Position: Graph learning will lose relevance due to poor benchmarks. In Singh, A., Fazel, M., Hsu, D., Lacoste-Julien, S., Berkenkamp, F., Maharaj, T., Wagstaff, K., and Zhu, J. (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 81067–81089. PMLR, 2025.

Bodlaender, H. L. and Koster, A. M. C. A. Combinatorial optimization on graphs of bounded treewidth. *The Computer Journal*, 51(3):255–269, 2007.

Bodlaender, H. L., Bonsma, P., and Lokshtanov, D. The fine details of fast dynamic programming over tree decompositions. In Hutchison, D., Kanade, T., Kittler, J., Kleinberg, J. M., Mattern, F., Mitchell, J. C., Naor, M., Nierstrasz, O., Pandu Rangan, C., Steffen, B., Sudan, M., Terzopoulos, D., Tygar, D., Vardi, M. Y., Weikum, G., Gutin, G., and Szeider, S. (eds.), *Parameterized and Exact Computation*, volume 8246 of *Lecture Notes in Computer Science*, pp. 41–53. Springer, Cham, Switzerland, 2013.

Böker, J., Chen, Y., Grohe, M., and Rattan, G. The complexity of homomorphism indistinguishability. In Rossmanith, P., Heggernes, P., and Katoen, J.-P. (eds.), *44th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 138 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 54:1–54:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019.

Coupette, C., Wayland, J., Simons, E., and Rieck, B. No metric to rule them all: Toward principled evaluations of graph-learning datasets. In Singh, A., Fazel, M., Hsu, D., Lacoste-Julien, S., Berkenkamp, F., Maharaj, T., Wagstaff, K., and Zhu, J. (eds.), *Proceedings of the 42nd International*

Conference on Machine Learning, volume 267 of *Proceedings of Machine Learning Research*, pp. 11405–11434. PMLR, 2025.

Curticapean, R., Dell, H., and Marx, D. Homomorphisms are a good basis for counting small subgraphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 210–223, 2017.

Díaz, J., Serna, M., and Thilikos, D. M. Counting H-colorings of partial k-trees. *Theoretical Computer Science*, 2002.

Dwivedi, V. P., Luu, A. T., Laurent, T., Bengio, Y., and Bresson, X. Graph neural networks with learnable structural and positional representations. In *International Conference on Learning Representations*, 2022.

Dwivedi, V. P., Joshi, C. K., Luu, A. T., Laurent, T., Bengio, Y., and Bresson, X. Benchmarking graph neural networks. *Journal of Machine Learning Research*, 24(43):1–48, 2023.

Guo, Y., Tang, H., Ma, J., Xu, H., and Wei, Z. Position: Spectral GNNs rely less on graph Fourier basis than conceived. In Singh, A., Fazel, M., Hsu, D., Lacoste-Julien, S., Berkenkamp, F., Maharaj, T., Wagstaff, K., and Zhu, J. (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 81471–81493. PMLR, 2025.

Hagberg, A. and Conway, D. Networkx: Network analysis with Python, 2020. URL https://networkx.github.io.

Hell, P. and Nesetril, J. *Graphs and Homomorphisms*. Oxford University Press, 2004.

Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. Open graph benchmark: Datasets for machine learning on graphs. *Advances in Neural Information Processing Systems*, 33:22118–22133, 2020.

Isufi, E., Gama, F., Shuman, D. I., and Segarra, S. Graph filters for signal processing and machine learning on graphs. *IEEE Transactions on Signal Processing*, 72:4745–4781, 2024.

Ivanov, S., Sviridov, S., and Burnaev, E. Understanding isomorphism bias in graph data sets. *arXiv preprint arXiv:1910.12091*, 2019.

Jin, E., Bronstein, M. M., Ceylan, I. I., and Lanzinger, M. Homomorphism counts for graph neural networks: All about that basis. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 22075–22098. PMLR, 2024.

Joshi, C. K., Bodnar, C., Mathis, S. V., Cohen, T., and Lio, P. On the expressive power of geometric graph neural networks. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J. (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 15330–15355. PMLR, 2023.

Kanatsoulis, C., Choi, E., Jegelka, S., Leskovec, J., and Ribeiro, A. Learning efficient positional encodings with graph neural networks. In *International Conference on Learning Representations*, 2025.

Kempton, M. and Tolbert, J. Isospectral reductions and quantum walks on graphs. *Algebraic Combinatorics*, 7(1):225–243, 2024.

Kloks, T. *Treewidth: computations and approximations*. Springer, 1994.

Lovász, L. Operations with structures. *Acta Mathematica Academiae Scientiarum Hungarica*, 18(3): 321–328, 1967.

Lovász, L. *Large Networks and Graph Limits*, volume 60 of *Colloquium Publications*. American Mathematical Society, Providence, RI, 2012.

Morris, C., Ritzert, M., Fey, M., Hamilton, W. L., Lenssen, J. E., Rattan, G., and Grohe, M. Weisfeiler and Leman go neural: Higher-order graph neural networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):4602–4609, 2019. doi: 10.1609/aaai.v33i01.33014602.

Morris, C., Kriege, N. M., Bause, F., Kersting, K., Mutzel, P., and Neumann, M. Tudataset: A collection of benchmark datasets for learning with graphs. *arXiv preprint arXiv:2007.08663*, 2020.

Morris, C., Lipman, Y., Maron, H., Rieck, B., Kriege, N. M., Grohe, M., Fey, M., and Borgwardt, K. Weisfeiler and Leman go machine learning: The story so far. *Journal of Machine Learning Research*, 24(333):1–59, 2023.

Nguyen, H. and Maehara, T. Graph homomorphism convolution. In Daumé III, H. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 7306–7316. PMLR, 2020.

Vasileiou, A., Cervino, J., Frossard, P., Kanatsoulis, C. I., Morris, C., Schaub, M. T., Vandergheynst, P., Wang, Z., Wolf, G., and Levie, R. Graph learning should move beyond restrictive views of spectral and message-passing GNNs. *arXiv preprint arXiv:2602.10031*, 2026.

Vayer, T., Chapel, L., Flamary, R., Tavenard, R., and Courty, N. Fused Gromov-Wasserstein distance for structured objects. *Algorithms*, 13(9):212, 2020.

Wang, Y. and Zhang, M. An empirical study of realized GNN expressiveness. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 52134–52155. PMLR, 2024.

Welke, P., Thiessen, M., Jogl, F., and Gärtner, T. Expectation-complete graph representations with homomorphisms. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J. (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 36910–36925. PMLR, 2023.

Wolf, H., Oeljeklaus, L., Kühner, P., and Grohe, M. Structural Node Embeddings with Homomorphism Counts. *arXiv preprint arXiv:2308.15283*, 2023.

Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019.

