
Recipe for a General, Powerful, Scalable Graph Transformer

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 We propose a recipe on how to build a general, powerful, scalable (GPS) graph
2 Transformer with linear complexity and state-of-the-art results on a diverse set
3 of benchmarks. Graph Transformers (GTs) have gained popularity in the field of
4 graph representation learning with a variety of recent publications but they lack a
5 common foundation about what constitutes a good positional or structural encoding,
6 and what differentiates them. In this paper, we summarize the different types of
7 encodings with a clearer definition and categorize them as being *local*, *global*
8 or *relative*. Further, GTs remain constrained to small graphs with few hundred
9 nodes, and we propose the first architecture with a complexity linear to the number
10 of nodes and edges $O(N + E)$ by decoupling the local real-edge aggregation
11 from the fully-connected Transformer. We argue that this decoupling does not
12 negatively affect the expressivity, with our architecture being a universal function
13 approximator for graphs. Our GPS recipe consists of choosing 3 main ingredients:
14 (i) positional/structural encoding, (ii) local message-passing mechanism, and (iii)
15 global attention mechanism. We build and open-source a modular framework that
16 supports multiple types of encodings and that provides efficiency and scalability
17 both in small and large graphs. We test our architecture on 11 benchmarks and
18 show very competitive results on all of them, show-casing the empirical benefits
19 gained by the modularity and the combination of different strategies.

20 1 Introduction

21 Graph Transformers (GTs) alleviate fundamental limitations pertaining to the sparse message passing
22 mechanism, e.g., over-smoothing [44], over-squashing [1], and expressiveness bounds [57, 42],
23 by allowing nodes to attend to all other nodes in a graph (*global attention*). This benefits several
24 real-world applications, such as modeling chemical interactions beyond the covalent bonds [59], or
25 graph-based robotic control [35]. Global attention, however, requires nodes to be better identifiable
26 within the graph and its substructures [13]. This has led to a flurry of recently proposed fully-
27 connected graph transformer models [13, 34, 59, 41, 29] as well as various positional encoding
28 schemes leveraging spectral features [13, 34, 36] and graph features [15, 8]. Furthermore, standard
29 global attention incurs quadratic computational costs $O(N^2)$ for a graph with N nodes and E edges,
30 that limits GTs to small graphs with up to a few hundred nodes.

31 Whereas various GT models focus on particular node identifiability aspects, a principled framework
32 for designing GTs is still missing. In this work, we address this gap and propose a recipe for building
33 general, powerful, and scalable (GPS) graph Transformers. The recipe defines (i) embedding modules
34 responsible for aggregating *positional encodings* (PE) and *structural encodings* (SE) with the node,
35 edge, and graph level input features; (ii) processing modules that employ a combination of local
36 message passing and global attention layers (see Figure 1).

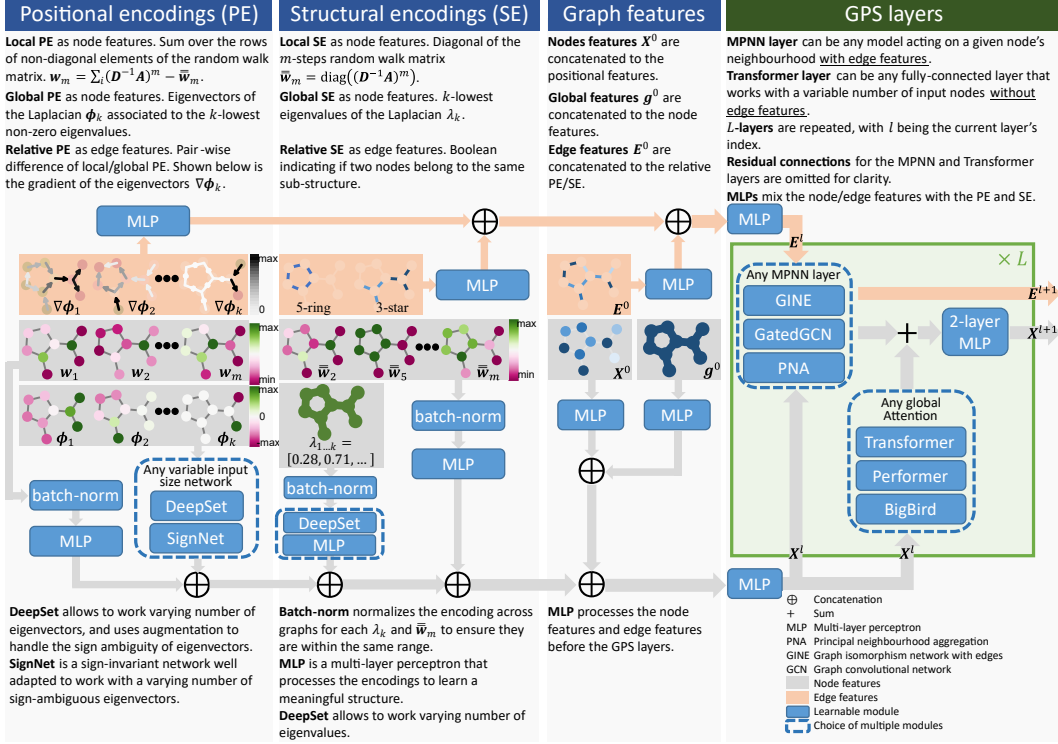


Figure 1: Modular GPS graph Transformer, with examples of PE and SE. Task specific layers for node/graph/edge-level predictions, such as pooling or output MLP, are omitted for simplicity.

37 The embedding modules organize multiple proposed PE and SE schemes into *local* and *global* levels
 38 serving as additional node features whereas positional and structural *relative* features contribute to
 39 edge features. The processing modules define a computational graph that allows to balance between
 40 message-passing graph neural networks (MPNNs) and Transformer-like global attention, including
 41 attention mechanisms *linear* in the number of nodes $O(N)$.

42 To the best of our knowledge, application of efficient attention models has not yet been thoroughly
 43 studied in the graph domain, e.g., only one work [10] explores the adaptation of Performer-style [11]
 44 attention approximation on small graphs. Particular challenges emerge with explicit edge features that
 45 are incorporated as attention bias in fully-connected graph transformers [34, 59]. Linear transformers
 46 do not materialize the attention matrix directly, hence incorporating edge features becomes a non-
 47 trivial task. In this work, we hypothesize that explicit edge features are not necessary for the *global*
 48 *graph attention* and adopt Performer [11] and BigBird [62] as exemplary linear attention mechanisms.

