



GraphiT: Encoding Graph Structure in Transformers

Grégoire Mialon, Dexiong Chen, Margot Selosse, Julien Mairal

► To cite this version:

Grégoire Mialon, Dexiong Chen, Margot Selosse, Julien Mairal. GraphiT: Encoding Graph Structure in Transformers. 2021. hal-03256708

HAL Id: hal-03256708

<https://hal.science/hal-03256708>

Preprint submitted on 10 Jun 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

GraphiT: Encoding Graph Structure in Transformers

Grégoire Mialon*	Dexiong Chen*
Inria ^{†‡}	Inria [†]
<code>gregoire.mialon@inria.fr</code>	<code>dexiong.chen@inria.fr</code>
Margot Selosse*	Julien Mairal
Inria [†]	Inria [†]
<code>margot.selosse@inria.fr</code>	<code>julien.mairal@inria.fr</code>

June 10, 2021

Abstract

We show that viewing graphs as sets of node features and incorporating structural and positional information into a transformer architecture is able to outperform representations learned with classical graph neural networks (GNNs). Our model, GraphiT, encodes such information by (i) leveraging relative positional encoding strategies in self-attention scores based on positive definite kernels on graphs, and (ii) enumerating and encoding local sub-structures such as paths of short length. We thoroughly evaluate these two ideas on many classification and regression tasks, demonstrating the effectiveness of each of them independently, as well as their combination. In addition to performing well on standard benchmarks, our model also admits natural visualization mechanisms for interpreting graph motifs explaining the predictions, making it a potentially strong candidate for scientific applications where interpretation is important.¹

1 Introduction

Graph-structured data are present in numerous scientific applications and are the subject of growing interest. Examples of such data are as varied as proteins in computational biology [29], which may be seen as a sequence of amino acids, but also as a graph representing their tertiary structure, molecules in chemoinformatics [11], shapes in computer vision and computer graphics [38], electronic health records, or communities in social networks. Designing graph representations for machine learning is a particularly active area of research, even though not new [4], with a strong effort currently focused on graph neural networks [6, 7, 18, 28, 37, 40]. A major difficulty is to find graph representations that are computationally tractable, adaptable to a given task, and capable of distinguishing graphs with different topological structures and local characteristics.

In this paper, we are interested in the transformer, which has become the standard architecture for natural language processing [36], and is gaining ground in computer vision [9] and computational biology [27]. The ability of a transformer to aggregate information across long contexts for sequence data makes it an interesting challenger for GNNs that successively aggregate local information from neighbors in a multilayer fashion. In contrast, a single self-attention layer of the transformer can potentially allow all the nodes of an input graph to communicate. The price to pay is that this core component is invariant to permutations of the input nodes, hence does not take the topological structure of the graph into account and loses crucial information.

*Equal contribution.

[†]Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France.

[‡]D.I., UMR 8548, École Normale Supérieure, Paris, France.

¹Code available at <https://github.com/inria-thoth/GraphiT>.

For sequence data, this issue is addressed by encoding positional information of each token and giving it to the transformer architecture. For graphs, the problem is more challenging as there is no single way to encode node positions. For this reason, there have been few attempts to use vanilla transformers for graph representation. To the best of our knowledge, the closest work to ours seems to be the graph transformer architecture of Dwivedi and Bresson [12], who propose an elegant positional encoding strategy based on the eigenvectors of the graph Laplacian [1]. However, they focus on applying attention to neighboring nodes only, as in GNNs, and their results suggest that letting all nodes communicate is not a competitive approach.

Our paper provides another perspective and reaches a slightly different conclusion; we show that even though local communication is indeed often more effective, the transformer with global communication can also achieve good results. For that, we introduce a set of techniques to encode the local graph structure within our model, GraphiT (encoding **graph** structure in **transformers**). More precisely, GraphiT relies on two ingredients that may be combined, or used independently. First, we propose relative positional encoding strategies for weighting attention scores by using positive definite kernels, a viewpoint introduced for sequences in [35]. This concept is particularly appealing for graphs since it allows to leverage the rich literature on kernels on graphs [21, 33], which are powerful tools for encoding the similarity between nodes. The second idea consists in computing features encoding the local structure in the graph. To achieve this, we leverage the principle of graph convolutional kernel networks (GCKN) of [7], which consists in enumerating and encoding small sub-structure (for instance, paths or subtree patterns), which may then be used as an input to the transformer model.