Appendix (Supplementary Materials)

A Extended Background **13**

A.1 Mathematical Preliminaries 13

A.2 Graph Theory 13

B Proofs of Homomorphism Distortion Properties **14**

C Proofs of Results on Distinction Capabilities **17**

D Further Experimental Details **21**

A Extended Background

We subsequently provide more information about the mathematical preliminaries and graph theory, which serve to make this paper self-contained.

A.1 Mathematical Preliminaries

Definition 5 (Pseudo-metric)**.** A pseudo-metric is a tuple (X, d) for a set X and a function $\mathrm{d}\colon X \times X \to \mathbb{R}_{\geq 0}$ where for all $x, y, z \in X$ the following hold:

1. $\mathrm{d}(x, x) = 0$
2. $\mathrm{d}(x, y) = \mathrm{d}(y, x)$
3. $\mathrm{d}(x, z) \leq \mathrm{d}(x, y) + \mathrm{d}(y, z)$

Definition 6 (Metric)**.** A metric is a tuple (X, d) for a set X and a function $\mathrm{d}\colon X \times X \to \mathbb{R}_{\geq 0}$ where for all $x, y, z \in X$ the following hold:

1. $\mathrm{d}(x, x) = 0$
2. $x \neq y \Rightarrow \mathrm{d}(x, y) > 0$
3. $\mathrm{d}(x, y) = \mathrm{d}(y, x)$
4. $\mathrm{d}(x, z) \leq \mathrm{d}(x, y) + \mathrm{d}(y, z)$

From the previous definitions a metric is always a pseudo-metric. Additionally, we call a a pseudo-metric *extended* if it can take values in $[0, \infty]$, i.e. additing infinity.

Definition 7 (Metric Topology)**.** Let (X, d) be a metric space. The open ball of radius $\epsilon > 0$ centered at $x \in X$ is defined as:

$$B_d(x, \epsilon) = \{y \in X \mid d(x, y) < \epsilon\}. \tag{8}$$

The *metric topology* induced by d, denoted $\mathcal{T}_d$, is the topology generated by the basis of all open balls. A set $U \subseteq X$ is open in $\mathcal{T}_d$ if and only if for every $x \in U$, there exists an $\epsilon > 0$ such that $B_d(x, \epsilon) \subseteq U$.

A.2 Graph Theory

We briefly expand on several nonstandard concepts from graph theory, namely *graph cores* and *treewidth*, both of which will be crucial components in the subsequent proofs.

Definition 8 (Tree decomposition)**.** A *tree decomposition* of a graph G is a pair $(\mathcal{T}, \{X_a\}_{a \in V(\mathcal{T})})$ where $\mathcal{T}$ is a tree and $\{X_a\}_{a \in V(\mathcal{T})}$ is the family of subsets—referred to as bags—of $V(G)$ where the following conditions hold:

1. Each $v \in V(G)$ is in at least one bag X_a.
2. For each pair $\{v, u\} \in E(G)$, there is at least one bag X_a such that $v, u \in X_a$.
3. For each $v \in V(G)$ the set $\mathcal{T}_v := \{a \in V(\mathcal{T}) | v \in X_a\}$ induces a connected subgraph of $\mathcal{T}$.

The width of a tree decomposition is $\max_{a \in V(\mathcal{T})} |X_a| - 1$, the maximum size of a *bag* minus 1. The minimum width among all possible decompositions is called the *treewidth* of G, denoted $\mathrm{tw}(G)$.

B Proofs of Homomorphism Distortion Properties

Lemma 1 (Homomorphism distortion is a pseudo-metric)**.** *The $\mathrm{d_{HD}}(\cdot,\cdot)$ is a pseudo-metric, i.e., the following properties hold for all attributed graphs G, G', G'':*

1. $\mathrm{d_{HD}}(G, G') \geq 0$ *and* $\mathrm{d_{HD}}(G, G) = 0$,
2. $\mathrm{d_{HD}}(G, G') = \mathrm{d_{HD}}(G', G)$,
3. $\mathrm{d_{HD}}(G, G') \leq \mathrm{d_{HD}}(G, G'') + \mathrm{d_{HD}}(G'', G')$.

Proof. Since norms are non-negative, we have that $\mathrm{d_{HD}}(G, G') \geq 0$. If $G = G'$, the identity mapping $f(v) = v$ is in $\mathrm{Hom}(G, G)$, minimizing the right-hand side of Definition 2, which evaluates to 0. Thus, Item 1 holds. As for the symmetry, Definition 2 evaluates terms involving $\mathrm{Hom}(G, G')$ and $\mathrm{Hom}(G', G)$ respectively, and taking their maximum. This operation is inherently symmetric, so switching G and G' has no effect. Thus, $\mathrm{d_{HD}}$ is symmetric and Item 2 holds. Finally, to prove the triangle inequality, i.e., Item 3, we make use of the fact that we can compose homomorphisms. Without loss of generality, we assume that G, G', and G'' are homomorphically equivalent, i.e., there is at least one homomorphism between arbitrary pairs of these graphs. Taking any two such homomorphisms $\varphi\colon G \to G''$ and $\varphi'\colon G'' \to G'$ their composition $\varphi' \circ \varphi$ is a homomorphism from G to G'. As such, it occurs in the calculation of $\mathrm{d_{HD}}(G, G')$. Due to the properties of the norm used in Definition 2, splitting the composition into φ and φ' is always lower bounded by taking the infimum over all homomorphisms in $\mathrm{d_{HD}}(G, G')$. Hence, the triangle inequality described in Item 3 holds. □

Proposition 1 (Lipschitz continuity)**.** *Given two attribute functions $f, g\colon V \to \mathcal{X}$ on the same graph G, we have $\mathrm{d_{HD}}((V, E, f), (V, E, g)) \leq \|f - g\|_\infty$ where $\|\cdot\|_\infty$ denotes the supremum norm induced by the norm of $\mathcal{X}$.*

Proof. Since we are dealing with the same graph, the identity is in $\mathrm{Hom}(G, G)$. If we evaluate Definition 2 with the identity, we have that $\max_{v \in V} \|f(v) - g(v)\|$ appears in the right-hand side minimization. Since the infimum is taken for calculating Definition 2 the claim follows. □

Theorem 1. *The $\mathrm{d_{HD}}$ is* at least as discriminative *as Ψ, i.e., $\mathrm{d_H}(V(G), V(H)) \leq \mathrm{d_{HD}}(G, H)$.*

Proof. Let Ψ be a function on $\mathcal{G}_\bullet$ to some normed space $\mathcal{X}$ as defined previously. For two attributed graphs G, H this is exactly

$$\|f_G(v) - f_H(w)\| = \|\Psi([(G, v)]) - \Psi([(H, w)])\|, \tag{9}$$

then, the $\mathrm{d_H}(V(G), V(H))$ is given by

$$\mathrm{d_H}(V(G), V(H)) = \max\left(\sup_{v \in V(G)} \inf_{w \in V(H)} \mathrm{d}_{\mathcal{X}}(v, w), \sup_{w \in V(H)} \inf_{v \in V(G)} \mathrm{d}_{\mathcal{X}}(v, w) \right)$$

and the one-way distortion is given by

$$\mathrm{dis}(G, H) = \inf_{\varphi \in \mathfrak{S}} \max_{v \in V(G)} \|f_G(v) - f_H(\varphi(v))\|.$$