49 Our contributions are as follows. (i) Provide a general, powerful, scalable (GPS) GT blueprint that
 50 incorporates positional and structural encodings with local message passing and global attention,
 51 visualized in Figure 1. (ii) Provide a better definition of PEs and SEs and organize them into *local*,
 52 *global*, and *relative* categories. (iii) Show that GPS with linear global attention, e.g., provided by
 53 Performer [11] or BigBird [62], scales to graphs with several thousand nodes and demonstrates
 54 competitive results even without explicit edge features within the attention module, whereas existing
 55 fully-connected GTs [34, 59] are limited to graphs of up to few hundred nodes. (iv) Conduct
 56 an extensive ablation study that evaluates contribution of PEs, local MPNN, and global attention
 57 components in perspective of several benchmarking datasets. (v) Finally, following the success of
 58 GraphGym [61] we implement the blueprint within a modular and performant GRAPHGPS package.

59 2 Related Work

60 **Graph Transformers (GT).** Considering the great successes of Transformers in natural language
 61 processing (NLP) [52, 30] and recently also in computer vision [16, 23, 22], it is natural to study their
 62 applicability in the graph domain as well. Particularly, they are expected to help alleviate the problems

of over-smoothing and over-squashing [1, 51] in MPNNs, which are analogous to the vanishing gradients and lack of long-term dependencies in NLP. Fully-connected Graph Transformer [13] was first introduced together with rudimentary utilisation of eigenvectors of the graph Laplacian as the node positional encoding (PE), to provide the otherwise graph-unaware Transformer a sense of nodes’ location in the input graph. Building on top of this work, SAN [34] implemented an invariant aggregation of Laplacian’s eigenvectors for the PE, alongside conditional attention for real and virtual edges of a graph, which jointly yielded significant improvements. Concurrently, Graphormer [59, 48] proposed using pair-wise graph distances (or 3D distances) to define relative positional encodings, with outstanding success on large molecular benchmarks. Further, GraphiT [41] used relative PE derived from diffusion kernels to modulate the attention between nodes. Finally, GraphTrans [29] proposed the first hybrid architecture, first using a stack of MPNN layers, before fully-connecting the graph. Since, the field has continued to propose alternative GTs: SAT [8], EGT [27], GRPE [45].

Positional and structural encodings. There have been many recent works on PE and SE, notably on Laplacian PE [13, 34, 3, 36], shortest-path and node degree centrality [59], kernel distance [41], random-walk SE [15], structure-aware [8, 6, 5], and more. Some works also propose dedicated networks to learn the PE/SE from an initial encoding [34, 15, 36, 8]. To better understand the different PE/SE and the contribution of each work, we categorize them in Table 1 and examine their effect in Section 3.2. In most cases, PE/SE are used as soft bias, meaning they are simply provided as input features. But in other cases, they can be used to direct the messages [3] or create *bridges* between distant nodes [33, 51].

Linear Transformers. The *quadratic* complexity of attention in the original Transformer architecture [52] motivated the search for more efficient attention mechanisms that would scale *linearly* with the sequence length. Most of such *linear transformers* are developed for language modeling tasks, e.g., Linformer [54], Reformer [32], Longformer [4], Performer [11], BigBird [62], and have a dedicated Long Range Arena benchmark [49] to study the limits of models against extremely long input sequences. Pyraformer [37] is an example of a linear transformer for time series data, whereas S4 [21] is a more general signal processing approach that employs the state space model theory without the attention mechanism. In the graph domain, linear transformers are not well studied. Choromanski et al. [10] are the first to adapt Performer-style attention kernelization to small graphs.

3 Methods

In this work we provide a general, powerful, scalable (GPS) architecture for graph Transformers, following our 3-part recipe presented in Figure 1. We begin by categorization of existing positional (PE) and structural encodings (SE), a necessary ingredient for graph Transformers. Next, we analyse how these encodings also increase expressive power of MPNNs. The increased expressivity thus provides double benefit to our hybrid MPNN+Transformer architecture, which we introduce in Section 3.3. Last but not least, we provide an extensible implementation of GPS in GRAPHGPS package, built on top of PyG [18] and GraphGym [61].

3.1 Modular positional and structural encodings

One of our contribution is to provide a modular framework for PE/SE. It was shown in previous works that they are one of the most important factors in driving the performance of graph Transformers. Thus, a better understanding and organization of the PE and SE will aid in building of a more modular architecture and in guiding of the future research.

We propose to organize the PE and SE into 3 categories: **local**, **global** and **relative** in order to facilitate the integration within the pipeline and facilitate new research directions. They are presented visually in Figure 1, with more details in Table 1. Although PE and SE can appear similar to some extent, they are different yet complementary. PE gives a notion of distance, while SE gives a notion of structural similarity. One can always infer certain notions of distance from large structures, or certain notions of structure from short distances, but this is not a trivial task, and the objective of providing PE and SE remains distinct, as discussed in the following subsections.

Despite presenting a variety of possible functions, we focus our empirical evaluations on the **global PE**, **relative PE** and **local SE** since they are known to yield significant improvements. We leave the empirical evaluation of other encodings for future work.

Table 1: The proposed categorization of positional encodings (PE) and structural encodings (SE). Some encodings are assigned to multiple categories in order to show their multiple expected roles.

Encoding type	Description	Examples
Local PE node features	Allow a node to know its position and role within a local cluster of nodes. <i>Within a cluster, the closer two nodes are to each other, the closer their local PE will be, such as the position of a word in a sentence (not in the text).</i>	<ul style="list-style-type: none"> Sum each column of non-diagonal elements of the m-steps random walk matrix. Distance between a node and the centroid of a cluster containing the node.
Global PE node features	Allow a node to know its global position within the graph. <i>Within a graph, the closer two nodes are, the closer their global PE will be, such as the position of a word in a text.</i>	<ul style="list-style-type: none"> Eigenvectors of the Adjacency, Laplacian [14, 34] or distance matrices. Distance from the graph’s centroid. Unique identifier for each connected component of the graph.
Relative PE edge features	Allow two nodes to understand their distances or directional relationships. <i>Edge embedding that is correlated to the distance given by any global or local PE, such as the distance between two words.</i>	<ul style="list-style-type: none"> Pair-wise node distances from heat kernels, random-walks, Green’s function, graph geodesic [3, 34, 41], or any local/global PE. Gradient of eigenvectors [3, 34] or any local/global PE. Boolean indicating if two nodes are in the same cluster.
Local SE node features	Allow a node to understand what sub-structures it is a part of. <i>Given an SE of radius m, the more similar the m-hop subgraphs around two nodes are, the closer their local SE will be.</i>	<ul style="list-style-type: none"> Degree of a node [59]. Diagonal of the m-steps random-walk matrix [15]. Time-derivative of the heat-kernel diagonal (gives the degree at $t = 0$). Enumerate or count predefined structures such as triangles, rings, etc. [6, 64]. Ricci curvature [51].
Global SE graph features	Provide the network with information about the global structure of the graph. <i>The more similar two graphs are, the closer their global SE will be.</i>	<ul style="list-style-type: none"> Eigenvalues of the Adjacency or Laplacian matrices [34]. Graph properties: diameter, girth, number of connected components, number of nodes, number of edges, nodes-to-edges ratio.
Relative SE edge features	Allow two nodes to understand how much their structures differ. <i>Edge embedding that is correlated to the difference between any local SE.</i>	<ul style="list-style-type: none"> Gradient of any local SE. Boolean indicating if two nodes are in the same sub-structure [5] (similar to the gradient of sub-structure enumeration).