We demonstrate the effectiveness of our approach on several classification and regression benchmarks, showing that GraphiT with global or local attention layer can outperform GNNs in various tasks, and also show that basic visualization mechanisms allow us to automatically discover discriminative graph motifs, which is potentially useful for scientific applications where interpretation is important.

2 Related work

Graph kernels. A classical way to represent graphs for machine learning tasks consists in defining a high-dimensional embedding for graphs, which may then be used to perform prediction with a linear models (*e.g.*, support vector machines). Graph kernels typically provide such embeddings by counting the number of occurrences of local substructures that are present within a graph [4]. The goal is to choose substructures leading to expressive representations sufficiently discriminative, while allowing fast algorithms to compute the inner-products between these embeddings. For instance, walks have been used for such a purpose [14], as well as shortest paths [5], subtrees [15, 24, 31], or graphlets [32]. Our work uses short paths, but other substructures could be used in principle. Note that graph kernels used in the context of comparing graphs, is a line of work different from the kernels on graphs that we will introduce in Section 3 for computing embeddings of nodes.

Graph neural networks. Originally introduced in [28], GNNs have been derived as an extension of convolutions for graph-structured data: they use a message passing paradigm in which vectors (messages) are exchanged (passed) between neighboring nodes whose representations are updated using neural networks. Many strategies have been proposed to aggregate features of neighboring nodes (see, *e.g.*, [6, 11]). The graph attention network (GAT) [37] is the first model to use an attention mechanism for aggregating local information. Recently, hybrid approaches between graph neural networks and graph kernels were proposed in [7, 10]. Diffusion processes on graphs that are related to the diffusion kernel we consider in our model were also used within GNNs in [20].

Transformers for graph-structured data. Prior to [12], there were some attempts to use transformers in the context of graph-structured data. The authors of [22] propose to apply attention to all nodes, yet without position encoding. In [42], a transformer architecture called Graph-BERT is fed with sampled subgraphs of a single graph in the context of node classification and graph clustering. They also propose to encode positions by aggregating different encoding schemes. However, these encoding schemes are either

impossible to use in our settings as they require having sampled subgraphs of regular structures as input, or less competitive than Laplacian eigenvectors as observed in [12]. The transformer model introduced in [41] needs to first transform an heterogeneous graph into a new graph structure via meta-paths, which does not directly operate on node features. To the best of our knowledge, our work is the first to demonstrate that vanilla transformers with appropriate node position encoding can compete with GNNs in graph prediction tasks.

3 Preliminaries about Kernels on Graphs

Spectral graph analysis. The Laplacian of a graph with n nodes is defined as $L = D - A$, where D is a $n \times n$ diagonal matrix that carries the degrees of each node on the diagonal and A is the adjacency matrix. Interestingly, L is a positive semi-definite matrix such that for all vector u in \mathbb{R}^n , $u^\top L u = \sum_{i \sim j} (u[i] - u[j])^2$, which can be interpreted as the amount of “oscillation” of the vector u , hence its “smoothness”, when seen as a function on the nodes of the graph.

The Laplacian is often used via its eigenvalue decomposition $L = \sum_i \lambda_i u_i u_i^\top$, where the eigenvalue $\lambda_i = u_i^\top L u_i$ characterizes the amount of oscillation of the corresponding eigenvector u_i . For this reason, this decomposition is traditionally viewed as the discrete equivalent to the sine/cosine Fourier basis in \mathbb{R}^n and associated frequencies. Note that very often the normalized Laplacian $I - D^{-\frac{1}{2}} A D^{-\frac{1}{2}}$ is used instead of L , which does not change the above interpretation.

