

HOMOMORPHISM COUNTS AS STRUCTURAL ENCODINGS FOR GRAPH LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Graph Transformers are popular neural networks that extend the well-known Transformer architecture to the graph domain. These architectures operate by applying self-attention on graph nodes and incorporating graph structure through the use of positional encodings (e.g., Laplacian positional encoding) or structural encodings (e.g., random-walk structural encoding). The quality of such encodings is critical, since they provide the necessary *graph inductive biases* to condition the model on graph structure. In this work, we propose *motif structural encoding* (MoSE) as a flexible and powerful structural encoding framework based on counting graph homomorphisms. Theoretically, we compare the expressive power of MoSE to random-walk structural encoding and relate both encodings to the expressive power of standard message passing neural networks. Empirically, we observe that MoSE outperforms other well-known positional and structural encodings across a range of architectures, and it achieves state-of-the-art performance on widely studied molecular property prediction datasets.

1 INTRODUCTION

Graph Neural Networks (GNNs) have been the prominent approach for graph machine learning in the last decade. Most conventional GNNs fall under the framework of Message Passing Neural Networks (MPNNs) and incorporate graph structure – the so-called *graph inductive bias* – in the learning and inference process by iteratively computing node-level representations using “messages” that are passed from neighboring nodes (Gilmer et al., 2017). More recently, Transformer architectures (Vaswani et al., 2017) have been applied to the graph domain and achieved impressive empirical performance (Rampásek et al., 2022; Ma et al., 2023; Ying et al., 2021), especially in molecular property prediction.

While MPNNs operate via exchanging messages between adjacent nodes in a graph, Transformers can be seen as a special type of GNN that operates on complete graphs. On the one hand, this allows for direct communication between all node pairs in a graph, regardless of whether there exists an edge between two nodes in the original input graph. On the other hand, since the node adjacency information is dropped, the Transformer lacks any “built-in” graph inductive bias. Instead, the underlying graph structure is usually provided by combining Transformer layers with message-passing layers (Yun et al., 2019; Rampásek et al., 2022; Bar-Shalom et al., 2024) or by incorporating additional pre-computed features that encode the topological context of each node. These additional features are referred to as *positional or structural encodings*. Common encodings include the Laplacian positional encoding (LapPE) (Dwivedi et al., 2023) and the random-walk structural encoding (RWSE) (Dwivedi et al., 2022). The quality of structural or positional encodings are reported as the key ingredients in the success of Transformers on graphs (Dwivedi et al., 2022; Rampásek et al., 2022).

Motivation. While empirical studies have observed the impact of structural or positional encodings on model performance (Dwivedi et al., 2022; Rampásek et al., 2022), our theoretical understanding of the expressive power of different positional encodings remains rather limited. This represents an important gap in the literature, especially since the expressive power of Transformer-based architectures heavily rely on the specific choice of the structural or positional encodings (Rosenbluth et al., 2024). Let us consider RWSE, which has been empirically reported as the most successful encoding on molecular benchmarks (Rampásek et al., 2022). We now illustrate a serious limitation of RWSE in its power of distinguishing nodes.

054 **How expressive is RWPE?** The expressive power
 055 of MPNNs is upper bounded by the *1-dimensional*
 056 *Weisfeiler Leman graph isomorphism test (1-*
 057 *WL)* (Xu et al., 2019; Morris et al., 2019). This
 058 means that the node invariants computed by
 059 MPNNs are at most as powerful as the node in-
 060 variants computed by 1-WL, and as a result, an
 061 MPNN can distinguish two nodes *only if* 1-WL
 062 can distinguish them. Some MPNN architectures
 063 such as Graph Isomorphism Networks (Xu et al.,
 064 2019) can also match this expressiveness bound
 065 and distinguish any pair of nodes that can be dis-
 066 tinguished by 1-WL. Can we relate the expressive
 067 power of RWPE to that of 1-WL in a similar fash-
 068 ion? We show that the node invariants given by
 069 RWSE are strictly weaker than the node invari-
 070 ants computed by 2-WL¹ (Proposition 4) and even
 071 incomparable to node invariants computed by 1-
 072 WL (Proposition 6). In fact, there are simple node
 073 pairs which can be distinguished by 1-WL but not
 074 by RWSE (and vice versa).

074 *Example 1.* Consider the nodes from the graphs G
 075 and H from Figure 1 (a). Observe that the nodes
 076 u_1 and u_2 can be distinguished by 1-WL. Inter-
 077 estingly, however, they are indistinguishable by
 078 RWSE for *all* lengths ℓ of the considered random
 079 walk. This observation also applies to the nodes v_1
 080 and v_2 . Moreover, this limitation is not merely of
 081 theoretical interest, as it readily applies to real-world
 082 molecules, where the use of RWSE is prominent. In
 083 practical terms, this implies that RWSE cannot dis-
 084 tinguish a Dinitrile group from a Morpholine group. \triangle


084 **Motif structural encodings as a flexible and powerful approach.** In this paper, we propose *motif*
 085 *structural encodings* (MoSE) as structural encodings that leverage homomorphism count vectors to
 086 capture the graph structure. Homomorphism counts have been investigated in the context of MPNNs
 087 to overcome well-known theoretical limitations in their expressivity (Barceló et al., 2021; Jin et al.,
 088 2024) with promising theoretical and empirical findings (see Section 2). Building on the existing
 089 literature, we show that MoSE is a parameterizable, flexible, and powerful alternative to existing
 090 positional and structural encoding schemes in the context of Transformers. In fact, we show that –
 091 unlike RWSE – MoSE cannot be located within the WL hierarchy (Proposition 3) and can provide
 092 non-trivial expressiveness gains as MoSE is more expressive than RWSE (Theorem 3).

092 *Example 2.* Consider the graphs G and H from our running example, and observe that H has a cycle
 093 of length six, while G has no cycles. The node-level homomorphism counts $\text{Hom}(\text{C}_6, G, v_1) \neq$
 094 $\text{Hom}(\text{C}_6, H, v_2)$ (introduced formally in Section 4.2) are not equal, which provides sufficient infor-
 095 mation to distinguish these nodes. Analogous statements can be made also for u_1 and u_2 . \triangle


096 **Contributions.** Our key contributions can be summarized as follows:

- 098 • We introduce MoSE and detail the expressiveness guarantees that can be achieved by using
 099 homomorphism counts as a graph inductive bias (Section 4.1).
- 100 • We compare MoSE to RWSE in terms of expressivity and related them to the expressive power of
 101 MPNNs through the well-known WL hierarchy (Section 4.2).
- 102 • We empirically validate the theoretical findings and demonstrate efficacy of MoSE on a wide
 103 variety of benchmarks (Section 5). We report state-of-the-art performance on the widely studied
 104 ZINC-12k dataset by simply replacing a random-walk based encoding with MoSE in an existing
 105 Transformer model. We also achieve state-of-the-art results on the QM9 molecular dataset for
 106 almost all properties. Finally, we experiment on a new synthetic dataset that further highlights the
 107 wide-ranging applicability of homomorphism counts as a graph inductive bias.

¹Throughout the paper, we refer to the folklore version of the WL hierarchy (Grohe & Otto, 2015).



(a) The nodes u_1 and u_2 (resp., v_1 and v_2) can be distinguished by 1-WL, but not by RWSE.



(b) A molecule from the ZINC dataset, with Carbon, Nitrogen, and Oxygen atoms. The nodes u_1 and u_2 are indistinguishable by RWSE and so are v_1 and v_2 .

Figure 1: RWSE is weaker than 1-WL in distinguishing nodes (a), which holds for *all* choices of the length of the random walk. This limitation surfaces in real-world molecular graphs (b).

2 RELATED WORK

(Graph) Transformers and Encodings. Vaswani et al. (2017) first highlighted the power of self-attention when they introduced their Transformer architecture. Dwivedi & Bresson (2020) generalized the concepts presented in Vaswani et al. (2017) to the graph domain. In recent years, several different graph transformer architectures have arisen. Yun et al. (2019) first combined message-passing layers with transformer layers, and other models have followed this approach (Rampásek et al., 2022; Bar-Shalom et al., 2024; Shirzad et al., 2023). Relatedly, Transformer architectures dedicated for knowledge graphs have also been developed (Zhang et al., 2023c; Liu et al., 2022). Shehzad et al. presents a survey of existing Graph Transformer models and how they compare to each other.

A key component to the success of graph transformer models is the use of effective graph inductive biases (Dwivedi et al., 2022). Laplacian positional encodings (LapPE), which were introduced by Dwivedi & Bresson (2020), are a popular choice, but they break invariance², as they are *not* invariant to the changes of signs and basis of eigenvectors (Lim et al., 2023b;a). Rampásek et al. (2022) presented a framework to build general and scalable Graph Transformer models, and performed ablations with a suite of different encoding types, including random-walk structural encodings (RWSE). Ying et al. (2021) introduced Graphormer, which uses an attention mechanism that is based on shortest-path (and related) distances in a graph. Ma et al. (2023) presented GRIT as a novel framework that relies solely on Relative Random Walk Probabilities (RRWPs) for inductive bias.

Expressive power of MPNNs. MPNNs capture the vast majority of GNNs (Gilmer et al., 2017) and their expressive power is upper bounded by 1-WL (Xu et al., 2019; Morris et al., 2019). This limitation has motivated a large body of work, including *higher-order* GNN architectures (Morris et al., 2019; Maron et al., 2019a;b; Keriven & Peyré, 2019), GNNs with unique node features (Loukas, 2020), or GNNs with random features (Sato et al., 2021; Abboud et al., 2021) to achieve higher expressive power. Expressive power of GNNs has also been evaluated in terms of their ability to detect graph components, such as biconnected components (Zhang et al., 2023b), and fully expressive architectures designed for special graph classes, e.g., for planar graphs (Dimitrov et al., 2023).