We focus our attention on

$$\sup_{v \in V(G)} \inf_{w \in V(H)} \|f_G(v) - f_H(w)\|.$$

First, we clarify some notation. A relationship $R \subseteq A \times B$ is a set containing some number of pairigs of elements of A and B. For a relationship $R \subseteq A \times B$ and a fixed element $a \in A$, we denote all pairs with a as the first element with $R(a, B) = \{(a, b) : b \in B\} \subseteq R$. Alternatively, we can see it

as the set containing pairings of a with any $b \in B$. The Hausdorff distance operates on the following relationship

$$Q \coloneqq V(G) \times V(H),$$

i.e all possible pairings (u, v) for $v \in V(G), w \in V(H)$. In contrast, assume that G and H are homomorphic and let $\varphi \in \mathrm{Hom}(G, H)$ be any homomorphism $\varphi\colon V(G) \to V(H)$. The pairing relationship for φ is

$$R_\varphi \coloneqq \{(v, \varphi(v)) : u \in V(G)\}.$$

Naturally, $R_\varphi \subseteq Q$ for any φ since we are restricting ourselves to less pairings, i.e., the ones that fulfill the homomorphism criteria. Now fix any $v \in V(G)$, the contention of the pairings holds

$$R_\varphi(v, V(H)) \subseteq Q(v, V(H)).$$

The infimum can then be bounded in terms of these subsets in terms of v along with all $w \in V(H)$ as

$$\begin{aligned}\inf_{w \in V(H)} \|f_G(v) - f_H(w)\| &= \inf_{(v,w) \in Q(v,V(H))} \|f_G(v) - f_H(w)\| \\ &\leq \inf_{(v,\varphi(v)) \in R_\varphi(v,V(H))} \|f_G(v) - f_H(w)\| \\ &= \|f_G(v) - f_H(\varphi(v))\|,\end{aligned}$$

where we arrive to the final equality since $R_\varphi(v, V(H))$ is the set with one element, namely $\{(v, \varphi(v))\}$. Since we chose v arbitrarily this holds for any $v \in V(G)$. Now, we can evaluate the outer supremum noting

$$\|f_G(v) - f_H(\varphi(v))\| \leq \sup_{u \in V(G)} \|f_G(u) - f_H(\varphi(u))\|$$

thus for $v \in V(G)$ it holds that

$$\inf_{w \in V(H)} \|f_G(v) - f_H(w)\| \leq \|f_G(v) - f_H(\varphi(v))\| \leq \sup_{u \in V(G)} \|f_G(u) - f_H(\varphi(u))\|$$

since this holds for all $v \in V(G)$, then it also holds for the supremum

$$\sup_{u \in V(G)} \inf_{w \in V(H)} \|f_G(u) - f_H(w)\| \leq \sup_{u \in V(G)} \|f_G(u) - f_H(\varphi(u))\| = \max_{v \in V(G)} \|f_G(u) - f_H(\varphi(u))\|$$

Since $V(G)$ is finite we have equality between $\sup$ and $\max$. Furthermore the above equation holds for any $\varphi \in \mathrm{Hom}(G, H)$ thus we may conclude

$$\sup_{u \in V(G)} \inf_{w \in V(H)} \|f_G(u) - f_H(w)\| \leq \inf_{\varphi \in \mathfrak{G}} \max_{v \in V(G)} \|f_G(u) - f_H(\varphi(u))\|.$$

If $\mathrm{Hom}(G, H) = \emptyset$ then we get ∞ and the bound still holds. We can provide a similar argument for the other direction and taking the max over both preserves the maximum relationship. The claim follows from this final fact. □

Proposition 2 (Permutation invariance)**.** *The* $\mathrm{d_{HD}}$ *is permutation-invariant, i.e., for two attributed graphs* G, H *with permutations* $\rho\colon V(G) \to V(G)$ *and* $\rho'\colon V(H) \to V(H)$*, we have* $\mathrm{d_{HD}}(\rho(V(G)), \rho'(V(H))) = \mathrm{d_{HD}}(G, H)$.

Proof. Let $G = (V, E, f_G)$ and $H = (V', E', f_H)$. For arbitrary permutations $\rho \in \mathrm{S}_{|V|}$ and $\gamma \in \mathrm{S}_{|V'|}$, let $G' = (\rho(V), E_{G'}, f_{G'})$ and $H' = (\gamma(V'), E_{H'}, f_{H'})$. Since $f_{G'} = \Psi \circ \iota_{G'}$ and $\iota_{G'} = \iota_G \circ \rho^{-1} = \iota_G$ we have $f_{G'} = \Psi \circ \iota_{G'} = \Psi \circ \iota_G = f_G$. Likewise we have $f_H = f_{H'}$.

The application of a permutation on a graph results in an isomorphic graph thus, if $G \simeq H$ then $\mathrm{d_{HD}}(G, H) = \mathrm{d_{HD}}(G', H') = 0$. Suppose $G \not\simeq H$ and $\varphi\colon G \to H$ is any homomorphism. Note $\rho^{-1} : G' \to G$ and $\gamma^{-1} : H' \to H$ thus we may define $\psi = \gamma \circ \varphi \circ \rho^{-1} \in \mathrm{Hom}(G', H')$.

For any vertex $w \in \rho(V)$ there exists a vertex $v \in V$ such that $f_{G'}(w) = f_{G'}(\rho(v)) = f_G(v)$ since $f_{G'}$ is permutation invariant. Additionally $f_{H'}(\psi(w)) = f_{H'}(\gamma \circ \varphi \circ \rho^{-1}(\rho(v))) = f_H(\varphi(v))$. Combining these we obtain

$$\max_{w \in \rho(V)} \|f_{G'}(w) - f_{H'}(\psi(w))\| = \max_{v \in V} \|f_G(\rho(v)) - f_H(\gamma \circ \varphi(v)\| = \max_{v \in V} \|f_G(v) - f_H(\varphi(v))\|.$$

Let $\mathfrak{G} = \mathrm{Hom}(G, H)$ and $\mathfrak{G}' = \mathrm{Hom}(G', H')$. Recall

$$\mathrm{dis}(G, H) = \inf_{\varphi \in \mathfrak{G}} \max_{v \in V} \|f_G(v) - f_H(\varphi(v))\|.$$

Assume for the fixed graph homomorphism $\varphi' \in \mathfrak{G}$, φ' defines the distortion i.e.,

$$\mathrm{dis}(G, H) = \max_{v \in V} \|f_G(v) - f_H(\varphi'(v))\|.$$

If we consider $\psi' = \gamma \circ \varphi' \circ \rho^{-1} \in \mathfrak{G}'$ then

$$\begin{aligned}\mathrm{dis}(G, H) = \max_{v \in V} \|f_G(v) - f_H(\varphi'(v))\| &= \max_{w \in \rho(V)} \|f_{G'}(w) - f_H(\psi'(w))\| \\ &\geq \inf_{\psi \in \mathfrak{G}'} \max_{w \in \rho(V)} \|f_{G'}(w) - f_{H'}(\psi(w))\| = \mathrm{dis}(G', H').\end{aligned}$$