Interestingly, it is then possible to define a whole family of positive definite kernels on the graph [33] by applying a regularization function r to the spectrum of L . We therefore get a rich class of kernels

$$K_r = \sum_{i=1}^m r(\lambda_i) u_i u_i^\top, \quad (1)$$

associated with the following norm $\|f\|_r^2 = \sum_{i=1}^m (f_i^\top u_i)^2 / r(\lambda_i)$ from a reproducing kernel Hilbert space (RKHS), where $r : \mathbb{R} \mapsto \mathbb{R}_*^+$ is a non-increasing function such that smoother functions on the graph would have smaller norms in the RKHS. We now introduce two examples.

Diffusion Kernel [21]. It corresponds to the case $r(\lambda_i) = e^{-\beta \lambda_i}$, which gives:

$$K_D = \sum_{i=1}^m e^{-\beta \lambda_i} u_i u_i^\top = e^{-\beta L} = \lim_{p \rightarrow +\infty} \left(I - \frac{\beta}{p} L \right)^p. \quad (2)$$

The diffusion kernel can be seen as a discrete equivalent of the Gaussian kernel, a solution of the heat equation in the continuous setting, hence its name. Intuitively, the diffusion kernel between two nodes can be interpreted as the quantity of some substance that would accumulate at the first node after a given amount of time (controlled by β) if we injected the substance at the other node and let it diffuse through the graph. It is related to the random walk kernel that will be presented below.

p-step random walk kernel. By taking $r(\lambda_i) = (1 - \gamma \lambda_i)^p$, we obtain a kernel that admits an interpretation in terms of p steps of a particular random walk on the graph:

$$K_{pRW} = (I - \gamma L)^p. \quad (3)$$

It is related to the diffusion kernel by choosing $\gamma = \beta/p$ and taking the limit with $p \rightarrow +\infty$, according to (2). Unlike the diffusion kernel which yields a dense matrix, the random walk kernel is sparse and has limited connectivity, meaning that two nodes are connected in K_{pRW} only if there exists a path of length p between them. As these kernels on graphs reflect the structural similarity of the nodes independently of their features, it is natural to use the resulting Gram matrix to encode such a structure within the transformer model, as detailed next.

4 Encoding Graph Structure in Transformers

In this section, we detail the architecture and structure encoding strategies behind GraphiT. In particular, we build upon the findings of [35] for our architectural choices and propose new strategies for encoding structural information in the transformer architecture.

4.1 Transformer Architectures for Graphs

We process graphs with a vanilla transformer encoder architecture [36] for solving classification and regression tasks, by seeing graphs as sets of node features. We first present the transformer without encoding the graph structure, before introducing mechanisms for encoding nodes positions in Section 4.2, and then topological structures in Section 4.3. Specifically, a transformer is composed of a sequence of layers, which process an input set of d_{in} features represented by a matrix X in $\mathbb{R}^{n \times d_{\text{in}}}$, and compute another set in $\mathbb{R}^{n \times d_{\text{out}}}$. A critical component is the attention mechanism:

$$\text{Attention}(Q, V) = \text{softmax}\left(\frac{Q Q^\top}{\sqrt{d_{\text{out}}}}\right) V \in \mathbb{R}^{n \times d_{\text{out}}}, \quad (4)$$

with $Q^\top = W_Q X^\top$ is called the query matrix, $V^\top = W_V X^\top$ the value matrix, and W_Q, W_V in $\mathbb{R}^{d_{\text{out}} \times d_{\text{in}}}$ are projection matrices that are learned. Note that following the recommendation of [35], we use the same matrix for keys and queries. This allows us to define a symmetric and positive definite kernel on pairs of nodes and thus to combine with other kernels on graph. This also reduces the number of parameters in our models without hurting the performance in practice. During the forward pass, the feature map X is updated in a residual fashion (with either $d_{\text{in}} = d_{\text{out}}$ or with an additional projection matrix when the dimensions do not match, omitted here for simplicity) as follows:

$$X = X + \text{Attention}(Q, V).$$

Note that transformers without position encoding and GNNs are tightly connected: a GNN can be seen as a transformer where aggregation is performed on neighboring nodes only, and a transformer as a GNN processing a fully-connected graph. However, even in this case, differences would remain between common GNNs and Transformers, as the latter use for example LayerNorm and skip-connections. In our paper, we will either adopt the local aggregation strategy, or let all nodes communicate in order to test this inductive bias in the context of graphs, where capturing long range interactions between nodes may be useful for some tasks.

4.2 Encoding Node Positions

The output of the transformer is invariant to permutations in the input data. It is therefore crucial to provide the transformer with information about the data structure. In this section, we revisit previously proposed strategies for sequences and graphs and devise new methods for positional encoding. Note that a natural baseline is simply to adopt a local aggregation strategy similar to GAT [37] or [12].

4.2.1 Existing Strategies for Sequences and Graphs

Absolute and relative positional encoding in transformers for sequences. In NLP, positional information was initially encoded by adding a vector based on sinusoidal functions to each of the input token features [36]. This approach was coined as absolute position encoding and proved to be also useful for other data modalities [9, 12]. In relative position encoding, which was proposed later [26, 30], positional information is added to the attention logits and in the values. This information only depends on the relative distance between the two considered elements in the input sequence.

Absolute positional encoding for graphs. Whereas positional encodings can be hand-crafted or learned relatively easily for sequences and images, which respectively admit a chain or grid structure, this task becomes non-trivial for graphs, whose structure may vary a lot inside a data set, besides the fact that the concept of node position is ill-defined. To address these issues, an absolute position encoding scheme was recently proposed in [12], by leveraging the graph Laplacian (LapPE). More precisely, each node of each graph is assigned a vector containing the first k coordinates of the node in the eigenbasis of the graph normalized Laplacian sorted in ascending order of the eigenvalues. Since the first eigenvector associated to the eigenvalue 0 is constant, the first coordinate is omitted.

As detailed in Section 3, these eigenvectors oscillate more and more and the corresponding coordinates are often interpreted as Fourier coefficients representing frequency components in increasing order. Note that eigenvectors of the Laplacian computed on different graphs could not be compared to each other in principle, and are also only defined up to a ± 1 factor. While this raises a conceptual issue for using them in an absolute positional encoding scheme, it is shown in [12]—and confirmed in our experiments—that the issue is mitigated by the Fourier interpretation, and that the coordinates used in LapPE are effective in practice for discriminating between nodes in the same way as the position encoding proposed in [36] for sequences. Yet, because the eigenvectors are defined up to a ± 1 factor, the sign of the encodings needs to be randomly flipped during the training of the network. In the next section, we introduce a novel strategy in the spirit of relative positional encoding that does not suffer from the previous conceptual issue, and can also be combined with LapPE if needed.

4.2.2 Relative Position Encoding Strategies by Using Kernels on Graphs

Modulating the transformer attention. To avoid the issue of transferability of the absolute positional encoding between graphs, we use information on the nodes structural similarity to bias the attention scores. More precisely, and in the fashion of [25, 35] for sequences, we modulate the attention kernel using the Gram matrix of some kernel on graphs described in Section 3 as follows:

$$\text{PosAttention}(Q, V, K_r) = \text{normalize} \left(\exp \left(\frac{Q Q^\top}{\sqrt{d_{\text{out}}}} \right) \odot K_r \right) V \in \mathbb{R}^{n \times d_{\text{out}}}, \quad (5)$$

with the same Q and V matrices, and K_r a kernel on the graph. “exp” denotes the elementwise exponential of the matrix entries and normalize means ℓ_1 -normalization on rows such that $\text{normalization}(\exp(u)) = \text{softmax}(u)$. The reason for multiplying the exponential of the attention logits before ℓ_1 -normalizing the rows is that it corresponds to a classical kernel smoothing. Indeed, if we consider the PosAttention output for a node i :