Our work is most closely related to approaches that design more expressive architectures by injecting the counts of certain substructures, widely referred to as motifs. Bouritsas et al. (2023) present an early work that enhances node feature encodings for GNNs by counting adjacent subgraph isomorphism of different patterns. Barceló et al. (2021) follow up on this work by proposing to instead count the number of homomorphism into a set of adjacent substructures. Several other works have identified homomorphism counts as a key tool for understanding the expressive power of MPNNs (Neuen, 2023; Lanzinger & Barceló, 2024; Wang & Zhang, 2024). Zhang et al. (2024) also recently proposed a new framework that characterizes the expressiveness of GNNs in terms of homomorphism counts. The method for obtaining these counts is outlined in (Curticapean et al., 2017) and (Bressan et al., 2023), who show when exactly homomorphisms can be counted efficiently in undirected and directed graphs, respectively. As such, Jin et al. (2024) built upon these theoretical observations to establish a general framework for using homomorphism counts in MPNNs. In particular, they show that homomorphism counts that capture certain “bases” of graph mappings provide a theoretically well-founded approach to enhance the expressivity of GNNs.

3 PRELIMINARIES

Graphs and Homomorphism Counts. An undirected *graph* is a set of *nodes* $V(G)$ and a set of *edges* $E(G) \subseteq V(G) \times V(G)$ which satisfy symmetry: $(u, v) \in E(G)$ if and only if $(v, u) \in E(G)$. Unless otherwise stated, we take all graphs to be finite, and we take all graphs to be *simple*: $(v, v) \notin E(G)$ for all $v \in V(G)$ and there exists at most one edge (up to edge symmetry) between any pair of nodes. We describe nodes $u, v \in V(G)$ as *adjacent* if $(u, v) \in E(G)$. The set of all nodes adjacent to $v \in V(G)$ is the *neighborhood* of v , notated $\mathcal{N}(v)$. The number of nodes $|\mathcal{N}(v)|$ is the *degree* of v , notated $d(v)$.

²By sacrificing invariance, it becomes much easier to design very expressive architectures. In fact, Graph Transformer architectures with LapPE are universal (Kreuzer et al., 2021), but so is a 2-layer MLP, or a single-layer GNN under LapPE (see Proposition 3.1 of Rosenbluth et al. (2024)).

A *homomorphism* from graph G to graph H is a function $f : V(G) \rightarrow V(H)$ such that $(u, v) \in E(G)$ implies $(f(u), f(v)) \in E(H)$. We say a homomorphism is an *isomorphism* if the function f is bijective, and if it additionally satisfies $(u, v) \in E(G)$ if and only if $(f(u), f(v)) \in E(H)$. In this case, we describe the graphs G and H as being *isomorphic*, denoted $G \cong H$. We write $\text{Hom}(G, H)$ for the number of homomorphisms from G to H . If we restrict to counting only homomorphisms which maps a particular node $g \in V(G)$ to the node $h \in V(H)$, we denote this *rooted homomorphism count* as $\text{Hom}_{g \rightarrow h}(G, H)$. Sometimes the choice of g is unimportant, so we use $\text{Hom}_{\rightarrow h}(G, H)$ to notate the rooted homomorphism count $\text{Hom}_{g \rightarrow h}(G, H)$ where g can be fixed as any arbitrary node in $V(G)$. We use $\text{Hom}(G, \cdot)$ to denote the function which maps any graph H to the integer $\text{Hom}(G, H)$, and we define rooted homomorphism count mappings analogously. If \mathcal{G} is a collection of graphs, then we treat $\text{Hom}(\mathcal{G}, \cdot)$ as a function that maps any input graph H to the ordered³ tuple (or integer vector) defined as $\text{Hom}(\mathcal{G}, H) = (\text{Hom}(G, H))_{G \in \mathcal{G}}$, and analogously for rooted counts.

Curticapean et al. (2017) describe *graph motif parameters* as functions $\Gamma(\cdot)$ that map graphs into \mathbb{Q} , such that there exists a *basis* of graphs $\text{Supp}(\Gamma) = \{F_i\}_{i=1}^\ell$ and corresponding coefficients $\{\alpha_i\}_{i=1}^\ell \subseteq \mathbb{Q} \setminus \{0\}$ which decompose Γ as a finite linear combination $\Gamma(\cdot) = \sum_{i=1}^\ell \alpha_i \cdot \text{Hom}(F_i, \cdot)$. Given G , the function which maps any H to the number of times that G appears as a subgraph⁴ of H is a graph motif parameter (Lovász, 2012) whose basis is known as the *Spasm*. Many mappings of interest can be described as graph motif parameters, so the theory of such bases bolsters homomorphism counts as broadly informative and powerful; we refer readers to Jin et al. (2024) for details.

Nguyen & Maehara (2020) present a generalization of homomorphism counts that allow for “vertex weighting”. For graph G , let $\omega : V(G) \rightarrow \mathbb{R}_{\geq 0}$ describe *node weights*. The *weighted homomorphism count* from a graph F into G is given as:

$$\omega\text{-Hom}(F, G) = \sum_{f \in \mathcal{H}} \left(\prod_{v \in V(F)} \omega(f(v)) \right)$$

where \mathcal{H} is the set of all homomorphisms from F to G . We define the node-rooted version $\omega\text{-Hom}_{u \rightarrow v}(F, G)$ by restricting \mathcal{H} to only those homomorphisms which map $u \in V(F)$ to $v \in V(G)$. The mapping $\omega\text{-Hom}_{\rightarrow(\cdot)}(F, \cdot)$ is defined analogously to the un-weighted case: the graph-node pair (G, v) gets mapped to $\omega\text{-Hom}_{u \rightarrow v}(F, G)$ where $u \in V(F)$ is fixed arbitrarily. Setting $\omega(v) = 1$ for all $v \in V(G)$ recovers the un-weighted count $\omega\text{-Hom}(F, G) = \text{Hom}(F, G)$, and similarly for the node-rooted version. Weighted homomorphism counts are well-studied in the context of their connection to graph isomorphism (Freedman et al., 2004; Cai & Govorov, 2021), and in the context of universal approximation capabilities as well as empirical performance of graph classifiers which use $\omega\text{-Hom}(F, \cdot)$ counts as a graph embedding (Nguyen & Maehara, 2020).

Expressive Power. A *graph invariant* is a function $\xi(\cdot)$ which acts on graphs, satisfying $\xi(G) = \xi(H)$ whenever $G \cong H$ for all graphs G and H . For indexed families of graph invariants A_i and B_j (where $i \in I$ and $j \in J$ respectively), we define the *expressivity* relation $A \preceq B$ to denote: for any choice of $i \in I$, there exists a choice of $j \in J$ such that $B_j(G) = B_j(H)$ implies that $A_i(G) = A_i(H)$ for all G and H . If $A \preceq B$ and $B \preceq A$, then we say that $A \simeq B$. We write $A \prec B$ when B has strictly greater expressive power, $A \preceq B$ but $A \not\simeq B$. When we reference \preceq for some fixed graph invariant $\xi(\cdot)$, we interpret ξ as a family which contains only one graph invariant (hence the indexing is trivial). Noting that a GNN architecture can be treated as a family of graph invariants indexed by the model weights, the relation \preceq generalises common notions the graph-distinguishability expressive power for GNNs (Zhang et al., 2023a). As we will see in Section 4, the \preceq relation naturally extends to expressivity comparisons between structural/positional encoding schemes as well.

MPNNs and Weisfeiler-Lehman Tests. Given graph G , the 1-WL test induces the node coloring $c^{(t)} : V(G) \rightarrow \Sigma$ for all t up until some termination step T , at which point the graph-level 1-WL label is defined as the multiset of final node colors $1\text{-WL}(G) = \{\!\!\{ c^{(T)}(v) : v \in V(G) \}\!\!\}$. Graphs G and H are considered *1-WL indistinguishable* they have identical graph labels $1\text{-WL}(G) = 1\text{-WL}(H)$. Crucially, it holds that $\text{MPNN} \preceq 1\text{-WL}$ where we treat “MPNN” as a family of graph invariants

³If the graphs in \mathcal{G} are not ordered initially, we arbitrarily fix some indexing of the elements in \mathcal{G} so that they become ordered.

⁴We count the number of times one can find a graph H' with $V(H') \subseteq V(H)$ and $E(H') \subseteq E(H)$ s.t. $G \cong H'$.

indexed by both the choice of message-passing architecture and the model parameters, and we treat 1-WL as a singleton family containing only the graph invariant $G \mapsto 1\text{-WL}(G)$ (Xu et al., 2019).

We extend the Weisfeiler-Lehman test to “higher dimensions” by coloring k -tuples of nodes. The k -WL test induces node tuple coloring $c^{(t)} : V(G)^k \rightarrow \Sigma$ for $t = 1, \dots, T$ iterations, and then it aggregates a final graph-level k -WL label analogously to 1-WL (Huang & Villar, 2021). In this work, we refer to the *folklore k -WL test*, which is provably equivalent to the alternative *olibiuous k -WL* formulation up to a re-indexing of k (Grohe & Otto, 2015). It has been shown that the k -WL tests form a well-ordered hierarchy of expressive power, where $k\text{-WL} \prec (k+1)\text{-WL}$ (Cai et al., 1989). Here, we are treating k -WL as a singleton family (trivial indexing) for any particular choice of k .