In a similar manner we can reason

$$\mathrm{dis}(H, G) \geq \mathrm{dis}(H', G').$$

From which it follows

$$\mathrm{d_{HD}}(G, H) \geq \mathrm{d_{HD}}(G', H').$$

With an identical argument as above we have for any $\psi \in \mathfrak{G}'$ there exists a map $\varphi = \gamma^{-1} \circ \psi \circ \rho \in \mathfrak{G}$ such that $\max_{v \in V} \|f_G(v) - f_H(\varphi(v))\| = \max_{w \in \rho(V)} \|f_{G'}(w) - f_{H'}(\psi(w))\|$. Thus, we can deduce

$$\mathrm{d_{HD}}(G', H') \geq \mathrm{d_{HD}}(G, H).$$

To which we may conclude, $\mathrm{d_{HD}}(G, H) = \mathrm{d_{HD}}(G', H')$. □

Corollary 1. *The following statements for Ψ hold:*

1. *If Ψ is constant, then $\mathrm{d_{HD}}(G, H) = 0 \iff \mathrm{core}(G) \simeq \mathrm{core}(H)$.*
2. *If Ψ is injective, then $\mathrm{d_{HD}}(G, H) = 0 \iff G \simeq H$.*
3. *If Ψ is injective, then $\mathrm{d_{HD}}$ is a metric on the quotient.*

Proof. Let $\mathfrak{G} = \mathrm{Hom}(G, H)$ and $\mathfrak{G}' = \mathrm{Hom}(H, G)$. We proceed case by case:

1. Assume $\mathrm{d_{HD}}(G, H) = 0 \implies \max\{\mathrm{dis}(G, H), \mathrm{dis}(H, G)\} = 0$.
   Since $\max\{\mathrm{dis}(G, H), \mathrm{dis}(H, G)\} = 0$ we have $\mathfrak{G}, \mathfrak{G}' \neq \emptyset \implies \mathrm{core}(G) \simeq \mathrm{core}(H)$.

   Assume $\mathrm{core}(G) \simeq \mathrm{core}(H) \implies \mathfrak{G}, \mathfrak{G}' \neq \emptyset$. For any $\psi \in \mathfrak{G}$

   $$\max_{v \in V} \|f_G(v) - f_H(\psi(v))\| = \max_{v \in V} \|\Psi \circ \iota_G(v) - \Psi \circ \iota_H(\psi(v))\| = 0 \quad (\Psi \text{ is constant})$$

   Since $\mathfrak{G} \neq \emptyset$ we have $\mathrm{dis}(G, H) = 0$. Similarly we can conclude $\mathrm{dis}(H, G) = 0$. Therefore $\mathrm{d_{HD}}(G, H) = 0$.

2. Suppose $\mathrm{d_{HD}}(G, H) = 0$ and assume to the contrary that $G \not\simeq H$.
   Case 1: $\mathrm{core}(G) \not\simeq \mathrm{core}(H)$. Without loss of generality assume $\mathfrak{G} = \emptyset$. This would then give us $\mathrm{dis}(G, H) = \infty \implies \mathrm{d_{HD}}(G, H) = \infty \neq 0 = \mathrm{d_{HD}}(G, H)$. This is a contradiction.
   Case 2: $\mathrm{core}(G) \simeq \mathrm{core}(H) \implies \mathfrak{G}, \mathfrak{G}' \neq \emptyset$. For any $\psi \in \mathfrak{G}$ and any vertex $v \in V(G)$ we have $f_G(v) = \Psi \circ \iota_G(v)$. Since $G \not\simeq H$ it follows that there does not exist any rooted isomorphism between (G, v) and $(H, \psi(v))$. Thus $\iota_G(v) \neq \iota_H(\psi(v))$ and since Ψ is injective we have $f_G(v) = \Psi \circ \iota_G \neq \Psi \circ \iota_H(\psi(v)) = f_H(\psi(v))$. Putting this together

   $$\forall \psi \in \mathfrak{G}, \quad \max_{v \in V(G)} \|f_G(v) - f_H(\psi(v))\| > 0 \implies \mathrm{dis}(G, H) > 0.$$

   Since $\mathfrak{G}' \neq \emptyset$ we can similarly reason $\mathrm{dis}(H, G) > 0$. This gives us the contradiction $\mathrm{d_{HD}}(G, H) > 0$.

   On the other hand if we assume $G \simeq H$ by an isomorphism $\psi \in \mathfrak{G}$. For any vertex $v \in V(G)$, $f_G(v) = \Psi \circ \iota_G(v) = \Psi \circ \iota_H(\psi(v)) = f_H(\psi(v))$. Putting this together

   $$0 \leq \mathrm{dis}(G, H) \leq \|f_G(v) - f_H(\psi(v))\| = 0 \implies \mathrm{dis}(G, H) = 0.$$

   There exists an isomorphism $\varphi \in \mathfrak{G}'$ hence, in a similar manner we can reason $\mathrm{dis}(H, G) = 0$. Therefore $\mathrm{d_{HD}}(G, H) = \max\{\mathrm{dis}(G, H), \mathrm{dis}(H, G)\} = 0$.

3. From Proposition 2 the metric $\mathrm{d_{HD}}$ is well defined on the quotient. $\mathrm{d_{HD}}((G_1, v_1), (H_1, w_1)) = \mathrm{d_{HD}}((G_2, v_2), (H_2, w_2)) = \mathrm{d_{HD}}([(G, v)], [(H, w)])$, for any $(G_1, v_1), (G_2, v_2) \in [(G, v)]$ and $(H_1, w_1), (H_2, w_2) \in [(H, w)]$. $((\mathcal{G}_\bullet / \simeq_\bullet), \mathrm{d_{HD}})$ inherits non-negativity, symmetry, and the triangle inequality directly from the metric $(\mathcal{G}_\bullet, \mathrm{d_{HD}})$. Additionally 2. describes non-degeneracy confirming $((\mathcal{G}_\bullet / \simeq_\bullet), \mathrm{d_{HD}})$ is a metric space.