115 **Positional encodings (PE)** are meant to provide an idea of the *position in space* of a given node
116 within the graph. Hence, when two nodes are close to each other within a graph or subgraph, their PE
117 should also be close. A simple approach is to compute the pair-wise distance between each pairs of
118 nodes or their eigenvectors as proposed in [59, 34], but this is not compatible with linear Transformers
119 since it requires to materialize the full attention matrix [11]. Instead, we want the PE to either be
120 features of the nodes or real edges of the graph, thus a better fitting solution is to use the eigenvectors
121 of the graph Laplacian or their gradient [14, 3, 34]. More PE examples are available in Table 1.

122 **Structural encodings (SE)** are meant to provide an embedding of the structure of graphs or subgraphs
123 to help increase the expressivity and the generalizability of graph neural networks (GNN). Hence,
124 when two nodes share similar subgraphs, or when two graphs are similar, their SE should also be close.
125 Simple approaches are to identify pre-defined patterns in the graphs as one-hot encodings, but they
126 require expert knowledge of graphs [6, 5]. Instead, using the diagonal of the m -steps random-walk
127 matrix encodes richer information into each node [15], such as for odd m it can indicate if a node is a
128 part of an m -long cycle. Structural encodings can also be used to define the *global* graph structure,
129 for instance using the eigenvalues of the Laplacian, or as *relative* edge features to identify if nodes
130 are contained within the same clusters, with more examples in Table 1.

131 3.2 Why do we need PE and SE in MPNN?

132 As reviewed earlier, several recent GNNs make use of positional encodings (PE) and structural
133 encodings (SE) as soft biases to improve the model expressivity (summarized in Table 1), which
134 also leads to better generalization. In this section, we present an examination of PE and SE by
135 showing how message-passing networks, despite operating on the graph structure, remain blind to
136 the information encapsulated by the PE and SE.

137 **1-Weisfeiler-Leman test (1-WL).** It is well known that standard MPNNs are as expressive as the
 138 1-WL test, meaning that they fail to distinguish non-isomorphic graphs under a 1-hop aggregation.
 139 We argue that the selected *local*, *global* and *relative* PE/SE allow MPNNs to become more expressive
 140 than the 1-WL test, thus making them fundamentally more expressive at distinguishing between nodes
 141 and graphs. To this end, we study the following two types of graphs (Figure 2 and Appendix C.1).

142 **Circular Skip Link (CSL) graph.** In a CSL graph-pair [43], we
 143 want to be able to distinguish the two non-isomorphic graphs. Since
 144 the 1-WL algorithm produces the same color for every node in both
 145 graphs, also every MPNN will fail to distinguish them. However,
 146 using a *global* PE (e.g., Laplacian PE [14]) assigns each node a unique
 147 initial color and makes the CSL graph-pair distinguishable. This
 148 demonstrates that an MPNN cannot learn such a PE from the graph
 149 structure alone. Next, using a *local* SE (e.g., diagonals of m -steps
 150 random walk) can successfully capture the difference in the skip links
 151 of the two graphs [39], resulting in their different node coloring [15].

152 **Decalin molecule.** In the bicyclic Decalin graph, Figures 2b and C.1b,
 153 the node a is isomorphic to node b , and so is the node c to node d . A
 154 1-WL coloring of the nodes, and analogously MPNN, would generate
 155 one color for the nodes a, b and another color for c, d . The same
 156 applies to the aforementioned *local* SE [15]. In case of link prediction,
 157 this causes potential links (a, d) and (b, c) to be indistinguishable [63].
 158 Using a distance-based *relative* PE on the edges or an eigenvector-
 159 based *global* PE, however, would allow to differentiate the two links.




Figure 2: Example graphs with anonymous nodes without distinguishing features.

160 3.3 GPS layer: an MPNN+Transformer hybrid

161 In this section we introduce the GPS layer, which is a hybrid MPNN+Transformer layer. First we
 162 argue how it alleviates the limitation of a closely related work. Next, we list the layer update equations
 163 which can be instantiated with a variety of MPNN and Transformer layers. Finally, we present its
 164 characteristics in terms of modularity, scalability and expressivity.

165 **Preventing early smoothing.** Why not use an architecture like GraphTrans [29] comprising of a few
 166 layers of MPNNs before the Transformer? Since MPNNs are limited by problems of over-smoothing,
 167 over-squashing, and low expressivity against the WL test [1, 51], these layers could *irreparably* fail to
 168 keep some information in the early stage. Although they could make use of PE/SE or more expressive
 169 MPNNs [3, 15], they are still likely to lose information. An analogous 2-stage strategy was successful
 170 in computer vision [16, 22] thanks to the high expressivity of convolutional layers on grids, but we
 171 do not expect it to achieve the same success on graphs due to the limitations of message-passing.

172 **Update function.** At each layer, the features are updated by aggregating the output of an MPNN layer
 173 with that of a global attention layer, as shown in Figures 1 and D.1, and described by the equations
 174 below. Note that the edge features are only passed to the MPNN layer, and that residual connections
 175 with batch normalization [28] are omitted for clarity. Both the MPNN and GlobalAttn layers are
 176 modular, i.e., MPNN can be any function that acts on a local neighborhood and GlobalAttn can be
 177 any fully-connected layer.

$$\mathbf{X}^{\ell+1}, \mathbf{E}^{\ell+1} = \text{GPS}^{\ell}(\mathbf{X}^{\ell}, \mathbf{E}^{\ell}, \mathbf{A}) \quad (1)$$

$$\text{computed as } \mathbf{X}_M^{\ell+1}, \mathbf{E}^{\ell+1} = \text{MPNN}_e^{\ell}(\mathbf{X}^{\ell}, \mathbf{E}^{\ell}, \mathbf{A}), \quad (2)$$

$$\mathbf{X}_T^{\ell+1} = \text{GlobalAttn}^{\ell}(\mathbf{X}^{\ell}), \quad (3)$$

$$\mathbf{X}^{\ell+1} = \text{MLP}^{\ell}(\mathbf{X}_M^{\ell+1} + \mathbf{X}_T^{\ell+1}), \quad (4)$$

178 where $\mathbf{A} \in \mathbb{R}^{N \times N}$ is the adjacency matrix of a graph with N nodes and E edges; $\mathbf{X}^{\ell} \in \mathbb{R}^{N \times d_{\ell}}$, $\mathbf{E}^{\ell} \in$
 179 $\mathbb{R}^{E \times d_{\ell}}$ are the d_{ℓ} -dimensional node and edge features, respectively; MPNN_e^{ℓ} and GlobalAttn^{ℓ} are
 180 instances of an MPNN with edge features and of a global attention mechanism at the ℓ -th layer with
 181 their corresponding learnable parameters, respectively; MLP^{ℓ} is a 2-layer MLP block.