Self-Attention and Positional/Structural Encoding. Central to any Transformer architecture is the *self-attention* mechanism. When applied to graphs, a single “head” of self-attention learns a t^{th} -layer hidden representation $\mathbf{h}_v^{(t)} \in \mathbb{R}^d$ of every node $v \in V(G)$ by taking the weighted sum $\mathbf{h}_v^{(t)} = \phi^{(t)}(\sum_{u \in V(G)} \alpha_{v,u}^{(t)} \mathbf{h}_u^{(t-1)})$ where the *attention coefficients* are $\alpha_{v,u}^{(t)} = \psi^{(t)}(\mathbf{h}_v^{(t-1)}, \mathbf{h}_u^{(t-1)})$. Here, $\phi^{(t)}$ and $\psi^{(t)}$ are learnably-parameterized functions. The most common implementation of $\phi^{(t)}$ and $\psi^{(t)}$ is *scaled dot product attention*, as defined by Vaswani et al. (2017); although some Transformers such as GRIT (Ma et al., 2023) deviate from this. In most Transformers, we utilize multiple attention heads in parallel (whose outputs are concatenated at each layer), and we interweave self-attention layers with fully-connected feed-forward layers.

Since basic self-attention does not receive the adjacency matrix as an input, many graph Transformer choose to inject a graph’s structural information into the model inputs by way of node positional/structural encoding. For example, the widely used *k -length random walk structural encoding* (RWSE $_k$) assigns to each node its corresponding diagonal entry of the degree normalized adjacency matrix⁵ $(\mathbf{D}^{-1} \mathbf{A})^i$ for all powers $i = 1, \dots, k$ (Rampásek et al., 2022; Dwivedi et al., 2022; Ma et al., 2023). Each node’s RWSE $_k$ vector is then concatenated or added to the node’s initial feature vector, and the resulting node embedding is passed into the Transformer as an input. More broadly, we can define *node positional or structural encoding* to refer to any isomorphism-invariant mapping pe which takes in a node-graph pair (v, G) and outputs a vector label $pe(v, G) \in \mathbb{R}^d$. Then, the graph-level PE label is defined as the multiset of node-level labels $\text{PE}(G) = \{\{pe(v, G) : v \in V(G)\}\}$. In this way, we can treat RWSE $_k$ as a family of graph invariants (mapping into multisets with elements in \mathbb{R}^d) indexed by k . Hence, we can compare the graph-level expressivity of RWSE $_k$ and any other general positional/structural encoding scheme PE under \preceq .

4 HOMOMORPHISM COUNTS AS A GRAPH INDUCTIVE BIAS

In this section we formally define motif structural encodings (MoSE). Just as with other vertex level structural or positional encodings, our encoding provides a way to encode the structural characteristics of each graph vertex in terms of a numerical vector. Such encodings can then provide graph inductive bias to machine learning architectures, such as Transformers, that cannot (or only in a limited fashion) take graph structure into account. Additional proof details are provided in Appendix A.

4.1 MOTIF STRUCTURAL ENCODING (MOSE)

The *motif structural encoding* (MoSE) scheme is parameterized by a choice of finite⁶ pattern graph family $\mathcal{G} = \{G_1, \dots, G_d\}$, as well as a choice of node weighting scheme ω which sends any (v, H) to a non-negative weight $\omega(v, H) \in \mathbb{R}_{\geq 0}$. For each node v of graph H , the node-level MoSE $_{\mathcal{G}, \omega}$ label of v is:

$$\text{MoSE}_{\mathcal{G}, \omega}(v, H) = \left[\omega\text{-Hom}_{\rightarrow v}(G_i, H) \right]_{i=1}^d \in \mathbb{R}^d \quad (1)$$

We will notate $e_v^{\mathcal{G}} := \text{MoSE}_{\mathcal{G}, \omega}(v, H)$ for shorthand when ω and H are either clear from context, or any arbitrary choice can be made. Unless otherwise specified, we use the constant ω weighting which sends all nodes in all graphs to 1, aligning MoSE $_{\mathcal{G}, \omega}$ with the un-weighted homomorphism counts used in previous works of Jin et al. (2024); Barceló et al. (2021). In this case, we omit ω from our

⁵Here, \mathbf{D} notates the diagonal degree matrix, and \mathbf{A} the adjacency matrix.

⁶We require that \mathcal{G} be a finite set, and that each graph within \mathcal{G} have finite size.

notation, using $\text{MoSE}_{\mathcal{G}}$ to denote our encoding. When we reference $\text{MoSE}_{\mathcal{G},\omega}$ at the graph-level, we mean the multiset of node-labels $\text{MoSE}_{\mathcal{G},\omega}(H) = \{\{e_v^{\mathcal{G}} : v \in V(H)\}\}$. It holds that $\text{MoSE}_{\mathcal{G},\omega}(\cdot)$ is a graph invariant because for any two graphs H and F with isomorphism $\iota : V(H) \rightarrow V(F)$, we have $e_v^{\mathcal{G}} = e_{\iota(v)}^{\mathcal{G}}$ for every vertex $v \in V(H)$ (Nguyen & Maehara, 2020).

We stated above that MoSE offers a very flexible solution to structural encodings. This flexibility comes from the two key parameters, the graphs \mathcal{G} , and the vertex weight function. Both of these parameters can be adapted to fit the problem domain and architecture as desired. The choice of \mathcal{G} can be informed in precise terms by desired levels of expressiveness as shown below in Proposition 3. Furthermore, the choice of \mathcal{G} can build on a range of empirical studies on structural information in MPNNs (see e.g., Barceló et al. (2021); Jin et al. (2024); Wang & Zhang (2024); Bouritsas et al. (2023)). Additional choice of a non-trivial weight function ω adds further power and flexibility. In particular we show that even with the simple weight function that maps nodes to the reciprocal of their degree is enough to exactly express RWSE in terms of MoSE (Proposition 5).

In the following we will provide insight into the expressivity of MoSE by leveraging connections to the established relations between homomorphism counts and MPNNs. For transformer architectures, similar frameworks of expressivity are not yet established. Recent work by Rosenbluth et al. (2024) shows that transformer architectures do not inherently contribute to expressivity, but that positional and structural encodings are the key ingredient to the expressivity of these architectures. We therefore study the expressivity of the positional encodings themselves, which in turn also translate to models that combine an architecture with MoSE. In particular, any typical Graph Transformer architecture with MoSE will naturally be at least as expressive as MoSE, and thus inherit all lower bounds from our analysis.

Our first two observations establish that MoSE is in a sense incomparable to the WL-hierarchy. For every fixed set of graphs \mathcal{G} that induces an encoding, we can provide an upper bound in terms of k -WL for some k dependent on \mathcal{G} . However, at the same time, the expressiveness of even MoSE encodings induced by a single graph cannot be confined within a level of the WL hierarchy (Proposition 2).

Proposition 1. *Let \mathcal{G} be a finite set of graphs and let k be the maximum treewidth of a graph in \mathcal{G} . Then $\text{MoSE}_{\mathcal{G}}$ is at most as distinguishing as k -WL. That is, $\text{MoSE}_{\mathcal{G}} \preceq k\text{-WL}$ ⁷.*

Although we cannot distinguish *all* graphs distinguished by k -WL using $\text{MoSE}_{\mathcal{G}}$ with a finite \mathcal{G} , we can distinguish any particular pair of graphs distinguished by k -WL. In fact, we only need a single graph in \mathcal{G} in order to do this.

Proposition 2. *For $k \geq 1$, let G and H be two graphs that are equivalent under k -WL. Then there exists a graph F such that $\text{MoSE}_{\{F\}}$ distinguishes G and H .*

Lanzinger & Barceló (2024) and Jin et al. (2024) studied the expressivity of MPNNs with respect to functions that map graphs to numbers. For a large class of these, so-called *graph motif parameters* (we ask the reader to refer to Jin et al. (2024) for details), the distinguishing power of such a function relates closely to homomorphism counts from its homomorphism basis. Building on these results we establish a very broad lower bounds for Transformer architectures with MoSE that provide a practical basis for choice of \mathcal{G} .

Proposition 3. *Let f be a graph motif parameter with basis $\mathcal{G}_f \subseteq \mathcal{G}$. Then $\text{MoSE}_{\mathcal{G}}$ is at least as distinguishing as f . That is, for any choice of graph motif parameter f , we have $f \preceq \text{MoSE}$.*

4.2 COMPARING THE EXPRESSIVITY OF MOSE TO RWSE

From above we see MoSE, is on the one hand closely aligned to the k -WL hierarchy, and we can use this alignment to inform the choice of \mathcal{G} . But on the other hand, it is not possible to contain MoSE within the hierarchy. In the following we show that this is not the case for RWSE. In fact, the expressiveness of RWSE (regardless of length parameter) is fully contained within 2-WL.

Proposition 4. *For every $\ell \geq 2$, any graph which can be distinguished by RWSE_{ℓ} can be distinguished by 2-WL.*

This in itself already has wide-ranging consequences, e.g., it is known that 2-WL cannot distinguish strongly regular graphs (see e.g., Fuhlbrück et al. (2021)), and therefore neither can RWSE. For

⁷With respect to the definition of \preceq , both sides in this case represent singleton families of graph invariants.

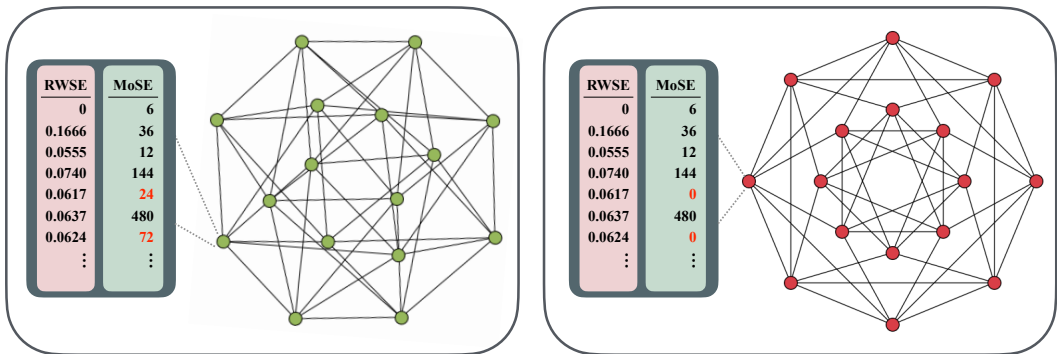


Figure 2: The 4×4 Rook’s Graph (left) and the Shrikhande Graph (right) are non-isomorphic strongly regular graphs that have the same parameters. RWSE produces the same vector on every vertex of both graphs whereas a simple construction of MoSE using $\text{Spasm}(C_7 \cup C_8)$ homomorphism counts easily distinguishes the two graphs.