□

Proposition 5. *Let G, H be two attributed graphs and G', H' subgraphs of each one. Then*

$$\mathrm{d_{HD}}(G, H) \geq \mathrm{d_{HD}}(G', H') - (\mathrm{d_{HD}}(G, G') + \mathrm{d_{HD}}(H, H')) .$$

Proof. We start with $\mathrm{d_{HD}}(G', H')$ where by triangle inequality

$$\mathrm{d_{HD}}(G', H') \leq \mathrm{d_{HD}}(G', G) + \mathrm{d_{HD}}(G, H')$$

then

$$\mathrm{d_{HD}}(G, H') \leq \mathrm{d_{HD}}(G, H) + \mathrm{d_{HD}}(H, H')$$

combine together and we get that

$$\mathrm{d_{HD}}(G', H') \leq \mathrm{d_{HD}}(G, G') + \mathrm{d_{HD}}(G, H) + \mathrm{d_{HD}}(H, H')$$

we move the distances to the subgraphs to the left hand side

$$\mathrm{d_{HD}}(G', H') - (\mathrm{d_{HD}}(G, G') + \mathrm{d_{HD}}(H, H')) \leq \mathrm{d_{HD}}(G, H),$$

then the claim follows if we flip arround the inequality. □

C Proofs of Results on Distinction Capabilities

Lemma 3 (Equal colors under WL coloring equivalence)**.** *Let G and G' be two graphs of order n whose attribute functions c and c' correspond to their respective WL colorings. Assuming identical color sets $\mathcal{C} = \mathcal{C}'$, let $N_k(G)$ denote the number of nodes in G with color k. Then, $N_k(G) = N_k(G')$ for each color $k \in \mathcal{C}$.*

Proof. The stable coloring of nodes obtained through the WL algorithm results in a partition of vertices defined by their colors: $\{V_1, ...V_m\}$ which is a *equitable partition*, which means that, $\forall v_i \in V_i, d_{ij} := |\{u \in \mathcal{N}_g(V) : c(u) = j\}|$ is fixed. This means that the amount of edges between nodes of color i and color j is fixed and strictly determined by the colors themselves. This implies that, given a node of color i, all the values $d_{i,j}$ are uniquely defined for each j, which also implies that its neighborhood structure is fixed, in terms of the amount of neighboring nodes of each color. This means though that both G and G' satisfy the following system of equations in terms of the number of their nodes belonging to the different classes, $\mathbf{N} = [N_1, ..., N_m]$:

$$\begin{cases} N_i d_{i,j} = N_j d_{j,i} & \forall i, j \\ \sum_i N_i = n \end{cases} \tag{10}$$

We now demonstrate that the solution $\mathbf{N}$ to this system is unique by (1) showing that it is unique up to a scalar factor, by looking at the first equation and analyzing the associated transition matrix, and (2) exploiting the second equation to obtain the actual uniqueness.
Let $D_i = \sum_j d_{i,j}$ be the total degree of color class i. We define a matrix T of size $m \times m$ describing the random walk on the color classes: $T_{i,j} = \frac{d_{i,j}}{D_i}$, which is a row-stochastic transition matrix, i.e., rows express the probabilities of jumping from one color class to all other, through edge connections. We introduce a variable $\boldsymbol{\pi}$ to represent the total volume of each class (the total number of edges from such class to others), which will represent the stationary distribution of our discrete Markov chain. In particular, let $\pi_i = N_i D_i$. Substituting $d_{i,j} = T_{i,j} D_i$ into the detailed balance equation, i.e., the first row of Equation (10), we obtain:

$$N_i(T_{i,j} D_i) = N_j d_{j,i} \implies \pi_i T_{i,j} = N_j d_{j,i}.$$

Summing this equation over all possible starting classes i yields the global balance for class j:

$$\sum_{i=1}^{m} \pi_i T_{i,j} = N_j \sum_{i=1}^{m} d_{j,i} = N_j D_j = \pi_j. \tag{11}$$

In vector notation, this is equivalent to:

$$\boldsymbol{\pi} T = \boldsymbol{\pi}. \tag{12}$$

Equation (12) identifies the vector $\boldsymbol{\pi}$ as the left eigenvector of the transition matrix T corresponding to the eigenvalue $\lambda = 1$.
Since G and G' are assumed to be connected, the induced quotient graph is strongly connected, which in matrix terms implies that T is irreducible. The reason why this is relevant is that, according to the Perron-Frobenius Theorem, an irreducible stochastic matrix T has a unique 1-eigenvector, hence $\boldsymbol{\pi}$ is unique up to scalar normalization. Since $\pi_i = N_i D_i$ and the degrees D_i are fixed structural constants identical for both graphs, the uniqueness of $\boldsymbol{\pi}$ implies the uniqueness of the frequency vector $\mathbf{N}$ up to a scalar factor α:

$$\mathbf{N} = \alpha \cdot \mathbf{v}_{base},$$

where $\mathbf{v}_{base}$ is a fixed reference vector determined solely by the matrix T. This directly implies that also the solution to the first equation in Equation (10) is unique.
Finally, the normalization constraint $\sum N_i = n$ uniquely determines the scalar α:

$$\alpha = \frac{n}{\sum_i (\mathbf{v}_{base})_i}$$

Since the transition matrix T (derived from $d_{i,j}$) and the total count n are identical for both graphs, the frequency vector $\mathbf{N}$ is uniquely determined. Therefore:

$$N_k(G) = N_k(G') \quad \forall k \in \mathcal{C}.$$

□

Proposition 6 (Generalization of WL distinguishability)**.** *Let G and H be two non-isomorphic graphs whose attribute functions $f_G = c$ and $f_H = c'$ correspond to their respective WL colorings. If WL can distinguish them, then $\mathrm{d_{HD}}(G, H) > 0$.*

Proof. Since we are analyzing the WL isomorphism test between two graphs, we can assume that their cardinality is the same, that is, $|V| = |V'| = n$. Let G and H be two non-isomorphic graphs, and suppose WL can distinguish them. We can then consider two cases:

1. Suppose $\mathfrak{G} = \mathrm{Hom}(G, H) = \emptyset$ or $\mathfrak{H} = \mathrm{Hom}(H, G) = \emptyset$. Then, the graphs are not isomorphic, and we have
$$\mathfrak{G} = \emptyset \Rightarrow \inf_{\varphi \in \mathfrak{G}} = \infty,$$
$$\mathfrak{H} = \emptyset \Rightarrow \inf_{\psi \in \mathfrak{H}} = \infty,$$
which imply $\mathrm{d_{HD}}(G, H) = \infty > 0$.