182 **Modularity** is achieved by allowing drop-in replacement for a number of module choices, including
 183 the initial PE/SE types, the networks that processes those PE/SE, the MPNN and global attention
 184 layers that constitute a GPS layer, and the final task-specific prediction head. Further, as research
 185 advances in different directions, GRAPHGPS allows to easily implement new PE/SE and other layers.

186 **Scalability** is achieved by allowing for a computational complexity linear in both the number of
 187 nodes and edges $O(N + E)$; excluding the potential precomputation step required for various PE,
 188 such as Laplacian eigen-decomposition. By restricting the PE/SE to real nodes and edges, and
 189 by excluding the edge features from the global attention layer, we can avoid materializing the full
 190 quadratic attention matrix. Therefore we can utilize a linear Transformer with $O(N)$ complexity,
 191 while the complexity of an MPNN is $O(E)$. For sparse graphs such as molecular graphs, regular
 192 graphs, and knowledge graphs, the edges are practically proportional to the nodes $E = \Theta(N)$,
 193 meaning the entire complexity can be considered linear in the number of nodes $O(N)$. Empirically,
 194 even on small molecular graphs, our architecture reduces computation time compared to other GT
 195 models, e.g., a model of ~6M parameters requires 196s per epoch on the ogbg-molpcba [25] dataset,
 196 compared to 883s for SAN [34] on the same GPU type.

197 **Expressivity** in terms of sub-structure identification and the Weisfeiler-Leman (WL) test is achieved
 198 via providing a rich set of PE/SE, as proposed in various works [3, 34, 15, 5, 6] and detailed in Section
 199 3.1. Further, the Transformer allows to resolve the expressivity bottlenecks caused by over-smoothing
 200 [34] and over-squashing [1] by allowing information to spread across the graph via full-connectivity.
 201 Finally, in Section 3.4, we demonstrate that, given the right components, the proposed architecture
 202 does not lose edge information and is a universal function approximator on graphs.

203 3.4 Theoretical expressivity

204 In this section, we first discuss how the MPNN layer allows to propagate edge and neighbor informa-
 205 tion on the nodes. Then, we show that the proposed model is a universal function approximator on
 206 graphs, similarly to the SAN architecture [34].

207 **Preserving edge information in the Transformer layer.** Most GTs do not fully utilize edge features
 208 of the input graph. The Graph Transformer [13], SAN [34] and Graphormer [59] only use edge
 209 features to condition the attention between a pair of nodes, that is, they influence the attention gating
 210 mechanism but are not explicitly involved in updating of the node representations. GraphiT [41]
 211 does not consider edge features at all. Recent 2-step methods GraphTrans [29] and SAT [8] can use
 212 edge features in their first MPNN step, however this step is applied only once and typically includes
 213 several k rounds of message passing. Therefore this latter approach may suffer from initial over-
 214 smoothing, as k -hop neighborhoods together with the respective edge features need to be represented
 215 in a fixed-sized node representation.

216 On the other hand, in GPS, interleaving one round of local neighborhood aggregation via an MPNN
 217 layer with global self-attention mechanism reduces the initial representation bottleneck and enables
 218 iterative local and global interactions. In the attention, the key-query-value mechanism only explicitly
 219 depends on the node features, but assuming efficient representation encoding by the MPNN, the node
 220 features can implicitly encode edge information, thus edges can play a role in either the key, query, or
 221 values. In Appendix C.2, we give a more formal argument on how, following an MPNN layer, node
 222 features can encode edge features alongside information related to node-connectivity.

223 **Universal function approximator on graphs.** Kreuzer et al. [34][Sec. 3.5] demonstrated the
 224 universality of graph Transformers. It was shown that, given the full set of Laplacian eigenvectors, the
 225 model was a universal function approximator on graphs and could provide an approximate solution to
 226 the isomorphism problem, making it more powerful than any Weisfeiler-Leman (WL) isomorphism
 227 test given enough parameters. Here, we argue that the same holds for our architecture since we can
 228 also use the full set of eigenvectors, and since all edge information can be propagated to the nodes.

229 4 Experiments

230 We perform ablation studies on 4 datasets to evaluate the contribution of the message-passing module,
 231 the global attention module, and the positional or structural encodings. Then, we evaluate GPS on a
 232 diverse set of 11 benchmarking datasets, and show state-of-the-art (SOTA) results in many cases.

233 We test on datasets from different sources to ensure diversity, providing their detailed description in
 234 Appendix A.1. From the Benchmarking GNNs [14], we test on the ZINC, PATTERN, CLUSTER,
 235 MNIST, CIFAR10. From the open graph benchmark (OGB) [25], we test on all graph-level datasets:
 236 ogbg-molhiv, ogbg-molpcba, ogbg-code2, and ogbg-ppa, and from their large-scale challenge we
 237 test on the OGB-LSC PCQM4Mv2 [26]. Finally, we also select MalNet-Tiny [19] with 5000 graphs,
 238 each of up to 5000 nodes, since the number of nodes provide a scaling challenge for Transformers.

Table 2: Summary of the ablation studies. Details of the architectural choices, parameters, standard deviation, and computation times are presented in Appendix B.

(a) Ablation of the Transformer and MPNN layers. We observe a major drop when using only a Transformer without an MPNN. Further, most datasets benefit from using a Transformer, without any negative impact.

(b) Ablation of the PE and SE types. RWSE provides consistent gains at relatively low computational cost, while SignNet^{DeepSets} is the single best performing encoding, albeit at increased computational cost.