MoSE there exist no such limitations (cf. Figure 2), beyond those dependent on choice of \mathcal{G} stated in Proposition 1.

While MoSE cannot fully capture 2-WL, it is in fact possible to express RWSE_ℓ for every $\ell \geq 2$ as a special case of MoSE. In contrast to previous results, our result here requires a vertex-weighting scheme that is not constant. However, even straightforward vertex-weighting by degrees – which comes at no cost in practice – is enough for our proof. We note that \mathcal{G} consisting of only ℓ graphs is sufficient to do so (see Appendix A for details).

Proposition 5. *For any $\ell \in \mathbb{N}$, there exists a finite family of graphs \mathcal{G} and a vertex-weighting scheme ω such that the node-level $\text{RWSE}_\ell(v, H)$ label is uniquely determined by $\text{MoSE}_{\mathcal{G}, \omega}(v, H)$ for any node v in any graph H . In particular, this holds already for the vertex-weighting scheme $\omega : v \mapsto 1/\text{degree}(v)$.*

Combining Proposition 4 and Proposition 5 immediately shows $\text{RWSE} \preceq \text{MoSE}$. The example from Figure 2 demonstrates that this inclusion is in fact strict.

Theorem 3. *Motif structural encoding is strictly more expressive than random walk structural encoding: $\text{RWSE} \prec \text{MoSE}$.*

Finally, we note that on node-level there are even cases where RWSE is weaker than 1-WL. Specifically, there are nodes that for which 1-WL assigns different labels, but RWSE_ℓ assigns the same vectors to both nodes for every $\ell \geq 1$ (cf., Figure 1).

Proposition 6. *RWSE is incomparable to 1-WL in terms of node-level expressiveness.*

5 EXPERIMENTS

We empirically show the efficacy of using MoSE as a graph inductive bias on several real-world datasets across a range of models. First, we provide an in-depth comparison of MoSE with other positional and structural encoding methods on the ZINC molecular dataset (Irwin et al., 2012; Dwivedi et al., 2023) in Section 5.1. Next, we specifically contrast the performance of MoSE and RWSE on the QM9 dataset (Wu et al., 2018) in Section 5.2. Finally, we conduct a synthetic experiment to highlight the flexibility of MoSE across different domains in Section 5.3. In an attempt to establish the viability of MoSE in the most general sense we avoid the use of non-trivial vertex-weighting schemes in our experiments (that is, ω always maps every vertex to 1). Full experimental details, including code instructions for reproduction are provided in Appendix B.

5.1 GRAPH REGRESSION ON ZINC

ZINC (Irwin et al., 2012; Dwivedi et al., 2023) is a molecular dataset, where the goal is to predict the constrained solubility of a given molecule. Each graph in the dataset represents an individual molecule,

Table 1: We report the mean absolute error (MAE) for GPS with various different encoding methods on the ZINC-12k dataset (with edge features). MoSE outperforms several popular positional encoding methods.

Encoding Type	MAE ↓	Rel. Change
<i>none</i>	0.113±0.007	± 0 %
PEG ^{LapEig}	0.161±0.006	+ 29.81 %
LapPE	0.116±0.009	+ 2.59 %
SignNet ^{MLP}	0.090±0.007	- 25.56 %
SignNet ^{DeepSets}	0.079±0.006	-43.04%
RWSE	0.070±0.002	-61.43%
MoSE (<i>ours</i>)	0.062 ±0.002	-82.26%

Table 2: We achieve state-of-the-art results on ZINC-12k when swapping out the random walk encoding from GRIT to MoSE. The MP column denotes whether or not the architecture contains a message passing component.

Model	MAE ↓	MP
GSN	0.101±0.010	✓
CIN (small)	0.094±0.004	✓
GPS	0.070±0.002	✓
Subgraphormer+PE	0.063±0.001	✓
GT	0.226±0.014	✗
GRIT+RRWP	0.059±0.002	✗
GRIT+MoSE (<i>ours</i>)	0.056 ±0.001	✗

Table 3: We report mean absolute error (MAE) for graph regression on ZINC-12K (with edge features) across several models. MoSE yields substantial improvements over RWSE for every architecture. The best feature encoding is highlighted in **bold**.

Model	Encoding Type		
	<i>none</i>	RWSE	MoSE
MLP-E	0.606±0.002	0.361±0.010	0.347 ±0.003
GIN-E	0.243±0.006	0.122±0.003	0.118 ±0.007
GIN-E+VN	0.151±0.006	0.085±0.003	0.068 ±0.004
GT-E	0.195±0.025	0.104±0.025	0.089 ±0.018
GPS	0.119±0.011	0.070±0.002	0.062 ±0.002

where the node features denote atom type, and edge features denote bond type. In accordance with Dwivedi et al., we use a subset of the ZINC dataset with edge features that contains 12,000 molecules and constrain all our models to roughly 500k parameters.

Experimental Setup. We begin by selecting the GPS (Rampásek et al., 2022) as our reference model and compare how MoSE performs against other popular encoding methods, such as RWSE, LapPE, etc. Then, we combine MoSE with GRIT (Ma et al., 2023), a recent graph transformer model that incorporates graph inductive biases without using message passing. GRIT utilizes relative random walk (RRWP) positional encodings at both the node representation level as well as the node-pair representation level. In our formulation, GRIT+MoSE, we remove all RRWP encodings for both the node and node-pair representations. For ZINC, we construct MoSE using the Spasm($C_7 \cup C_8$) homomorphism counts in accordance with Jin et al. (2024). All other experimental details can be found in Appendix B.1.

Results. We compare MoSE to other prominent positional and structural encoding methods. Our results are presented in Table 1 with baseline results for all cases except MoSE taken from Rampásek et al. (2022). We see that MoSE significantly outperforms the other encoding types, especially the spectral methods, such as PEG^{LapEig} (Wang et al., 2022) and LapPE (Dwivedi & Bresson, 2020). When compared to RWSE, the most prominent structural-based encoding method, MoSE still achieves over 10% relative improvement.

Furthermore, using MoSE in conjunction with the GRIT architecture (Ma et al., 2023) yields state of the art results on ZINC, as detailed in Table 2. Recall that when we substituted MoSE in place of the RRWP encoding for GRIT, we only did so at the node level. This suggests that even without explicit node-pair encodings, MoSE still captures a meaningful structural information about the graph. Not only does MoSE surpass the expressivity of RWSE (Section 4.2), but the information encoded in MoSE can aid the model in recovering specific substructure information (Jin et al., 2024).

Comparing MoSE to RWSE across Multiple Architectures. Because Rampásek et al. find that RWSE performs better for molecular datasets, we provide a detailed comparison of how MoSE

Table 4: We report MAE results for GPS with various feature enhancements on the QM9 dataset. The best model is highlighted in red, the second best is blue, and the third best is olive.

Property	R-GIN	R-GIN+FA	R-SPN	E-BasePlanE	GPS	
					RWSE	MoSE (<i>ours</i>)
mu	2.64±0.11	2.54±0.09	2.21±0.21	1.97±0.03	1.47±0.02	1.43±0.02
alpha	4.67±0.52	2.28±0.04	1.66±0.06	1.63±0.01	1.52±0.27	1.48±0.16
HOMO	1.42±0.01	1.26±0.02	1.20±0.08	1.15±0.01	0.91±0.01	0.91±0.01
LUMO	1.50±0.09	1.34±0.04	1.20±0.06	1.06±0.02	0.90±0.06	0.86±0.01
gap	2.27±0.09	1.96±0.04	1.77±0.06	1.73±0.02	1.47±0.02	1.45±0.02
R2	15.63±1.40	12.61±0.37	10.63±1.01	10.53±0.55	6.11±0.16	6.22±0.19
ZPVE	12.93±1.81	5.03±0.36	2.58±0.13	2.81±0.16	2.63±0.44	2.43±0.27
U0	5.88±1.01	2.21±0.12	0.89±0.05	0.95±0.04	0.83±0.17	0.85±0.08
U	18.71±23.36	2.32±0.18	0.93±0.03	0.94±0.04	0.83±0.15	0.75±0.03
H	5.62±0.81	2.26±0.19	0.92±0.03	0.92±0.04	0.86±0.15	0.83±0.09
G	5.38±0.75	2.04±0.24	0.83±0.05	0.88±0.04	0.83±0.12	0.80±0.14
Cv	3.53±0.37	1.86±0.03	1.23±0.06	1.20±0.06	1.25±0.05	1.02±0.04
Omega	1.05±0.11	0.80±0.04	0.52±0.02	0.45±0.01	0.39±0.02	0.38±0.01

performs against RWSE across a range of architectures. We select a simple 4-layer multi-layer perceptron (MLP-E), the self-attention Graph Transformer (GT-E) from (Dwivedi & Bresson, 2020), and GPS (Rampásek et al., 2022) as our models. All models are adapted to account for edge features. GRIT is omitted since its node-pair encoding mechanism does not make it directly comparable to the other models. For reference GNN models, we report results for GIN-E and its extension GIN-E+VN including a virtual node (Hu et al., 2020). As before, MoSE uses the set of homomorphism counts for Spasm($C_7 \cup C_8$). For RWSE, we follow the protocol from Rampásek et al. and set the random-walk length to 20.