$$\text{PosAttention}(Q, V, K_r)_i = \sum_{j=1}^n \frac{\exp(Q_i Q_j^\top / \sqrt{d_{\text{out}}}) \times K_r(i, j)}{\sum_{j'=1}^n \exp(Q_i Q_{j'}^\top / \sqrt{d_{\text{out}}}) \times K_r(i, j')} V_j \in \mathbb{R}^{d_{\text{out}}},$$

we obtain a typical smoothing, *i.e.*, a linear combination of features with weights determined by a non-negative kernel, here $k(i, j) := \exp(Q_i Q_j^\top / \sqrt{d_{\text{out}}}) \times K_r(i, j)$, and summing to 1. In fact, k can be considered as a new kernel between nodes i and j made by multiplying a kernel based on positions (K_r) and a kernel based on content (via Q). As observed in [35] for sequences, modulating the attention logits with a kernel on positions is related to relative positional encoding [30], where we bias the attention matrix with a term depending only on the relative difference in positions between the two elements in the input set. Moreover, during the forward pass, the feature map X is updated as follows:

$$X = X + D^{-\frac{1}{2}} \text{PosAttention}(Q, V, K_r), \quad (6)$$

where D is the matrix of node degrees. We found such a normalization with $D^{-1/2}$ to be beneficial in our experiments since it reduces the overwhelming influence of highly connected graph components. Note that, as opposed to absolute position encoding, we do not add positional information to the values and the model does not rely on the transferability of eigenvectors between different Laplacians.

Choice of kernels and parameters. Interestingly, the choice of the kernel enables to encode a priori knowledge directly within the model architecture, while the parameter has an influence on the attention span. For example, in the diffusion kernel, β can be seen as the duration of the diffusion process. The smaller it is, the more focused the attention on the close neighbors. Conversely, a large β corresponds to an homogeneous attention span. As another example, it is clear that the choice of p in the random walk kernel corresponds to visiting at best p -degree neighbors. In our experiments, the best kernel may vary across datasets, suggesting that long-range global interactions are of different importance depending on the task. For example, on the dataset PROTEINS (see Section 5), the vanilla transformer without using any structural information performs very well.

4.3 Encoding Topological Structures

Position encoding aims at adding positional only information to the feature vector of an input node or to the attentions scores. Substructures is a different type of information, carrying local positional information and content, which has been heavily used within graph kernels, see Section 2. In the context of graphs neural networks, this idea was exploited in the graph convolutional kernel network model (GCKN) of [7], which is a hybrid approach between GNNs and graph kernels based on substructure enumeration (*e.g.*, paths). Among different strategies we experimented, enriching nodes features with the output of a GCKN layer turned out to be a very effective strategy.

Graph convolutional kernel networks (GCKN). GCKNs [7] is a multi-layer model that produces a sequence of graph feature maps akin to a GNN. The main difference is that each layer enumerates local sub-structures at each node (here, paths of length k), encodes them using a kernel embedding, and aggregates the resulting sub-structure representations. This results in a feature map that carries more information about the graph structure than traditional neighborhood aggregation based GNNs, which is appealing for transformers since the vanilla version is blind to the graph structure.

Formally, let us consider a graph G with n nodes, and let us denote by $\mathcal{P}_k(u)$ the set of paths shorter than or equal to k that start with node u . With an abuse of notation, p in $\mathcal{P}_k(u)$ will denote the concatenation of all node features encountered along the path. Then, a layer of GCKN defines a feature map X in $\mathbb{R}^{n \times d}$ such that

$$X(u) = \sum_{p \in \mathcal{P}_k(u)} \psi(p),$$

where $X(u)$ is the column of X corresponding to node u and ψ is a d -dimensional embedding of the path features p . More precisely, the path features in [7] are embedded to a RKHS by using a Gaussian kernel, and a finite-dimensional approximation is obtained by using the Nyström method [39]. The embedding is parametrized and can be learned without or with supervision (see [7] for details). Moreover, path features of varying lengths up to a maximal length can be used with GCKN. In this work, we evaluate the strategy of encoding a node as the concatenation of its original features and those produced by one GCKN layer. This strategy has proven to be very successful in practice.