	Ablation	ZINC	PCQM4Mv2 subset	CIFAR10	MalNet -Tiny
		MAE ↓	MAE ↓	Acc. ↑	Acc. ↑
Global Attention	<i>none</i>	0.070	0.1213	69.95	92.23
	Transformer	0.070	0.1159	72.31	93.50
	Performer	0.071	0.1142	70.67	92.64
	BigBird	0.071	0.1237	70.48	92.34
MPNN	<i>none</i>	0.113	0.3294	68.86	73.90
	GINE	0.070	0.1284	71.11	92.27
	GatedGCN	0.086	0.1159	72.31	92.64
	PNA	0.070	0.1409	73.42	91.67

	Ablation	ZINC	PCQM4Mv2 subset	CIFAR10	MalNet -Tiny
		MAE ↓	MAE ↓	Acc. ↑	Acc. ↑
PE / SE	<i>none</i>	0.113	0.1355	71.49	92.64
	RWSE	0.070	0.1159	71.96	92.77
	LapPE	0.116	0.1201	72.31	92.74
	SignNet ^{MLP}	0.090	0.1158	71.74	92.57
	SignNet ^{DeepSets}	0.079	0.1144	72.37	93.13
	EquivStable ^{LapPE}	0.161	0.1209	72.10	92.27

*Encodings are color-coded by their **positional** or **structural** type.

4.1 Ablation studies

In this section, we evaluate multiple options for the three main components of our architecture in order to gauge their contribution to predictive performance and to better guide dataset-specific hyper-parameter optimization. First, we quantify benefits of the considered global-attention modules in 4 tasks. Then, we note that the MPNN layer is an essential part for high-performing models, and identify the layer type most likely to help. Finally, we observe when different global PE or local SE provide significant boost in the performance. All ablation results are averaged over multiple random seeds and summarized in Table 2, with additional information available in Appendix B.

Global-Attention module. Here we consider global attention implemented as $O(N^2)$ key-query-value Transformer attention or linear-time attention mechanisms of Performer or BigBird. We notice in Table 2a that using a Transformer is always beneficial, except for the ZINC dataset where no changes are observed. This motivates our architecture and the hypothesis that long-range dependencies are generally important. We further observe that Performer falls behind Transformer in terms of the predictive performance, although it provides a gain over the baseline and the ability to scale to very large graphs. Finally, BigBird in our setting offers no significant gain, while also being slower than Performer (see Appendix B).

Having no gain on the ZINC dataset is expected since the task is a combination of the computed octanol-water partition coefficient (cLogP) [56] and the synthetic accessibility score (SA-score) [17], both of which only count occurrences of local sub-structures. Hence, there is no need for a global connectivity, but a strong need for structural encodings.

Message-passing module. Next, we evaluate the effect of various message-passing architectures, Table 2a. It is apparent that they are fundamental to the success of our method: removing the layer leads to a significant drop in performance across all datasets. Indeed, without an MPNN, the edge features are not taken into consideration at all. Additionally, without reinforcing of the local graph structure, the network can overfit to the PE/SE. This reiterates findings of Kreuzer et al. [34], where considerably larger weights were assigned to the local attention.

We also find that although a vanilla PNA [12] generally outperforms GINE [24] and GatedGCN [14], adding the PE and SE results in major performance boost especially for the GatedGCN. This is consistent with results of Dwivedi et al. [15] and shows the importance of these encodings for gating.

Perhaps the necessity of a local message-passing module is due to the limited amount of graph data, and scaling to colossal datasets [46] that we encounter in language and vision could change that. Indeed, the Graphormer architecture [59] was able to perform very well on the full PCQM4Mv2 dataset without a local module. However, even large Transformer-based language models [7] and vision models [23] can benefit from an added local aggregation and outperform pure Transformers.

Table 3: Test performance in five benchmarks from [14]. Shown is the mean \pm s.d. of 10 runs with different random seeds. Highlighted are the top **first**, **second**, and **third** results.

Model	ZINC	MNIST	CIFAR10	PATTERN	CLUSTER
	MAE \downarrow	Accuracy \uparrow	Accuracy \uparrow	Accuracy \uparrow	Accuracy \uparrow
GCN [31]	0.367 \pm 0.011	90.705 \pm 0.218	55.710 \pm 0.381	71.892 \pm 0.334	68.498 \pm 0.976
GIN [57]	0.526 \pm 0.051	96.485 \pm 0.252	55.255 \pm 1.527	85.387 \pm 0.136	64.716 \pm 1.553
GAT [53]	0.384 \pm 0.007	95.535 \pm 0.205	64.223 \pm 0.455	78.271 \pm 0.186	70.587 \pm 0.447
GatedGCN [14]	0.282 \pm 0.015	97.340 \pm 0.143	67.312 \pm 0.311	85.568 \pm 0.088	73.840 \pm 0.326
GatedGCN-LSPE [15]	0.090 \pm 0.001	—	—	—	—
PNA [12]	0.188 \pm 0.004	97.94 \pm 0.12	70.35 \pm 0.63	—	—
DGN [3]	0.168 \pm 0.003	—	72.838 \pm 0.417	86.680 \pm 0.034	—
GSN [6]	0.101 \pm 0.010	—	—	—	—
CIN [5]	0.079 \pm 0.006	—	—	—	—
CRaWI [50]	0.085 \pm 0.004	97.944 \pm 0.050	69.013 \pm 0.259	—	—
GIN-AK+ [64]	0.080 \pm 0.001	—	72.19 \pm 0.13	86.850 \pm 0.057	—
SAN [34]	0.139 \pm 0.006	—	—	86.581 \pm 0.037	76.691 \pm 0.65
Graphormer [59]	0.122 \pm 0.006	—	—	—	—
K-Subgraph SAT [8]	0.094 \pm 0.008	—	—	86.848 \pm 0.037	77.856 \pm 0.104
EGT-SPE(+DO) [27]	0.154 \pm 0.011	97.722 \pm 0.222	67.004 \pm 0.624	86.730 \pm 0.036	77.909 \pm 0.245
GPS (ours)	0.070 \pm 0.004	98.051 \pm 0.126	72.298 \pm 0.356	90.324 \pm 0.132	77.95 \pm 0.305

Positional/Structural Encodings. Finally, we evaluate the effects of various PE/SE schemes, Table 2b. We find them generally beneficial to downstream tasks, in concordance to the vast literature on the subject (see Table 1). The benefits of the different encodings are very dataset dependant, with the random-walk structural encoding (RWSE) being more beneficial for molecular data and the Laplacian eigenvectors encodings (LapPE) being more beneficial in image superpixels. However, using SignNet with DeepSets encoding [36] as an improved way of processing the LapPE seems to be consistently successful across tasks. We hypothesize that SignNet can learn structural representation using the eigenvectors, for example, to generate local heat-kernels that approximate random walks [2].