Results. We see from Table 3 that MoSE performs better than RWSE across all models. Even though the graphs in ZINC are relatively small and a random-walk length of 20 already traverses most of the graph, our experiments show that the structural information provided by MoSE empirically performs better. We also note that GIN-E+VN becomes competitive with the state-of-the-art when using MoSE. This is consistent with our theoretical findings and also align well with recent empirical findings that suggest the use of virtual nodes can substantially improve model performance of MPNNs (Tönshoff et al., 2024).

5.2 GRAPH REGRESSION ON QM9

QM9 is a real-world molecular dataset that contains over 130,000 graphs (Wu et al., 2018; Brockschmidt, 2020). The node features include the atom’s type and other descriptive features. Unlike ZINC, where there is only one regression target, QM9 presents 13 different quantum chemical properties to regress over, making it a much more robust benchmark.

Experimental Setup. We select GPS (Rampásek et al., 2022) as our base model due to its modularity. Again, we compare the performance of GPS+RWSE to the performance of GPS+MoSE. We follow the protocol from Jin et al. and construct MoSE using the set of homomorphism counts on connected graphs with at most 5 vertices and the 6-cycle. This is because there are several regression targets for QM9, so GPS will benefit from a graph inductive bias that can account for a more diverse set of motifs (Proposition 3). We keep a random-walk length of 20 for RWSE to align with the hyperparameter setting for molecular datasets outlined by Rampásek et al.. Finally, we provide several leading GNN models on the QM9 dataset for comparison (Brockschmidt, 2020; Alon & Yahav, 2021; Abboud et al., 2022; Dimitrov et al., 2023).

Results. From Table 4, it is clear that MoSE outperforms RWSE consistently across multiple different regression targets. Specifically, GPS+MoSE achieves the same or a lower MAE on 11 of the 13 molecular properties. It also comfortably outperforms the leading GNN models, including R-SPN (Abboud et al., 2022) and E-BasePlane (Dimitrov et al., 2023). In contrast, GPS+RWSE is occasionally outperformed by these GNN architectures, despite being computationally more

486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539




Figure 3: Plot of the distribution of the fractional domination numbers in our synthetic dataset.




Figure 4: We report the MAE performance of different encodings with a 4-layer MLP and a GT on our synthetic dataset.

demanding than the GNN baselines. These findings further support our theoretical results and show that MoSE is a very promising alternative structural encoding scheme.

5.3 SYNTHETIC DATASET: PREDICTING FRACTIONAL DOMINATION NUMBER

To supplement our real-world dataset examples, we also generate a new synthetic dataset where the task is to predict the *fractional dominance number* of a graph. The dataset is comprised of 10,000 randomly-generated Erdős-Rényi graphs of various densities (Erdős & Rényi, 1959). The exact range of the number of vertices and density is provided in Appendix B.3.

We select the fractional dominance number of a graph as our target of interest for two reasons: first, the property is inherently global and relies on complex long-range interactions. This is unlike other popularly used metrics for synthetic experiments, e.g., clustering coefficients, which are aggregations of local properties. Second, the distribution of this parameter over our set of random graphs is complex (Figure 3), making the task more challenging. Furthermore, (fractional) dominating sets are important for a range of applications, such as optimising network coverage (Kutten & Peleg, 1995).

Experimental Setup. We focus on measuring the practical performance of MoSE compared to other positional and structural encoding methods. For this we select a 4-layer MLP and a self-attention Graph Transformer (Dwivedi & Bresson, 2020) as our reference models. Since neither model contains a message passing component, they both are considered architectures without inherent graph inductive biases. We use the homomorphism counts for $\text{Spasm}(C_7 \cup C_8)$ for MoSE, use a random walk length of 20 for RWSE, and set the dimension of LapPE to 8. Further experimental details are in Appendix B.3.

Results. Figure 4 shows that across both models, MoSE performs the best compared to the other graph inductive biases. This suggests that MoSE is able to capture more complex graph information in the broader context. We also note that MLP+LapPE performs significantly worse than all other configurations. This indicates that the self-attention mechanism in GT is a key component to allowing the model to leverage the spectral graph information provided by LapPE.

6 DISCUSSION AND FUTURE WORK

In this work, we propose motif structural encodings as a flexible new graph inductive bias that is based on counting graph homomorphisms. MoSE is supported by expressiveness guarantees, which transfer to architectures that use MoSE as their graph inductive bias. In particular, we show that MoSE is a more expressive encoding than RWSE and relate these encoding schemes to the expressive power of MPNNs and more generally to the WL hierarchy. We empirically validate that the theoretical advantages of MoSE also translate to performance gains across a set of real-world and synthetic datasets. Although our work takes a closer look at the expressivity of different graph inductive biases, there remain several open questions in this area of characterizing different encoding methods. Furthermore, we have yet to explore the full potential of the vertex-weighting parameter of our encoding scheme.

REFERENCES

- 540
541
542 Ralph Abboud, İsmail İlkan Ceylan, Martin Grohe, and Thomas Lukasiewicz. The surprising power
543 of graph neural networks with random node initialization. In *IJCAI*, 2021.
- 544 Ralph Abboud, Radoslav Dimitrov, and İsmail İlkan Ceylan. Shortest path networks for graph
545 property prediction. In *LoG*, 2022.
- 546
547 Uri Alon and Eran Yahav. On the bottleneck of graph neural networks and its practical implications.
548 In *ICLR*, 2021.
- 549 Guy Bar-Shalom, Beatrice Bevilacqua, and Haggai Maron. Subgraphormer: Unifying subgraph gnns
550 and graph transformers via graph products. In *NeurIPS*, 2024.
- 551
552 Pablo Barceló, Floris Geerts, Juan L. Reutter, and Maksimilian Ryschkov. Graph neural networks
553 with local graph parameters. In *NeurIPS*, 2021.
- 554
555 Giorgos Bouritsas, Fabrizio Frasca, Stefanos Zafeiriou, and Michael M. Bronstein. Improving graph
556 neural network expressivity via subgraph isomorphism counting. *IEEE Trans. Pattern Anal. Mach.*
557 *Intell.*, 45(1):657–668, 2023.
- 558 Marco Bressan, Matthias Lanzinger, and Marc Roth. The complexity of pattern counting in directed
559 graphs, parameterised by the outdegree. In *STOC*, 2023.
- 560
561 Marc Brockschmidt. GNN-FiLM: Graph neural networks with feature-wise linear modulation. In
562 *ICML*, 2020.
- 563 Jin-Yi Cai and Artem Gonorov. On a theorem of lovász that $\text{hom}(\cdot, h)$ determines the isomorphism
564 type of h , 2021. URL <https://arxiv.org/abs/1909.03693>.
- 565
566 Jin-Yi Cai, Martin Fürer, and Neil Immerman. An optimal lower bound on the number of variables for
567 graph identification. *30th Annual Symposium on Foundations of Computer Science*, pp. 612–617,
568 1989. URL <https://api.semanticscholar.org/CorpusID:261288588>.
- 569 Radu Curticapean, Holger Dell, and Dániel Marx. Homomorphisms are a good basis for counting
570 small subgraphs. In *STOC*, 2017.
- 571
572 Radoslav Dimitrov, Zeyang Zhao, Ralph Abboud, and İsmail İlkan Ceylan. PlanE: representation
573 learning over planar graphs. In *NeurIPS*, 2023.
- 574 Zdenek Dvorák. On recognizing graphs by numbers of homomorphisms. *J. Graph Theory*, 64(4):
575 330–342, 2010.
- 576
577 Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs.
578 *arXiv preprint arXiv:2012.09699*, 2020.
- 579 Vijay Prakash Dwivedi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Graph
580 neural networks with learnable structural and positional representations. In *ICLR*, 2022.
- 581
582 Vijay Prakash Dwivedi, Chaitanya K Joshi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and
583 Xavier Bresson. Benchmarking graph neural networks. *Journal of Machine Learning Research*, 24
584 (43):1–48, 2023.
- 585
586 P Erdős and A Rényi. On random graphs i. *Publ. math. debrecen*, 6(290-297):18, 1959.
- 587
588 M. Freedman, L. Lovasz, and A. Schrijver. Reflection positivity, rank connectivity, and homomor-
589 phism of graphs, 2004. URL <https://arxiv.org/abs/math/0404468>.
- 590
591 Frank Fuhlbrück, Johannes Köbler, Ilia Ponomarenko, and Oleg Verbitsky. The weisfeiler-leman
592 algorithm and recognition of graph properties. *Theor. Comput. Sci.*, 895:96–114, 2021. doi:
593 10.1016/J.TCS.2021.09.033. URL <https://doi.org/10.1016/j.tcs.2021.09.033>.
- Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural
message passing for quantum chemistry. In *ICML*, 2017.