5 Experiments

In this section, we evaluate instances of GraphiT as well as popular GNNs and position encoding baselines on various graph classification and regression tasks. We want to answer several questions:

- Q1:** Can vanilla transformers, when equipped with appropriate position encoding and/or structural information, outperform GNNs in graph classification and regression tasks?
- Q2:** Is kernel-based relative positional encoding more effective than the absolute position encoding provided by the eigenvectors of the Laplacian (LapPE)?
- Q3:** What is the most effective way to encode graph structure information within transformers?

We also discuss our results and conduct ablation studies. Finally, we demonstrate the ability of attention scores to highlight meaningful graph features when using kernel-based positional encoding.

5.1 Methodology

Benchmark and baselines. We benchmark our methods on various graph classification datasets with discrete node labels (MUTAG, PROTEINS, PTC, NCI1) and one regression dataset with discrete node labels (ZINC). These datasets can be obtained *e.g* via the Pytorch Geometric toolbox [13]. We compare our models to the following GNN models: Molecular Fingerprint (MF) [11], Graph Convolutional Networks (GCN) [18], Graph Attention Networks (GAT) [37], Graph Isomorphism Networks (GIN) [40] and finally Graph Convolutional Kernel Networks (GCKN) [7]. In particular, GAT is an important baseline as it uses attention to aggregate neighboring node information. We compare GraphiT to the transformer architecture proposed in [12] and also use their Laplacian absolute position encoding as a baseline for evaluating our graph kernel relative position encoding. All models are implemented in Pytorch and our code is available in the supplementary material.

Reporting scores. For all datasets except ZINC, we samples ten times random train/val/test splits, of size 80/10/10, respectively. For each split, we evaluate all methods by (i) training several models on the train fold with various hyperparameters; (ii) performing model selection on val, by averaging the validation accuracy of a model on its last 50 epochs; (iii) retraining the selected model on train+val; (iv) estimate the test score by averaging the test accuracy over the last 50 epochs. The results reported in our tables are then averaged over the ten splits. This procedure is a compromise between a double-nested cross validation procedure that would be too computationally expensive and reporting 10-fold cross validation scores that would overestimate the test accuracy. For ZINC, we use the same train/val/test splits as in [12], train GraphiT with 10 layers, 8 heads and 64 hidden dimensions as in [12], and report the average test mean absolute error on 4 independent runs.

Optimization procedure and hyperparameter search. Our models are trained with the Adam optimizer by decreasing the learning rate by a factor of 2 each 50 epochs. For classification tasks, we train about the same number (81) of models with different hyperparameters for each GNN and transformer method, thus spending a similar engineering effort on each method. For GNN models, we select the best type of global pooling, number of layers and hidden dimensions from three different values. Regarding transformers, we select the best number of heads instead of global pooling type for three different values. For all considered models, we also select the best learning rate and weight decay from three different values and the number of epochs is fixed to 300 to guarantee the convergence. For the ZINC dataset, we found that a standard warmup strategy suggested for transformers in [36] leads to more stable convergence for larger models. The rest of the hyperparameters remains the same as used in [12]. More details and precise grids for hyperparameter search can be found in Appendix A.

5.2 Results and Discussion

Comparison of GraphiT and baselines methods. We show our results in Table 1. For smaller datasets such as MUTAG, PROTEINS or PTC, our Transformer without positional encoding performs reasonably well compared to GNNs, whereas for NCI1 and ZINC, incorporating structural information into the model is key to good performance. On all datasets, GraphiT is able to perform as well as or better than the baseline GNNs. In particular on ZINC, GraphiT outperforms all previous baseline methods by a large margin. For this, it seems that the factor $D^{-1/2}$ in (6) is important, allowing to capture more information about the graph structure. The answer to **Q1** is therefore positive.