2. Suppose $\mathfrak{G} \neq \emptyset$ and $\mathfrak{H} \neq \emptyset$, and define the final graph coloring (obtained through the color refinement algorithm) as integers features, e.g. $c_i = i$, such that
$$\|c(v) - c'(\varphi(v))\| \iff c(v) \neq c'(\varphi(v)).$$
Suppose now that $\mathrm{d_{HD}}(G, H) = 0$. Then,
$$\exists \varphi \in \mathfrak{G} : \forall v \in V(G) \quad \|c(v) - c'(\varphi(v))\| = 0,$$
$$\exists \psi \in \mathfrak{H} : \forall w \in V(H) \quad \|c'(w) - c(\psi(w))\| = 0$$
Which implies that
$$\exists \varphi \in \mathfrak{G} : \forall v \in V(G) \quad c(v) = c'(\varphi(v)),$$
$$\exists \psi \in \mathfrak{H} : \forall w \in V(H) \quad c'(w) = c(\psi(w)).$$
Define now $\mathcal{C}(V) :=$ set of colors on $V(G)$. Then two previous two equations imply, respectively,
$$\mathcal{C}(V(G)) \subseteq \mathcal{C}(V(H)),$$

$$\mathcal{C}(V(H)) \subseteq \mathcal{C}(V(G)),$$

which combined imply $\mathcal{C}(V) = \mathcal{C}(V(H))$. Now, since the set of colors appearing on the two graphs are the same, the only way for the WL test to result in distinguishing them is that their histogram should be different, that is, that there exist at least one color the amount of the appearance of each color in the two graphs differ. To elaborate on this point, let's define, for each color $k \in \mathbb{N}$, $N_k(G)$ as the amount of nodes in G with color k. So, the color histogram for graph G is defined by $[N_1(G), N_2(G), ...N_K(G)]$, with $K = |\mathcal{C}(V(G))| = |\mathcal{C}(V(H))|$, and the one for H is defined by $[N_1(H), N_2(H), ...N_K(H)]$. At this point, we can exploit Lemma 3, obtaining the result that the histograms for G and H are the same. This would lead to WL not distinguishing the two graphs, which leads to a contradiction with the initial assumptions. This implies that the assumption $\mathrm{d_{HD}}(G, H) = 0$ cannot hold, hence $\mathrm{d_{HD}}(G, H) > 0$.

□

Theorem 2 (Discriminative power of coloring)**.** *Let G and H be two graphs and Ψ be a WL coloring. Then $\mathrm{d_{WL}}(G, H) \leq \mathrm{d_H}(V(G), V(H)) \leq \mathrm{d_{HD}}(G, H)$.*

Proof. We can address the two inequalities in order: (1) first of all, we will show $\mathrm{d_{WL}}(G, H) \leq \mathrm{d_H}(G, H)$, and then (2) $\mathrm{d_H}(G, H) \leq \mathrm{d_{HD}}(G, H)$.

1. If the histograms of the graphs are equal $\mathrm{h}(G) = \mathrm{h}(H)$, then both $\mathrm{d_{WL}}$ and $\mathrm{d_H}$ are 0, and because $0 \leq 0$, the inequality holds. In case the histograms differ, we will have $\mathrm{d_{WL}}(G, H) = 1$. Additionally, because the histograms of G and H differ, from Lemma 3 there will be at least one color c that appears in only one between $h(G)$ and $h(H)$. The existence of such color will imply by definition that $\mathrm{d_H}(G, H) \geq 1$, leading to $\mathrm{d_{WL}}(G, H) \leq \mathrm{d_H}(G, H)$.

2. Consider now the second inequality: $\mathrm{d_H}(G, G') \leq \mathrm{d_{HD}}(G, H)$. Suppose that $\mathrm{d_H}(G, H) =: \alpha$ and, without loss of generality, suppose that this value of the distance is given by a specific pair of nodes $v \in G \quad \text{and} \quad w \in H$, i.e., $\alpha = \|c(v) - c'(w)\|$. But also the following set of equalities holds, given a specific φ:

$$\begin{aligned}
\mathrm{d_H}(G, H) = & \\
&= \alpha \\
&= \|c(v) - c'(w)\| \\
&= \min_{\hat{v} \in H} \|c(v) - c'(\hat{v})\| \\
&\leq \inf_{\varphi \in \mathfrak{G}} \|c(v) - c'(\varphi(v))\| \\
&\leq \inf_{\varphi \in \mathfrak{G}} \{\max_{v \in G} \|c(v) - c'(\varphi(v))\|\} \\
&= \mathrm{dis}(G, H) \\
&\leq \mathrm{d_{HD}}(G, H)
\end{aligned}$$

which proves the statement.

□

Proposition 7 (Granularity of metric topologies)**.** *Let $\mathcal{T}_{\mathrm{HD}}$ and $\mathcal{T}_{\mathrm{WL}}$ be the metric topologies induced by $\mathrm{d_{HD}}$ and $\mathrm{d_{WL}}$, respectively. Then, on a set of attributed graphs $\mathcal{G}$ with Ψ as the WL coloring, we have $\mathcal{T}_{\mathrm{WL}} \subset \mathcal{T}_{\mathrm{HD}}$.*

Proof. We proceed to show the containment $\mathcal{T}_{\mathrm{WL}} \subseteq \mathcal{T}_{\mathrm{HD}}$ by the basis characterization

$$\mathcal{T}_{\mathrm{WL}} \subseteq \mathcal{T}_{\mathrm{HD}} \iff \forall G \in \mathcal{G} \text{ and } G \in B \in \mathcal{B}_{WL},\ \exists B' \in \mathcal{B}_{HD} \text{ such that } G \in B' \subseteq B.$$

For any $G \in \mathcal{G}$ the elements of $\mathcal{B}_{\mathrm{WL}}$ are defined as the open balls w.r.t. metric $\mathrm{d_{WL}}$, thus denote the open ball of radius $r > 0$ centered at G by $B = B_{\mathrm{WL}}(G, r)$.

1. If $r > 1$. We have $B = B_{\mathrm{WL}}(G, r) = \mathcal{G}$. Set
$$r' = \max_{H \in \mathcal{G}} \{\mathrm{d}_{\mathrm{HD}}(G, H)\} + \epsilon \quad , \text{ for } \epsilon > 0$$
then $B' = B_{\mathrm{HD}}(G, r') = \mathcal{G}$. Thus there exists $B' \in \mathcal{B}_{\mathrm{HD}}$ such that $G \in B' \subseteq B$.
2. If $r < 1$. We have $B = B_{\mathrm{WL}}(G, r) = \{H \in \mathcal{G} |\ \mathrm{d}_{\mathrm{WL}}(G, H) = 0\}$. Set
$$r' = \min_{H \in \mathcal{G}} \{\mathrm{d}_{\mathrm{HD}}(G, H)\}$$
then $B' = B_{\mathrm{HD}}(G, r') = \{H \in \mathcal{G} |\ \mathrm{d}_{\mathrm{HD}}(G, H) = 0\}$. We have $G \in B'$ and $B' \in \mathcal{B}_{\mathrm{HD}}$. For any graph $H \in B'$ from Theorem 2 we have $0 \leq \mathrm{d}_{\mathrm{WL}} GH \leq \mathrm{d}_{\mathrm{HD}}(G, H) = 0$ thus $\mathrm{d}_{\mathrm{WL}}(G, H) = 0$ and $H \in B$ therefore $B' \subseteq B$. There exists $B \in \mathcal{B}_{\mathrm{HD}}$ such that $G \in B' \subseteq B$.