- 594 Martin Grohe and Martin Otto. Pebble games and linear equations. *The Journal of Symbolic Logic*,
595 80(3):797–844, 2015.
- 596
- 597 Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta,
598 and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. In *NeurIPS*,
599 2020.
- 600 Ningyuan Teresa Huang and Soledad Villar. A short tutorial on the weisfeiler-lehman test and
601 its variants. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and*
602 *Signal Processing (ICASSP)*. IEEE, June 2021. doi: 10.1109/icassp39728.2021.9413523. URL
603 <http://dx.doi.org/10.1109/ICASSP39728.2021.9413523>.
- 604
- 605 John J Irwin, Teague Sterling, Michael M Mysinger, Erin S Bolstad, and Ryan G Coleman. Zinc: a
606 free tool to discover chemistry for biology. *Journal of chemical information and modeling*, 52(7):
607 1757–1768, 2012.
- 608 Emily Jin, Michael M. Bronstein, İsmail İlkan Ceylan, and Matthias Lanzinger. Homomorphism
609 counts for graph neural networks: All about that basis. In *ICML*, 2024.
- 610
- 611 Nicolas Keriven and Gabriel Peyré. Universal invariant and equivariant graph neural networks. In
612 *NeurIPS*, 2019.
- 613 Devin Kreuzer, Dominique Beaini, Will Hamilton, Vincent Létourneau, and Prudencio Tossou.
614 Rethinking graph transformers with spectral attention. In M. Ranzato, A. Beygelz-
615 imer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural In-*
616 *formation Processing Systems*, volume 34, pp. 21618–21629. Curran Associates, Inc.,
617 2021. URL [https://proceedings.neurips.cc/paper_files/paper/2021/](https://proceedings.neurips.cc/paper_files/paper/2021/file/b4fd1d2cb085390fbbadae65e07876a7-Paper.pdf)
618 [file/b4fd1d2cb085390fbbadae65e07876a7-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/b4fd1d2cb085390fbbadae65e07876a7-Paper.pdf).
- 619 Shay Kutten and David Peleg. Fast distributed construction of k-dominating sets and applications. In
620 *PODC*, 1995.
- 621
- 622 Matthias Lanzinger and Pablo Barceló. On the power of the weisfeiler-leman test for graph motif
623 parameters. In *ICLR*, 2024.
- 624
- 625 Moritz Lichter, Iliia Ponomarenko, and Pascal Schweitzer. Walk refinement, walk logic, and the
626 iteration number of the weisfeiler-leman algorithm. In *LICS*, 2019.
- 627 Derek Lim, Joshua Robinson, Stefanie Jegelka, and Haggai Maron. Expressive
628 sign equivariant networks for spectral geometric learning. In *NeurIPS*,
629 2023a. URL [http://papers.nips.cc/paper_files/paper/2023/hash/](http://papers.nips.cc/paper_files/paper/2023/hash/3516aa3393f0279e04c099f724664f99-Abstract-Conference.html)
630 [3516aa3393f0279e04c099f724664f99-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2023/hash/3516aa3393f0279e04c099f724664f99-Abstract-Conference.html).
- 631 Derek Lim, Joshua David Robinson, Lingxiao Zhao, Tess E. Smidt, Suvrit Sra, Haggai Maron, and
632 Stefanie Jegelka. Sign and basis invariant networks for spectral graph representation learning.
633 In *ICLR 2023*. OpenReview.net, 2023b. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=Q-UHqMorzil)
634 [Q-UHqMorzil](https://openreview.net/forum?id=Q-UHqMorzil).
- 635
- 636 Xiao Liu, Shiyu Zhao, Kai Su, Yukuo Cen, Jiezhong Qiu, Mengdi Zhang, Wei Wu, Yuxiao Dong,
637 and Jie Tang. Mask and reason: Pre-training knowledge graph transformers for complex logical
638 queries. In *SIGKDD*, 2022.
- 639
- 640 Andreas Loukas. What graph neural networks cannot learn: depth vs width. In *ICLR*, 2020.
- 641
- 642 László Lovász. *Large Networks and Graph Limits*, volume 60 of *Colloquium Publications*. American
643 Mathematical Society, 2012.
- 644
- 645 Liheng Ma, Chen Lin, Derek Lim, Adriana Romero-Soriano, Puneet K. Dokania, Mark Coates, Philip
646 H. S. Torr, and Ser-Nam Lim. Graph inductive biases in transformers without message passing. In
647 *ICML*, 2023.
- 648
- 649 Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably powerful graph
650 networks. In *NeurIPS*, 2019a.

- 648 Haggai Maron, Ethan Fetaya, Nimrod Segol, and Yaron Lipman. On the universality of invariant
649 networks. In *ICML*, 2019b.
- 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 Daniel Neuen. Homomorphism-distinguishing closedness for graphs of bounded tree-width. *arXiv*
655 *preprint arXiv:2304.07011*, 2023.
- 656
657 Hoang Nguyen and Takanori Maehara. Graph homomorphism convolution. In *ICML*, 2020.
- 658 Ladislav Rampásek, Michael Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Do-
659 minique Beaini. Recipe for a general, powerful, scalable graph transformer. In *NeurIPS*, 2022.
- 660
661 Eran Rosenbluth, Jan Tönshoff, Martin Ritzert, Berke Kisin, and Martin Grohe. Distinguished in
662 uniform: Self-attention vs. virtual nodes. In *ICLR*, 2024.
- 663 Ryoma Sato, Makoto Yamada, and Hisashi Kashima. Random features strengthen graph neural
664 networks. In *SDM*, 2021.
- 665
666 Edward R Scheinerman and Daniel H Ullman. *Fractional graph theory: a rational approach to the*
667 *theory of graphs*. Courier Corporation, 2013.
- 668 Ahsan Shehzad, Feng Xia, Shagufta Abid, Ciyuan Peng, Shuo Yu, Dongyu Zhang, and Karin Verspoor.
669 Graph transformers: A survey. *arXiv preprint arXiv:2407.09777*, 2024.
- 670
671 Hamed Shirzad, Ameya Velingker, Balaji Venkatachalam, Danica J Sutherland, and Ali Kemal Sinop.
672 Expformer: Sparse transformers for graphs. In *ICML*, 2023.
- 673
674 Jan Tönshoff, Martin Ritzert, Eran Rosenbluth, and Martin Grohe. Where did the gap go? reassessing
675 the long-range graph benchmark. *Trans. Mach. Learn. Res.*, 2024, 2024. URL <https://openreview.net/forum?id=Nm0WX86sKv>.
- 676
677 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
678 Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- 679
680 Haorui Wang, Haoteng Yin, Muhan Zhang, and Pan Li. Equivariant and stable positional encoding
681 for more powerful graph neural networks. In *ICLR*, 2022.
- 682
683 Yanbo Wang and Muhan Zhang. Towards better evaluation of gnn expressiveness with brec dataset.
684 In *ICML*, 2024.
- 685
686 Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S
687 Pappu, Karl Leswing, and Vijay Pande. Moleculenet: a benchmark for molecular machine learning.
688 *Chemical science*, 9(2):513–530, 2018.
- 689
690 Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural
691 networks? In *ICLR*, 2019.
- 692
693 Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and
694 Tie-Yan Liu. Do transformers really perform badly for graph representation? In *NeurIPS*, 2021.
- 695
696 Seongjun Yun, Minbyul Jeong, Raehyun Kim, Jaewoo Kang, and Hyunwoo J Kim. Graph transformer
697 networks. In *NeurIPS*, 2019.
- 698
699 Bingxu Zhang, Changjun Fan, Shixuan Liu, Kuihua Huang, Xiang Zhao, Jincui Huang, and Zhong
700 Liu. The expressive power of graph neural networks: A survey, 2023a. URL <https://arxiv.org/abs/2308.08235>.
- 701
702 Bohang Zhang, Shengjie Luo, Liwei Wang, and Di He. Rethinking the expressive power of Gnns via
703 graph biconnectivity. In *ICLR*, 2023b.
- 704
705 Bohang Zhang, Jingchu Gai, Yiheng Du, Qiwei Ye, Di He, and Liwei Wang. Beyond weisfeiler-
706 lehman: A quantitative framework for gnn expressiveness. In *ICLR*, 2024.


702 Wen Zhang, Yushan Zhu, Mingyang Chen, Yuxia Geng, Yufeng Huang, Yajing Xu, Wenting Song,
703 and Huajun Chen. Structure pretraining and prompt tuning for knowledge graph transfer. In *WWW*,
704 2023c.

707 A PROOFS FOR SECTION 4.2

709 *Proof for Proposition 1.* Dvorák (2010) showed that if G and H are equivalent under k -WL, then
710 every graph F with treewidth at most k has $\text{hom}(F, G) = \text{hom}(F, H)$. This extends also to the
711 vertex level. Let $v \in V(G)$ and $v' \in V(H)$ such that v and v' have equivalent k -WL labels. Then
712 also $\text{Hom}(F, G, v) = \text{Hom}(F, H, v')$ (see Lanzinger & Barceló (2024), Lemma 12). Hence, if the
713 maximum treewidth in \mathcal{G} is k , then equivalence under k -WL implies also that they are equivalent
714 under $\text{MoSE}_{\mathcal{G}}$. \square

716 *Proof Sketch of Proposition 6.* One direction is shown in ??, which shows two graphs G and H , with
717 highlighted red and blue nodes each. For every step length ℓ RWSE produces the same node features
718 for both red vertices and both blue vertices, respectively. At the same time, it is straightforward to see
719 that the 1-WL labeling of the red/blue nodes is different in the two graphs.

720 For the other direction, the classic example comparing the graph G consisting of two triangles and the
721 graph H containing the 6-vertex cycle as drawn below. It is well known that all nodes in both graphs
722 have the same 1-WL label, in contrast RWSE^3 labels the nodes in G differently from the nodes in H ,
723 the respective vectors are given below. \square



732 Figure 5: The 6 vertex cycle (left) and the disjoint sum of 2 triangles are indistinguishable by 1-WL,
733 but distinguishable by RWSE already after 3 steps. The RWSE vectors in the respective graphs are
734 all the same and are given next to the graphs.