Comparison of relative position encoding schemes. Here, we compare our transformer used with different relative positional encoding strategies, including adjacency (1-step RW kernel with $\gamma = 1.0$ corresponding to a normalized adjacency matrix $D^{-1/2}AD^{-1/2}$) which is symmetric but not positive semi-definite, 2 and 3-step RW kernel with $\gamma = 0.5$ and a diffusion kernel with $\beta = 1$. Unlike the vanilla transformer that works poorly on big datasets including NCI1 and ZINC, keeping all nodes communicate through our diffusion kernel positional encoding can still achieve performances close to encoding methods relying on local

Table 1: Average mean classification accuracy/mean absolute error.

Method / Dataset	MUTAG	PROTEINS	PTC	NCI1	ZINC (no edge feat.)
Size	188	1113	344	4110	12k
Classes	2	2	2	2	Reg.
Max. number of nodes	28	620	109	111	37
MF [11]	81.5±11.0	71.9±5.2	57.3±6.9	80.6±2.5	0.387±0.019
GCN [18]	78.9±10.1	75.8±5.5	54.0±6.3	75.9±1.6	0.367±0.011
GAT [37]	80.3±8.5	74.8±4.1	55.0±6.0	76.8±2.1	0.384±0.007
GIN [40]	82.6±6.2	73.1±4.6	55.0±8.7	81.7±1.7	0.387±0.015
GCKN-subtree [7]	87.8±9.4	72.0±3.7	62.1±6.4	79.6±1.8	0.474±0.001
[12]	79.3±11.6	65.8±3.1	58.4±8.2	78.9±1.1	0.359±0.014
[12] + LapPE	83.9±6.5	70.1±3.2	57.7±3.1	80.0±1.9	0.323±0.013
Transformers (T)	82.2±6.3	75.6±4.9	58.1±10.5	70.0±4.5	0.696±0.007
T + LapPE	85.8±5.9	74.6±2.7	55.6±5.0	74.6±1.9	0.507±0.003
T + Adj PE	87.2±9.8	72.4±4.9	59.9±5.9	79.7±2.0	0.243±0.005
T + 2-step RW kernel	85.3±6.9	72.8±4.5	62.0±9.4	78.0±1.5	0.243±0.010
T + 3-step RW kernel	83.3±6.3	76.2±4.4	61.0±6.2	77.6±3.6	0.244±0.011
T + Diffusion kernel	82.7±7.6	74.6±4.2	59.1±7.4	78.9±1.6	0.255±0.010
T + GCKN	84.4±7.8	69.5±3.8	61.5±5.8	78.1±5.1	0.274±0.011
T + GCKN + 2-step RW kernel	90.4±5.8	72.5±4.6	58.4±7.6	81.0±1.8	0.213±0.016
T + GCKN + Adj PE	90.5±7.0	71.1±6.9	57.9±4.2	81.4±2.2	0.211±0.010

communications such as adjacency or p-step kernel encoding. Interestingly, our adjacency encoding, which could be seen as a variant of the neighborhood aggregation of node features used in GAT, is shown to be very effective on many datasets. In general, sparse local positional encoding seems to be useful for our prediction tasks, which tempers the answer to **Q1**.

Comparison of structure encoding schemes in node features. In this paragraph, we compare different ways of injecting graph structures to the vanilla transformer, including Laplacian PE [12] and unsupervised GCKN-path features [7]. We observe that incorporating topological structures directly into the features of input nodes is a very useful strategy for vanilla transformers. This yields significant performance improvement on almost all datasets except PROTEINS, by either using the Laplacian PE or GCKN-path features. Among them, GCKN features are observed to outperform Laplacian PE by a large margin, except on MUTAG and PROTEINS (third column of Table 2). A possible reason for this exception is that PROTEINS seems to be very specific, such that prediction models do not really benefit from encoding positional information. In particular, GCKN brings a pronounced performance boost on ZINC, suggesting the importance of encoding substructures like paths for prediction tasks on molecules.