4.2 Benchmarking GPS

We compare GPS against a set of popular message-passing neural networks (GCN, GIN, GatedGCN, PNA, etc.), graph transformers (SAN, Graphormer, etc.), and other recent graph neural networks with SOTA results (CIN, CRaWL, GIN-AK+, ExpC). To ensure diverse benchmarking tasks, we use datasets from Benchmarking-GNNs [14], OGB [25] and its large-scale challenge [26], with more details given in Appendix A.1. We report the mean and standard deviation over 10 random seeds.

Benchmarking GNNs [14]. We first benchmark our method on 5 tasks from Benchmarking GNNs [14], namely ZINC, MNIST, CIFAR10, PATTERN, and CLUSTER, shown in Table 3. We observe that our GPS gives SOTA results on 4 datasets, showcasing the ability to perform very well on a variety of synthetic tasks designed to test the model expressivity.

Open Graph Benchmark [25]. Next, we benchmark on all 4 graph-level tasks from OGB, namely molhiv, molpcba, ppa, and code2, Table 4. On the molhiv dataset, we observed our model to suffer from overfitting, but to still outperform SAN, while other graph Transformers do not report results. For the molpcba, ppa, and code2, GPS always ranks among the top 3 models, highlighting again the versatility and expressiveness of the GPS approach. Further, GPS outperforms every other GT on all 4 benchmarks, except SAT on code2.

OGB-LSC PCQM4Mv2 [26]. The large-scale PCQM4Mv2 dataset has been a popular benchmark for recent GTs, particularly due to Graphormer [59] winning the initial challenge. We report the results in Table 5, observing significant improvements over message-passing networks at comparable parameter budget. GPS also outperforms GRPE [45], EGT [27], and Graphormer [59] with less than half their parameters, and with significantly less overfitting on the training set. Contrarily to Graphormer, we do not need to precompute spatial distances from approximate 3D molecular conformers [60], the RWSEs we utilize are graph-based only.

MalNet-Tiny. The MalNet-Tiny [19] dataset consists of function call graphs with up to 5,000 nodes. These graphs are considerably larger than previously considered inductive graph-learning benchmarks, which enables us to showcase scalability of GPS to much larger graphs than prior methods. Our GPS

Table 4: Test performance in graph-level OGB benchmarks [25]. Shown is the mean \pm s.d. of 10 runs. Models that were first pre-trained on another dataset or use an ensemble are not included here.

Model	ogbg-molhiv	ogbg-molpcba	ogbg-ppa	ogbg-code2
	AUROC \uparrow	Avg. Precision \uparrow	Accuracy \uparrow	F1 score \uparrow
GCN+virtual node	0.7599 \pm 0.0119	0.2424 \pm 0.0034	0.6857 \pm 0.0061	0.1595 \pm 0.0018
GIN+virtual node	0.7707 \pm 0.0149	0.2703 \pm 0.0023	0.7037 \pm 0.0107	0.1581 \pm 0.0026
GatedGCN-LSPE	—	0.267 \pm 0.002	—	—
PNA	0.7905 \pm 0.0132	0.2838 \pm 0.0035	—	0.1570 \pm 0.0032
DeeperGCN	0.7858 \pm 0.0117	0.2781 \pm 0.0038	0.7712 \pm 0.0071	—
DGN	0.7970 \pm 0.0097	0.2885 \pm 0.0030	—	—
GSN (directional)	0.8039 \pm 0.0090	—	—	—
GSN (GIN+VN base)	0.7799 \pm 0.0100	—	—	—
CIN	0.8094 \pm 0.0057	—	—	—
GIN-AK+	0.7961 \pm 0.0119	0.2930 \pm 0.0044	—	—
CRaWl	—	0.2986 \pm 0.0025	—	—
ExpC [58]	0.7799 \pm 0.0082	0.2342 \pm 0.0029	0.7976 \pm 0.0072	—
SAN	77.85 \pm 0.247	0.2765 \pm 0.0042	—	—
GraphTrans (GCN-Virtual)	—	0.2761 \pm 0.0029	—	0.1830 \pm 0.0024
K-Subtree SAT	—	—	0.7522 \pm 0.0056	0.1937 \pm 0.0028
GPS (ours)	0.7880 \pm 0.0101	0.2907 \pm 0.0028	0.8015 \pm 0.0033	0.1894 \pm 0.0024

Table 5: Evaluation on PCQM4Mv2 [26] dataset. For GPS evaluation, we treated the *validation* set of the dataset as a test set, since the *test-dev* set labels are private. For more details refer to Appendix A.

Model	PCQM4Mv2			
	Test-dev MAE \downarrow	Validation MAE \downarrow	Training MAE	# Param.
GCN	0.1398	0.1379	n/a	2.0M
GCN-virtual	0.1152	0.1153	n/a	4.9M
GIN	0.1218	0.1195	n/a	3.8M
GIN-virtual	0.1084	0.1083	n/a	6.7M
GRPE [45]	0.0898	0.0890	n/a	46.2M
EGT [27]	0.0872	0.0869	n/a	89.3M
Graphormer [48]	n/a	0.0864	0.0348	48.3M
GPS-small	n/a	0.0938	0.0653	6.2M
GPS-medium	n/a	0.0858	0.0726	19.4M

reaches 92.72% \pm 0.7pp test accuracy when using Performer global attention. Interestingly, using Transformer global attention leads to further improved GPS performance, 93.36% \pm 0.6pp (based on 10 runs), albeit at the cost of doubled run-time. In both cases, we used comparable architecture to Freitas et al. [19], with 5 layers and 64 dimensional hidden node representation, and outperform their best GIN model with 90% accuracy. See Appendix B for GPS ablation study on MalNet-Tiny.

5 Conclusion

Our work is setting the foundation for a unified architecture of graph neural networks, with modular and scalable graph Transformers and a broader understanding of the role of graphs with positional and structural encodings. In our ablation studies, we demonstrated the importance of each module: the Transformer, flexible message-passing, and rich positional and structural encodings all contributed to the success of GPS on a wide variety of benchmarks. Indeed, considering 5 Benchmarking-GNN tasks [14], 4 OGB tasks [25], and 2 others, we outperformed every graph Transformer on 10 out of 11 tasks while also achieving state-of-the-art on 7 of them. We further showed that the model can scale to very large graphs of several thousand nodes, far beyond any previous graph Transformer. By open-sourcing the GRAPHGPS package, we hope to accelerate the research in efficient and expressive graph Transformers, and move the field closer to a unified hybrid Transformer architecture for graphs.

Limitations. We find that graph transformers are sensitive to hyperparameters and there is no *one-size-fits-all* solution for all datasets. We also identify a lack of challenging graph datasets necessitating long-range dependencies where linear attention architectures could exhibit all scalability benefits.