737 A.1 THEOREM 3

738 *Proof of Proposition 5.* Take any graph H and notate $V(H) = \{1, 2, \dots, n\}$. For $v, v' \in V(H)$,
739 define $P(v \xrightarrow{i} v')$ to be the probability of starting a length i random walk at node v and ending it at
740 node v' . Here, “length i ” refers to the number of edges in the random walk where we allow repeat
741 edges, and “random walk” refers to a walk where each step (say we are currently at node v) assigns a
742 uniform probability $1/d(v)$ to moving to any of the neighbors in $\mathcal{N}(v)$.
743

744 As shorthand, define $M = D^{-1}A$ where D and A are the diagonal degree matrix and adjacency
745 matrix of H respectively. Let us first prove that $(M^i)_{v,v'} = P(v \xrightarrow{i} v')$ for any nodes $v, v' \in V$ and
746 for all powers $i \in \mathbb{N}$ by performing induction on i . The base case $i = 1$ holds trivially because we
747 assume that random steps are taken uniformly over neighbors. Assume as the inductive hypothesis
748 that our claim holds for some $i \geq 1$, and note that:

749

750

$$(M^{i+1})_{v,v'} = (M^1 M^i)_{v,v'} = \sum_{x=1}^n M_{v,x} \cdot M_{x,v'}^i$$

751

752

753 but our base case and inductive hypothesis tell us that

754

755

$$\sum_{x=1}^n M_{v,x} \cdot M_{x,v'}^i = \sum_{x=1}^n P(v \xrightarrow{1} x) \cdot P(x \xrightarrow{i} v')$$

By the law of total probability:

$$\sum_{x=1}^n P(v \xrightarrow{1} x) \cdot P(x \xrightarrow{i} v') = P(v \xrightarrow{i+1} v')$$

This completes the induction. Hence, for any node v in any graph H , the node-level RWSE_k vector is equivalent to:

$$\text{RWSE}_k(v, H) = \left[(M^i)_{v,v} \right]_{i=1}^k = \left[P(v \xrightarrow{i} v) \right]_{i=1}^k \in \mathbb{R}^k$$

Now, let us show that we can recover this random-walk vector using MoSE. Consider the family of all cycles with up to k nodes $\mathcal{G} = \{C_i\}_{i=1}^k$, and let ω be the vertex weighting which maps any node to the reciprocal of its degree. For each $C_i \in \mathcal{G}$, enumerate the nodes $V(C_i) = \{u_0, u_1, \dots, u_{i-1}\}$ such that $(u_\ell, u_{\ell+1}) \in E(C_i)$ for all $\ell = 0, \dots, i-2$ and $(u_{i-1}, u_0) \in E(C_i)$.

Take any node v in any graph H , and define \mathcal{H}_i to be the set of all homomorphisms from C_i to G which map node $u_0 \in C_i$ to node $v \in G$. Each homomorphism in $f \in \mathcal{H}_i$ can be constructed by making $i-1$ choices of node image $f(u_\ell) \in \mathcal{N}(f(u_{\ell-1}))$ for $\ell = 1, \dots, i-1$ such that $f(u_{i-1}) \in \mathcal{N}(f(u_0))$ where we have fixed $f(u_0) = v$. In other words, each homomorphism $f \in \mathcal{H}_i$ corresponds to a i -edge walk (allowing edge repeats) starting and ending at node v , which we will call W_f . This correspondence is described by taking the induced subgraph $W_f = G[f(C_i)]$ where $f(C_i)$ is the image of C_i under homomorphism f . In the reverse direction, any i -edge walk

$$W_f : v = v^{(0)}, v^{(1)}, \dots, v^{(i-1)}, v^{(i)} = v \quad \text{in } H$$

that starts and ends at node v corresponds to a homomorphism $f \in \mathcal{H}_i$ which maps $f(u_\ell) = v^{(\ell)}$. Hence, we have established a bijective correspondence between \mathcal{H}_i and i -edge walk starting/ending at v .

For any $f \in \mathcal{H}_i$, we have:

$$\prod_{v \in V(C_i)} \omega(f(v)) = \prod_{v \in W_f} 1/d(v) = \prod_{\ell=0}^{i-1} 1/d(v^{(\ell)})$$

Since the probability of stepping from node $v^{(\ell-1)} \in V(W_f)$ to node $v^{(\ell)} \in V(W_f)$ in a random walk is simply $1/d(v^{(\ell-1)})$, it holds that:

$$\prod_{\ell=0}^{i-1} 1/d(v^{(\ell)}) = P(W_f)$$

where $P(W_f)$ is the probability of a random walk starting in $v^{(0)}$ and ending in $v^{(i)}$ being exactly the walk W_f . Thus,

$$\omega\text{-Hom}_{u_0 \rightarrow v}(C_i, G) = \sum_{f \in \mathcal{H}_i} P(W_f) = P(v \xrightarrow{i} v)$$

giving us:

$$\text{MoSE}_{\mathcal{G}, \omega}(v, H) = \left[\omega\text{-Hom}_{\rightarrow v}(C_i, H) \right]_{i=1}^d = \left[P(v \xrightarrow{i} v) \right]_{i=1}^k = \text{RWSE}_k(v, H)$$

□

Proof of Proposition 4. We first recall the walk refinement procedure from Lichter et al. (2019) such that we can relate it directly to RWSE.

The scheme at its basis assumes a directed complete colored graph $G = (V, H, \chi)$ where the coloring χ always assigns different colors to self-loops than to other edges. For tuples of m vertices $(v_1, v_2, \dots, v_m) \in V^m$ define

$$\bar{\chi} = (v_1, v_2, \dots, v_m) = ((\chi(v_1, v_2), \chi(v_2, v_3), \dots, \chi(v_{m-1}, v_m)).$$

Ultimately, we want to use the scheme on undirected uncolored graphs. An undirected (and uncolored) graph G , is converted into the setting above by adding the coloring $\chi : V(G) \rightarrow \{-1, 0, 1\}$ as follows: for all self-loops χ assigns -1 , i.e., $\chi(v, v) = -1$. For all distinct pairs v, u , we set $\chi(v, u) = 1$ if $(v, u) \in E(G)$ and $\chi(v, u) = 0$ if $(v, u) \notin E(G)$. For the directed edges recall that we consider complete directed graphs and hence $E = V^2$.

With this in hand we can define the walk refinement procedure of Lichter et al. (2019). Formally, for $k \geq 2$, the k -walk refinement of a colored complete directed graph $G = (V, E, \chi)$ as the function that gives the new coloring $\chi_{W[k]}$ to every edge as follows

$$\chi_{W[k]}(v, u) = \{\{\bar{\chi}(v, w_1, \dots, w_{k-1}, u) \mid w_i \in V\}\}$$

Intuitively, the refinement takes into account all walks of length k in the graph. Note that walks of length shorter than k are also captured via self-loops (because of their different weights we are also always aware that a self-loop is taken and that a specific tuple in the multi-set represents a shorter walk).

As with the Weisfeiler-Leman color refinement procedure, k -walk refinement can be iterated until it reaches a stable refinement (after finitely many steps). Importantly, Lichter et al. (2019) (Lemma 4) show that the stable refinements of the k -walk refinement procedure and 2-WL produce the same partitioning of vertices. Intuitively, pairs of vertices that look the same to $W[k]$ iff they look the same to 2-WL.

All that is then left is to show that the stable $W[k]$ refinement uniquely determines the RWSE_k feature vector for every vertex v . For this it is enough to observe, that at the first step of $W[k]$, which we will denote as $\chi_{W[k]}^1$, both the degree of V , as well as all the number of k' -walks that start and end in v are determined by $\chi_{W[k]}^1(v, v)$ for every $v \in V$ and $k' \leq k \geq 2$. Furthermore, the number of k' -walks from v to v is directly obtainable from every level of the refinement, as every level of refinement of (v, u) preserves information about the color of the initial edge coloring $\chi(v, u)$ (recall, the initial edge coloring provides information about which edges exist in the original undirected G , and hence from every walk tuple we see whether the walk exists in G). Hence, in $\chi_{W[k]}^2$ from every tuple in the multiset for (v, v) , we can see whether the multiset is a k' -walk for every $k' \leq k$ (in the original undirected graph), with the labels in the tuple determining the degree of each vertex in the walk. From the definition of RWSE_k it is clear that this is enough information (see also the proof of Proposition 5) to determine all dimensions of the RWSE_k feature vectors from this information. \square

B ADDITIONAL EXPERIMENTAL DETAILS

We provide additional experimental details and hyperparameters relating to the results that were presented in Section 5. All code and instructions on how to reproduce our results are available at the following link: <https://github.com/a82641806/MoSE>.

Compute Resources. All experiments for ZINC, QM9, and the synthetic dataset were conducted on a cluster with 12 NVIDIA A10 GPUs (24 GB) and 4 NVIDIA H100s. Each node had 64 cores of Intel(R) Xeon(R) Gold 6326 CPU at 2.90GHz and 500GB of RAM. All experiments used at most 1 GPU at a time.

B.1 ZINC EXPERIMENTAL DETAILS

We use a subset of the ZINC molecular dataset that contains 12,000 graphs (Dwivedi et al., 2023). The dataset is split into 10,000 graphs for training, 1,000 graphs for validation, and 1,000 graphs for testing. Each graph in the dataset represents a molecule, where the node features indicate the atom type, and the edge features indicate the bond type between two nodes.

B.1.1 ZINC WITH EDGE FEATURES

In Table 5, we provide additional results on ZINC-12k with edge features. Observe that the trend for MoSE performing better than RWSE remains consistent. For completeness sake, we also perform a series of experiments where we combine RWSE+MoSE encodings. The combination does not have any significant effect on the results. All results are reported as the average of 4 runs at different seeds.

Table 5: We report mean absolute error (MAE) for graph regression on ZINC-12K (with edge features) across several models. MoSE yields substantial improvements over RWSE for every architecture.

Model	Encoding Type			
	<i>none</i>	RWSE	MoSE	RWSE+MoSE
MLP-E	0.606±0.002	0.361±0.010	0.347±0.003	0.349±0.003
R-GCN	0.413±0.005	0.207±0.007	0.197±0.004	0.188±0.005
GIN-E	0.243±0.006	0.122±0.003	0.118±0.007	0.117±0.005
GIN-E+VN	0.151±0.006	0.085±0.003	0.068±0.004	0.064±0.003
GT-E	0.195±0.025	0.104±0.025	0.089±0.018	0.090±0.005
GPS	0.119±0.011	0.070±0.002	0.062±0.002	0.065±0.002

B.1.2 ZINC WITHOUT EDGE FEATURES

Dwivedi et al. (2023) also present a set of experiments that uses the ZINC-12k dataset without edge features. This is because there are a number of models that do not take edge features into account. Therefore, we also perform a set of experiments on ZINC that do not include edge features and report the results in Table 6.