From 1. and 2. it follows that $\mathcal{T}_{\mathrm{WL}} \subseteq \mathcal{T}_{\mathrm{HD}}$. To show the strict containment consider the Shrikhande graph S and the 4×4 Rook's graph R (Arvind et al., 2020). In Arvind et al. (2020) authors show that for S and R color histograms $h(S) = h(G)$ and they are not homomorphic since S contains a 4-clique and S does not. Thus, $\mathrm{d}_{\mathrm{WL}}(S, R) = 0$. For any open set $U \in \mathcal{T}_{\mathrm{WL}}$ where $R \in U$ we can decompose U into its basis elements
$$U = \bigcup_{i \in I} B_i$$
here B_i are all basis elements and I is some indexing set. Thus for some graph $G \in \mathcal{G}$ with radius $r' > 0$ and index $\lambda \in I$ we have $R \in B_\lambda = B_{\mathrm{WL}}(G, r')$. Since $\mathrm{d}_{\mathrm{WL}}(S, R) = 0$ it follows $S \in B_\lambda \subseteq U$. Thus any open set in $\mathcal{T}_{\mathrm{WL}}$ which contains R contains S. On the other hand since
$$\mathrm{Hom}(R, S) = \emptyset \implies \mathrm{dis}(R, S) = \infty \implies \mathrm{d}_{\mathrm{HD}}(R, S) = \infty.$$
For any fixed $0 < r < \infty$ consider the open set $B = B_{\mathrm{HD}}(R, r) \in \mathcal{T}_{\mathrm{HD}}$. Note $R \in B$ and $S \notin B$ therefore $B \notin \mathcal{T}_{\mathrm{WL}}$. We can then conclude $\mathcal{T}_{\mathrm{WL}} \subset \mathcal{T}_{\mathrm{HD}}$. (QED)

□

Proposition 3. *Let $G = (V, E, f_G)$ be an attributed graph and $G' = (V', E', f_{G'})$ be an attributed subgraph of G. The following statements hold if $f_{G'}$ is a restriction of f_G to V':*

1. $\mathrm{dis}(G, G') = 0$,
2. *If additionally for any $n \in \mathbb{N}$, $G = K_n$ or $G = C_\ell$ for odd $\ell > 2 \in \mathbb{N}$, then* $\mathrm{dis}(G, G') > 0$.

Proof. We proof the statements one at a time:

1. This is a rather simple case and has a name. The function $\varphi' \colon G' \to G$ we are looking for is the *inclusion homomorphism*, where for every $v \in V'$ we have $\varphi'(v) = v$. φ' is a homomorphism since by definition of an induced subgraph if $\{u, v\} \in E'$ then it's also in E'. Note now that for φ', we have
$$\begin{aligned} \max_{v' \in V'} \|f_{G'}(v') - f_G(\varphi'(v'))\| &= \max_{v' \in V'} \|f_{G'}(v') - f_{G'}(\varphi'(v))\| \\ &= \max_{v' \in V'} \|f_{G'}(v') - f_G(v)\| \\ &= \max_{v' \in V'} 0 \\ &= 0 \end{aligned}$$
2. A retraction $r \colon G \to G'$ is a homomorphism that always satisfies $r(v) = v$. Note that to obtain a 0 we need a retraction $r \colon G \to G'$. Since for K_n and C_ℓ of odd ℓ only admit a retraction to themselves (Hell & Nesetril, 2004) and $G \neq G'$, the claim follows.

□

Lemma 2. *Let $G \in \mathcal{G}_n$ be an attributed graph. If $\mathcal{F} = \mathcal{G}_n$, then $\Gamma(G)_{\mathcal{F}}$ is (permutation) invariant.*

Proof. Let $\rho\colon V \to V$ be an isomorphism on G and let $H = (\rho(V), E', f_H)$ be the isomorphic graph resulting from applying the isomorphism, where E' is induced by permuting vertices V. We aim to show $\Gamma(G)_{\mathcal{F}} = \Gamma(H)_{\mathcal{F}}$. This equivalent to showing

$$[\mathrm{d_{HD}}(F, G)]_{F\in\mathcal{F}} = [\mathrm{d_{HD}}(F, H)]_{F\in\mathcal{F}}$$

Since H is isomorphic to G from Proposition 2 it follows that for any fixed $F \in \mathcal{F}$ we have $\mathrm{d_{HD}}(F, G) = \mathrm{d_{HD}}(F, H)$. With this fact it follows that $\Gamma(G)_{\mathcal{F}} = \Gamma(H)_{\mathcal{F}}$.

□

Proposition 4. *Let $G \in \mathcal{G}_n$ be an attributed graph. If $\mathcal{F} = \mathcal{G}_n$, then $\Gamma(G)_{\mathcal{F}}$ is complete up to the distinguishing power of f_G*

Proof. Let $G, H \in \mathcal{G}_n$ be a graph with attribute functions f_G and f_H respectively. We want to show that $G \simeq H$ if and only if $\Gamma(G)_{\mathcal{F}} = \Gamma(H)_{\mathcal{F}}$.Lemma 2 establishes the sufficiency of this condition. We need only to verify the necessity. The distinguishing power of f_G and f_H are determined by Ψ hence, assume Ψ distinguishes G, H, i.e., $\mathrm{d_H}(V(G), V(H)) > 0$, and $\Gamma(G)_{\mathcal{F}} = \Gamma(H)_{\mathcal{F}} \implies$ for any fixed graph $F \in \mathcal{F}$, $\mathrm{d_{HD}}(F, G) = \mathrm{d_{HD}}(F, H)$. Since $\mathcal{F} = \mathcal{G}_n$ it follows that $G, H \in \mathcal{F}$ thus for $F = H$ we have , $0 = \mathrm{d_{HD}}(H, H) = \mathrm{d_{HD}}(H, G)$. From Corollary 1 we have $G \simeq H$.

□

D Further Experimental Details

We provide brief details on our experiments, described in Section 5 in the main text.

Choices of $\mathcal{F}$. The family of cycle graphs $\mathcal{F}_n^C$ contains all cycle graphs up to C_n. Since the smallest cycle graph possible is C_3, then $|\mathcal{F}_n^C| = n - 2$.

Treewidth statistics. Table S.1 shows the distribution of treewidths along with the number of graphs and the maximum number of nodes, all elements required to measure the computational cost of running our method.

Table S.1: Distribution of treewidths (show in %) per dataset

Property	ZINC	NCI1	ogbg-molhiv
$\mathrm{tw}(G) = 1$	0.550	2.019	3.847
$\mathrm{tw}(G) = 2$	99.033	92.360	92.023
$\mathrm{tw}(G) = 3$	0.417	5.572	4.063
$\mathrm{tw}(G) = 4$	0.000	0.024	0.063
$\mathrm{tw}(G) = 5$	0.000	0.000	0.002
Max $\mathrm{tw}(G)$	3	4	5
Max $\lvert V(G)\rvert$	37	111	222

Algorithmic computations We provide the pseudocode for the implementation of the calculation of the homomorphism distortion. Primarily, this algorithm relies on a NICE-TREEWIDTHDECOMPOSITION, we then perform a dynamic programming technique from Bodlaender et al. (2013), originally proposed for the "Minimum Weighted Homomorphism Problem". the decomposition to build a DP table that will allow use to traverse the search space using a min-max approach. The solution relies on the construction of a DP table and a traversing it to solve a min-max problem by accounting for partial homomorphism from each of the vertices in a bag. Our implementation of the tree decomposition relies on finding an initial good approximation of treewidth with the NETWORKX.TREEWIDTH_MIN_DEGREE function from library NETWORKX(Hagberg & Conway, 2020), then refining said approximation to the optimal treewidth by an interpretation of the algorithmic procedure described in (Kloks, 1994) and finally performing Algorithm 1 process as a dynamic programming interpretation from Bodlaender et al. (2013).