Societal Impact. As a general graph representation learning method, we do not foresee immediate negative societal outcomes. However, its particular application, e.g., in drug discovery or computational biology, will have to be thoroughly examined for trustworthiness or malicious usage.

References

- [1] Uri Alon and Eran Yahav. On the bottleneck of graph neural networks and its practical implications. In *International Conference on Learning Representations*, 2021.
- [2] Sebastian Andres, Jean-Dominique Deuschel, and Martin Slowik. Heat kernel estimates for random walks with degenerate weights. *Electronic Journal of Probability*, 21:1–21, 2016.
- [3] Dominique Beaini, Saro Passaro, Vincent Létourneau, Will Hamilton, Gabriele Corso, and Pietro Liò. Directional graph networks. In *International Conference on Machine Learning*, pages 748–758. PMLR, 2021.
- [4] Iz Beltagy, Matthew E. Peters, and Arman Cohan. Longformer: The long-document transformer. *CoRR*, abs/2004.05150, 2020.
- [5] Cristian Bodnar, Fabrizio Frasca, Nina Otter, Yuguang Wang, Pietro Lio, Guido F Montufar, and Michael Bronstein. Weisfeiler and Lehman go cellular: CW networks. *Advances in Neural Information Processing Systems*, 34:2625–2640, 2021.
- [6] Giorgos Bouritsas, Fabrizio Frasca, Stefanos P Zafeiriou, and Michael Bronstein. Improving graph neural network expressivity via subgraph isomorphism counting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [7] Ivan Chelombiev, Daniel Justus, Douglas Orr, Anastasia Dietrich, Frithjof Gressmann, Alexandros Kolioussis, and Carlo Luschi. Groupbert: Enhanced transformer architecture with efficient grouped structures. *arXiv:2106.05822*, 2021.
- [8] Dexiong Chen, Leslie O’Bray, and Karsten Borgwardt. Structure-aware transformer for graph representation learning. *arXiv:2202.03036*, 2022.
- [9] Zhengdao Chen, Lei Chen, Soledad Villar, and Joan Bruna. On the equivalence between graph isomorphism testing and function approximation with gnns. *Advances in Neural Information Processing Systems*, 2019.
- [10] Krzysztof Choromanski, Han Lin, Haoxian Chen, Tianyi Zhang, Arijit Sehanobish, Valerii Likhoshesterov, Jack Parker-Holder, Tamas Sarlos, Adrian Weller, and Thomas Weingarten. From block-Toeplitz matrices to differential equations on graphs: towards a general theory for scalable masked transformers. *arXiv:2107.07999*, 2021.
- [11] Krzysztof Marcin Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamás Sarlós, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukasz Kaiser, David Benjamin Belanger, Lucy J. Colwell, and Adrian Weller. Rethinking attention with performers. In *9th International Conference on Learning Representations*, 2021.
- [12] Gabriele Corso, Luca Cavalleri, Dominique Beaini, Pietro Liò, and Petar Veličković. Principal neighbourhood aggregation for graph nets. *Advances in Neural Information Processing Systems*, 33:13260–13271, 2020.
- [13] Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs. *arXiv:2012.09699*, 2020.
- [14] Vijay Prakash Dwivedi, Chaitanya K Joshi, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Benchmarking graph neural networks. *arXiv:2003.00982*, 2020.
- [15] Vijay Prakash Dwivedi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Graph neural networks with learnable structural and positional representations. In *International Conference on Learning Representations*, 2022.
- [16] Stéphane d’Ascoli, Hugo Touvron, Matthew L Leavitt, Ari S Morcos, Giulio Biroli, and Levent Sagun. Convit: Improving vision transformers with soft convolutional inductive biases. In *International Conference on Machine Learning*, pages 2286–2296. PMLR, 2021.
- [17] Peter Ertl and Ansgar Schuffenhauer. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *Journal of cheminformatics*, 1(1):1–11, 2009.
- [18] Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- [19] Scott Freitas, Yuxiao Dong, Joshua Neil, and Duen Horng Chau. A large-scale database for graph representation learning. In *35th Conference on Neural Information Processing Systems: Datasets and Benchmarks Track*, 2021.

- [20] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*, pages 1263–1272. PMLR, 2017.
- [21] Albert Gu, Karan Goel, and Christopher Re. Efficiently modeling long sequences with structured state spaces. In *International Conference on Learning Representations*, 2022.
- [22] Jianyuan Guo, Kai Han, Han Wu, Chang Xu, Yehui Tang, Chunjing Xu, and Yunhe Wang. CMT: Convolutional neural networks meet vision transformers. *arXiv:2107.06263*, 2021.
- [23] Kai Han, Yunhe Wang, Hanting Chen, Xinghao Chen, Jianyuan Guo, Zhenhua Liu, Yehui Tang, An Xiao, Chunjing Xu, Yixing Xu, et al. A survey on vision transformer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [24] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. Strategies for pre-training graph neural networks. *arXiv:1905.12265*, 2019.
- [25] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open Graph Benchmark: Datasets for Machine Learning on Graphs. *34th Conference on Neural Information Processing Systems*, 2020.
- [26] Weihua Hu, Matthias Fey, Hongyu Ren, Maho Nakata, Yuxiao Dong, and Jure Leskovec. OGB-LSC: A large-scale challenge for machine learning on graphs. In *35th Conference on Neural Information Processing Systems: Datasets and Benchmarks Track*, 2021.
- [27] Md Shamim Hussain, Mohammed J Zaki, and Dharmashankar Subramanian. Edge-augmented graph transformers: Global self-attention is enough for graphs. *arXiv:2108.03348*, 2021.
- [28] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015.
- [29] Paras Jain, Zhanghao Wu, Matthew Wright, Azalia Mirhoseini, Joseph E Gonzalez, and Ion Stoica. Representing long-range context for graph neural networks with global attention. *Advances in Neural Information Processing Systems*, 34, 2021.
- [30] Katikapalli Subramanyam Kalyan, Ajit Rajasekharan, and Sivanesan Sangeetha. Ammus: A survey of transformer-based pretrained models in natural language processing. *arXiv:2108.05542*, 2021.
- [31] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv:1609.02907*, 2016.
- [32] Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. In *8th International Conference on Learning Representations*, 2020.
- [33] Ioannis Koutis and Huong Le. Spectral modification of graphs for improved spectral clustering. *Advances in Neural Information Processing Systems*, 32, 2019.
- [34] Devin Kreuzer, Dominique Beaini, William L. Hamilton, Vincent Létourneau, and Prudencio Tossou. Rethinking graph transformers with spectral attention. In *Advances in Neural Information Processing Systems*, 2021.
- [35] Vitaly Kurin, Maximilian Igl, Tim Rocktäschel, Wendelin Boehmer, and Shimon Whiteson. My body is a cage: the role of morphology in graph-based incompatible control. *arXiv:2010.01856*, 2020.
- [36] Derek Lim, Joshua Robinson, Lingxiao Zhao, Tess Smidt, Suvrit Sra, Haggai Maron, and Stefanie Jegelka. Sign and basis invariant networks for spectral graph representation learning. *arXiv:2202.13013*, 2022.
- [37] Shizhan Liu, Hang Yu, Cong Liao, Jianguo Li, Weiyao Lin, Alex X. Liu, and Schahram Dustdar. Pyraformer: Low-complexity pyramidal attention for long-range time series modeling and forecasting. In *International Conference on Learning Representations*, 2022.
- [38] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.
- [39] Andreas Loukas. What graph neural networks cannot learn: depth vs width. In *International Conference on Learning Representations*, 2020.
- [40] Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably powerful graph networks. *arXiv:1905.11136*, 2019.