Table 6: We report additional MAE results for various models on ZINC-12k without edge features.

Model	<i>none</i>	RWSE	MoSE	+RWSE+MoSE
MLP	0.663±0.002	0.263±0.006	0.218±0.005	0.202±0.001
GIN	0.294±0.012	0.190±0.004	0.158±0.004	0.168±0.005
GT	0.674±0.001	0.217±0.007	0.209±0.020	0.184±0.001
GPS	0.178±0.016	0.116±0.001	0.102±0.001	0.105±0.006

B.1.3 HYPERPARAMETERS

We detail a rough summary of the hyperparameters that we use for the ZINC experiments in Table 7. These correspond to the results that are shown in Table 5. Note that all models are restricted to a parameter budget of roughly 500k.

For all models, we perform the same grid-search on each feature enhancement (none, RWSE, MoSE) independently. For MLPe, we grid search depth \times width:

$$\{2, 4, 8\} \times \{156, 220, 284, 348\}$$

and each of these models is searched with StepLR scheduler and Cosine Annealing scheduler. For GT and GPS, we search depth \times width \times number of heads:

$$\{(8, 74, 4), (10, 64, 4), (10, 92, 4), (4, 104, 8)\}$$

We use configurations from Jin et al. (2024) for GINE.

B.2 QM9 EXPERIMENTAL DETAILS

The QM9 dataset is another molecular dataset that consists of 130,831 graphs Wu et al. (2018). There are split into 110,831 graphs for training, 10,000 for validation, and 10,000 for testing. Node features indicate the atom type and other additional atom features, such as its atomicnumber and the number of Hydrogens connected to it, etc. The edge features indicate bond type between two nodes. The task in this dataset is to predict 13 different quantum chemical properties, ranging from a molecule’s dipole moment (μ) to its free energy (G). Note that for this dataset, we use the edge features provided for all of our experiments.

Table 7: Hyperparameters for ZINC experiments reported in Table 5.

Model	Layers	Hidden Dim	Batch Size	Learning Rate	Epochs	PE dim	#Heads
MLP-E	4	128	128	0.001	1000	64	N/A
GIN-E	4	128	128	0.001	1000	64	N/A
GIN-E+VN	4	128	128	0.001	1000	64	N/A
GT-E	4	92	64	0.005	2000	64	4
GPS	4	92	64	0.005	2000	64	4

Table 8: We report MAE results for GPS with various feature enhancements on the QM9 dataset.

Property	GPS	+RWSE	+MoSE	+RWSE+MoSE
mu	1.52±0.02	2.64±0.11	1.47±0.02	1.43±0.02
alpha	2.62±0.38	1.52±0.27	1.48±0.16	1.72±0.11
HOMO	1.17±0.41	0.91±0.01	0.91±0.01	0.92±0.01
LUMO	0.92±0.01	0.90±0.06	0.86±0.01	0.88±0.16
gap	1.46±0.02	1.47±0.02	1.45±0.02	1.48±0.01
R2	6.82±0.31	6.11±0.16	6.22±0.19	6.01±0.03
ZPVE	2.25±0.18	2.63±0.44	2.43±0.27	2.23±0.25
U0	0.96±0.34	0.83±0.17	0.85±0.08	0.80±0.05
U	0.81±0.05	0.83±0.15	0.75±0.03	0.78±0.03
H	0.81±0.26	0.86±0.15	0.83±0.09	0.87±0.16
G	0.77±0.04	0.83±0.12	0.80±0.14	0.72±0.03
Cv	2.56±0.72	1.25±0.05	1.02±0.04	1.03±0.08
Omega	0.40±0.01	0.39±0.02	0.38±0.01	0.39±0.01

B.2.1 ADDITIONAL RESULTS ON QM9

We provide additional results on the QM9 dataset for GPS in Table 8. All reported results are the average of 5 different runs at different seeds. Additionally, we provide the same analysis with a 4-layer MLP-E and the GT (Dwivedi & Bresson, 2020) model. These results are presented in Table 9 and Table 10 respectively. We consistently see that across all models, MoSE outperforms RWSE. The division between using MoSE alone and a combination of RWSE+MoSE is less trivial than in the case of ZINC though. This is because in the QM9 dataset, due to the large number of tasks to regress over, some tasks may not be influenced by topological structure as much. Therefore, the benefit of using MoSE is less significant.

Table 9: We report MAE results for MLP-E with different structural encodings on the QM9 dataset.

Property	MLP-E	+RWSE	+MoSE	+RWSE+MoSE
mu	5.31±0.03	4.91±0.06	4.88±0.05	4.84±0.03
alpha	6.76±0.04	5.08±0.18	5.11±0.11	4.83±0.09
HOMO	2.96±0.03	2.46±0.03	2.36±0.04	2.34±0.02
LUMO	3.17±0.04	2.76±0.3	2.65±0.03	2.57±0.03
gap	4.34±0.02	3.58±0.02	3.45±0.05	3.35±0.04
R2	43.69±1.08	26.69±0.31	26.50±0.43	24.97±0.45
ZPVE	14.36±0.60	10.42±1.24	8.94±0.83	10.28±0.35
U0	8.06±0.78	5.39±0.58	5.30±0.41	4.74±0.28
U	8.33±0.56	5.34±0.59	5.19±0.39	5.10±0.62
H	8.19±0.33	5.17±0.28	4.77±0.16	5.08±0.33
G	7.58±0.14	5.83±0.46	5.28±0.54	5.03±0.36
Cv	6.06±0.14	3.80±0.15	3.73±0.25	3.72±0.44
Omega	1.61±0.04	1.29±0.02	1.17±0.02	1.24±0.05

Table 10: We report MAE results for GT-E with different structural encodings on the QM9 dataset.

Property	GT	+RWSE	+MoSE	+RWSE+MoSE
mu	4.12±0.05	2.53±0.06	2.43±0.01	2.31±0.06
alpha	8.80±4.65	1.43±0.03	2.00±0.08	1.32±0.01
HOMO	1.72±0.01	1.15±0.01	1.11±0.01	1.08±0.01
LUMO	1.63±0.01	1.05±0.02	1.02±0.01	0.99±0.01
gap	2.38±0.03	1.67±0.01	1.62±0.01	1.57±0.02
R2	26.73±8.59	8.30±0.32	10.69±1.56	7.80±0.17
ZPVE	10.39±1.35	2.46±0.04	4.66±2.09	2.37±0.03
U0	6.27±1.95	0.94±0.10	2.56±1.66	0.98±0.57
U	5.38±0.17	0.92±0.03	3.10±2.75	1.40±0.72
H	5.30±0.17	0.92±0.06	3.49±2.02	0.96±0.51
G	5.44±0.51	0.88±0.08	2.72±1.21	0.63±0.03
Cv	4.32±0.07	1.25±0.01	1.43±0.04	1.15±0.01
Omega	1.25±0.01	0.50±0.01	0.47±0.02	0.46±0.01

Table 11: Hyperparameters for QM9 experiments.

Model	Layers	Hidden Dim	Batch Size	Learning Rate	Epochs	PE Dim	#Heads
MLP-E	4	220	128	0.001	800	64	N/A
GT	8, 10, 12	64, 128, 256	128	0.0001, 0.0001	1200	64	4, 8
GPS	8, 10, 12	64, 128, 256	128	0.0001, 0.0001	1200	64	4, 8

B.2.2 HYPERPARAMETERS

We detail a rough summary of the hyperparameters that we use for the QM9 experiments in Table 11. Different tasks use different hyperparameters. The precise setup can be found in the public github repository.

For MLP-E, we do not perform any grid searching, and simply use the same hyperparameters for every task and every feature enhancement. For GPS and GT we grid-search each feature enhancement independently (on a reduced number of epochs: 700). We first search the tasks R2, ZPVE, U0, and GAP on the following MPNN-type \times depth \times width \times head \times learning rate values:

$$\{\text{RGCN, GINE}\} \times \{(8, 128, 8), (10, 64, 4), (12, 256, 8)\} \times \{0.0001, 0.00001\}$$

Then, the best two models from each results is searched on the corresponding tasks for each injection:

Initial Search	Corresponding tasks
R2	MU, ALPHA
ZPVE	OMEGA, Cv
U0	U, H, G
GAP	HOMO, LUMO

B.3 SYNTHETIC DATASET DETAILS

We generate a new synthetic dataset where the goal is to predict the fractional domination number of a graph. The dataset contains 10,000 randomly-generated Erdős-Rényi graphs. The density range for the graph generation is between [0.25,0.75], and the number of nodes varies from [16,32].

A *fractional dominating set* of a graph G is a weight assignment $\alpha : V(G) \rightarrow [0, 1]$, s.t., for every vertex v , the weight assigned to it and its neighbours is at least 1. The weight of a fractional dominating set α is $\sum_{v \in V(G)} \alpha(v)$. The fractional domination number of G is the least weight of a fractional dominating set of G (Scheinerman & Ullman, 2013).

For the MLP, we again use a 4-layer MLP, and in order to focus solely on the power of the graph inductive biases, we select a standard self-attention GraphTransformer (Dwivedi & Bresson, 2020).

1026 The percentage mean absolute error results for each model and configuration are reported in Table 12.
1027 For each model, we do not perform any hyperparameter tuning and use the defaults outlined above.
1028 Please refer to our code for the specific configuration files.
1029

1030 Table 12: Percentage MAE results for prediction the fractional dominance number on our synthetic
1031 dataset. MoSE outperforms LapPE and RWSE in both an MLP and GT model.
1032

	LapPE	RWSE	MoSE
1033 MLP	48.22% \pm .31	7.53% \pm .07	5.79% \pm .09
1034 GT	8.35% \pm .06	6.05% \pm .08	5.46% \pm .10

1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079