Computational Complexity Given two graphs $G = (V, E), H = (W, E')$ and $|V| = n, |W| = m$, let $\mathrm{tw}(G) = k$. We present the computational complexity of finding out k, the tree decomposition for k and performing the DP. First, assume we do not know k, the algorithm in NETWORKX gives an approximate treewidth k' such that $k \leq k'$ in $\mathcal{O}(V^3)$ time. Then, Kloks (1994) gives a $\mathcal{O}(k \cdot n)$ time algorithm for constructing the final *nice* tree decomposition. Finally, the complexity of the DP is $\mathcal{O}(m^{k+1} \cdot k^2 \log k \cdot n)$ by (Bodlaender et al., 2013), where $\mathcal{O}(k \log k)$ is the amount of nodes in the tree, $\mathcal{O}(k \log k)$ is the sorting of processing the DP table at certain nodes an $m := q$ in the reference.

Algorithm 1 MINDISTORTIONHOMDP(G, H, f_G, f_H) — min–max distortion over $\mathrm{Hom}(G, H)$ via a nice tree decomposition of G

Require: Graphs G, H, features f_G, f_H

$\delta(v, w) \leftarrow \|f_G(v) - f_H(w)\|$ ▷ precompute all-pairs distance matrix
$N_H(w) \leftarrow \{w' : (w, w') \in E(H)\} \cup \{w : (w, w) \in E(H)\}$ for each $w \in V(H)$
$(T, \{X_t\}, r) \leftarrow$ NICETREEDECOMPOSITION(G) ▷ T is the tree, X_t is the bag at node t and r is the root
$c_t \leftarrow \{\}$ for all $t \in T$ ▷ DP tables of $\phi : X_t \to V(H)$, where each key is $\{(u, w) \mid u \in X_t \text{ and } \in V(H))$ such that there exist a homomorphism $\phi(u) = w\}$
for $t \in$ POSTORDER(T, r) **do** ▷ POSTORDER returns a stack that guarantees that when a node is seen all of its children have been already seen.
  if $\mathrm{kind}(t) =$ LEAF **then**
    $c_t(\varnothing) \leftarrow 0$
  else if $\mathrm{kind}(t) =$ INTRODUCE(v), child t' **then**
    for $\phi' \in \mathrm{dom}(c_{t'})$ **do** ▷ Each $\phi \in \mathrm{dom}(c_{t'})$ is a pair $(u \in X_t, w \in V(H))$ where $\phi(u) = w$ a homomorphism.
      for $w \in V(H)$ **do**
        ok $\leftarrow$ **true**
        for $u \in X_t \setminus \{v\}$ **do**
          if $(u, v) \in E(G)$ **and** $\phi'(u) \notin N_H(w)$ **then**
            ok $\leftarrow$ **false**
          end if
        end for
        if ok **then**
          $\phi \leftarrow \phi' \cup \{v \mapsto w\}$ ▷ This is the pair (u, w)
          $c_t(\phi) \leftarrow \min\big(c_t(\phi),\ \max(c_{t'}(\phi'),\ \delta(v, w))\big)$
        end if
      end for
    end for
  else if $\mathrm{kind}(t) =$ FORGET(v), child t' **then**
    for $\phi \in \mathrm{dom}(c_{t'})$ **do**
      $\phi_{-v} \leftarrow \phi \setminus \{v \mapsto \phi(v)\}$
      $c_t(\phi_{-v}) \leftarrow \min\big(c_t(\phi_{-v}),\ c_{t'}(\phi)\big)$
    end for
  else ▷ JOIN, children t_1, t_2
    for $\phi \in \mathrm{dom}(c_{t_1})$ **do**
      if $\phi \in \mathrm{dom}(c_{t_2})$ **then**
        $c_t(\phi) \leftarrow \max\big(c_{t_1}(\phi),\ c_{t_2}(\phi)\big)$ ▷ max since $V_{t_1} \cap V_{t_2} = \varnothing$
      end if
    end for
  end if
end for
if $c_r = \varnothing$ **then**
  return $+\infty$ ▷ no homomorphism exists
end if
return $\min_{\phi \in \mathrm{dom}(c_r)} c_r(\phi)$

Hardware setup. For BREC experiments and precomputation of the homomorphism encoding we use a compute node with ×2 AMD EPYC 9554 64-CORE. For model training, we use ×4 AMD EPYC 9654 96-CORE PROCESSOR and ×8 NVIDIA RTX PRO 6000 BLACKWELL SERVER EDITION with CUDA version 13.0, and Pytorch version 2.11.0.

Hyperparameters. Table S.2 shows the hyperparameters of the homomorphism distortion encoding. Table S.3 shows the fixed hyperparamters shared accross all the training setup, Table S.4 shows the hyperparamters for the GCN and GAT models when run on ZINC and Table S.5 shows the hyperparamters for the GIN model for ZINC, ogbg-molhiv and NCI1.

Table S.2: Encoding Hyperparamters.

Hyperparameter	Values
Walk length	$\{10, 75, 150, 200\}$
tw	$\{1, 2\}$
$\mathcal{F}^C$	$\{C_8, C_9\}$

Table S.3: Fixed hyperparameters.

Parameter	Value
Residual	True
Normalization	Batch
Initial Lr	0.001
Batch Size	128
Min. LR	10^{-5}

Table S.4: Hyperparameters for GCN and GAT on `ZINC`.

Parameter	GCN	GAT
Architecture		
Layer	16	16
Hidden Dimension	172	22
Hidden Output	–	176
Heads	–	8
Input Dim	28	28
Readout	Mean	Mean
Residual	True	True
Normalization	Batch	Batch
Dropout	$\{0.0, 0.5\}$	–
Training		
Max Epoch	1000	1000

Table S.5: Hyperparameters on GIN in `ZINC`, `ogbg-molhiv`, and `NCI1`.

Parameter	`ZINC`	`ogbg-molhiv`	`NCI1`
Architecture			
Layers	16	5	5
Hidden Dimension	124	300	64
Input Dimension	28	300	37
Readout	Sum	Sum	Sum
Dropout	$\{0.0, 0.5\}$	$\{0.0, 0.5\}$	$\{0.0, 0.5\}$
Training			
Max Epoch	1000	150	350