- [41] Grégoire Mialon, Dexiong Chen, Margot Selosse, and Julien Mairal. GraphiT: Encoding graph structure in transformers. *arXiv:2106.05667*, 2021.
- [42] Christopher Morris, Martin Ritzert, Matthias Fey, William L. Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. Weisfeiler and Leman go neural: Higher-order graph neural networks. In *The Thirty-Third AAAI Conference on Artificial Intelligence*, pages 4602–4609. AAAI Press, 2019.
- [43] Ryan Murphy, Balasubramaniam Srinivasan, Vinayak Rao, and Bruno Ribeiro. Relational pooling for graph representations. In *International Conference on Machine Learning*, pages 4663–4673. PMLR, 2019.
- [44] Kenta Oono and Taiji Suzuki. Graph neural networks exponentially lose expressive power for node classification. In *International Conference on Learning Representations*, 2020.
- [45] Wonpyo Park, Woonggi Chang, Donggeon Lee, Juntae Kim, and Seung won Hwang. GRPE: Relative positional encoding for graph transformer. *arXiv:22201.12787*, 2022.
- [46] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020.
- [47] Ryoma Sato. A survey on the expressive power of graph neural networks. *arXiv:2003.04078*, 2020.
- [48] Yu Shi, Shuxin Zheng, Guolin Ke, Yifei Shen, Jiacheng You, Jiyan He, Shengjie Luo, Chang Liu, Di He, and Tie-Yan Liu. Benchmarking graphormer on large-scale molecular modeling datasets. *arXiv:2203.04810*, 2022.
- [49] Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. Long range arena: A benchmark for efficient transformers. In *International Conference on Learning Representations*, 2021.
- [50] Jan Toenshoff, Martin Ritzert, Hinrikus Wolf, and Martin Grohe. Graph learning with 1d convolutions on random walks. *arXiv:2102.08786*, 2021.
- [51] Jake Topping, Francesco Di Giovanni, Benjamin Paul Chamberlain, Xiaowen Dong, and Michael M Bronstein. Understanding over-squashing and bottlenecks on graphs via curvature. *arXiv:2111.14522*, 2021.
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017.
- [53] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.
- [54] Sinong Wang, Belinda Z. Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv:2006.04768*, 2020.
- [55] Boris Weisfeiler and Andrei Leman. The reduction of a graph to canonical form and the algebra which appears therein. *NTI, Series*, 2(9):12–16, 1968.
- [56] Scott A Wildman and Gordon M Crippen. Prediction of physicochemical parameters by atomic contributions. *Journal of chemical information and computer sciences*, 39(5):868–873, 1999.
- [57] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019.
- [58] Mingqi Yang, Renjian Wang, Yanming Shen, Heng Qi, and Baocai Yin. Breaking the expression bottleneck of graph neural networks. *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [59] Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. Do transformers really perform badly for graph representation? In *Advances in Neural Information Processing Systems*, 2021.
- [60] Chengxuan Ying, Mingqi Yang, Shuxin Zheng, Guolin Ke, Shengjie Luo, Tianle Cai, Chenglin Wu, Yuxin Wang, Yanming Shen, and Di He. First place solution of KDD Cup 2021 & OGB large-scale challenge graph prediction track. *arXiv:2106.08279*, 2021.

- 487 [61] Jiaxuan You, Rex Ying, and Jure Leskovec. Design space for graph neural networks. In
488 *Advances in Neural Information Processing Systems*, 2020.
- 489 [62] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti,
490 Santiago Ontañón, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed.
491 Big Bird: Transformers for longer sequences. In *Advances in Neural Information Processing*
492 *Systems*, 2020.
- 493 [63] Muhan Zhang, Pan Li, Yinglong Xia, Kai Wang, and Long Jin. Labeling trick: A theory of using
494 graph neural networks for multi-node representation learning. *Advances in Neural Information*
495 *Processing Systems*, 34, 2021.
- 496 [64] Lingxiao Zhao, Wei Jin, Leman Akoglu, and Neil Shah. From stars to subgraphs: Uplifting any
497 GNN with local structure awareness. In *International Conference on Learning Representations*,
498 2022.

Checklist

1. For all authors...

- (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
- (b) Did you describe the limitations of your work? [Yes] We discuss limitations in Section 5
- (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section 5
- (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]

2. If you are including theoretical results...

- (a) Did you state the full set of assumptions of all theoretical results? [Yes]
- (b) Did you include complete proofs of all theoretical results? [Yes] We provide a sketch in Section 3.4 and more details in Appendix C.

3. If you ran experiments...

- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Code, data, and instructions are available in the supplementary material. We also include the performance traces/logs from our benchmarking experiments, supporting the results reported herein.
- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] We describe the datasets in Appendix A.1, splits in Appendix A.2, hyperparameters in Appendix A.3. Full configuration files are provided in the supplementary material.
- (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] We include standard deviations over several random seeds depending on the dataset evaluation protocol, more details are in Appendix A.2.
- (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] We elaborate on the compute and used resources in Appendix A.4.

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

- (a) If your work uses existing assets, did you cite the creators? [Yes]
- (b) Did you mention the license of the assets? [Yes] For datasets see Appendix A.1, for software see Appendix A.4.
- (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] The source code of GRAPHGPS is available in the supplementary material.
- (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
- (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] The used benchmarking datasets come from: the molecular domain with no personal information, anonymized source code (ogbg-code2 and MalNet-Tiny), anonymized images (MNIST, CIFAR10), or are statistically generated.

5. If you used crowdsourcing or conducted research with human subjects...

- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
- (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]