Combining relative position encoding and structure encoding in node features. Table 2 reports the ablation study of the transformer used with or without structure encoding and coupled or not with a relative positional encoding. The results show that relative PE outperforms the topological Laplacian PE, suggesting a positive answer to **Q2**. However, using both simultaneously improves the results considerably, especially on ZINC. In fact, combining relative position encoding and a structure encoding scheme globally improves the performances. In particular, using the GCKN-path layer features works remarkably well for all datasets except PROTEINS. More precisely, we see that the combination of GCKN with the adjacent matrix PE yields the best results among the other combinations for MUTAG and NCI1. In addition, the GCKN coupled with the 3-step RW kernel achieves the second best performance for ZINC. The answer to **Q3** is therefore combining a structure encoding scheme in node features, such as GCKN, with relative positional encoding.

Table 2: Ablation: comparison of different structural encoding schemes and their combinations.

Dataset	Structure encoding in node features	Relative positional encoding in attention scores				
		T vanilla	T + Adj	T + 2-step	T + 3-step	T + Diffusion
MUTAG	None	82.2±6.3	87.2±9.8	85.3±6.9	83.3±6.3	82.7±7.6
	LapPE	85.8±5.9	86.0±4.2	84.7±4.7	83.5±5.2	84.2±7.2
	GCKN	84.4±7.8	90.5±7.0	90.4±5.8	90.0±6.3	90.0±6.8
PROTEINS	None	75.6±4.9	72.4±4.9	72.8±4.5	76.2±4.4	74.6±4.2
	LapPE	74.6±2.7	74.7±5.2	75.0±4.7	74.3±5.3	74.7±5.3
	GCKN	69.5±3.8	71.1±6.9	72.5±4.6	70.0±5.1	72.4±4.9
PTC	None	58.1±10.5	59.9±5.9	62.0±9.4	61.0±6.2	59.1±7.4
	LapPE	55.6±5.0	57.1±3.8	58.8±6.6	57.1±5.3	57.3±7.8
	GCKN	61.5±5.8	57.9±4.2	58.4±7.6	55.2±8.8	55.9±8.1
NCI1	None	70.0±4.5	79.7±2.0	78.0±1.5	77.6±3.6	78.9±1.6
	LapPE	74.6±1.9	78.7±1.9	78.4±1.3	78.7±1.5	77.8±1.0
	GCKN	78.1±5.1	81.4±2.2	81.0±1.8	81.0±1.8	81.0±2.0
ZINC	None	0.696±0.007	0.243±0.005	0.243±0.010	0.244±0.011	0.255±0.010
	LapPE	0.507±0.003	0.202±0.011	0.227±0.030	0.210±0.003	0.221±0.038
	GCKN	0.274±0.011	0.211±0.010	0.213±0.016	0.203±0.011	0.218±0.006

Discussion. Combining transformer and GCKN features results in substantial improvement over the simple sum or mean global poolings used in original GCKN models on ZINC dataset as shown in Table 1, which suggests that transformers can be considered as a very effective method for aggregating local substructure features on large datasets at the cost of using much more parameters. This point of view has also been explored in [25], which introduces a different form of attention for sequence and text classification. A potential limitation of GraphiT is its application to large graphs, as the complexity of the self-attention mechanism scales quadratically with the size of the input sequence. However, a recent line of work coined as efficient transformers alleviated these issues both in terms of memory and computational cost [2, 19]. We refer the reader to the following survey [34].

5.3 Visualizing attention scores for the Mutagenicity dataset.

We now show that the attention yields visualization mechanisms for detecting important graph motifs.

Mutagenicity. In chemistry, a mutagen is a compound that causes genetic mutations. This important property is related to the molecular substructures of the compound. The Mutagenicity dataset [17] contains 4337 molecules to be classified as mutagen or not, and aims at better understanding which substructures cause mutagenicity. We train GraphiT with diffusion position encoding on this dataset with the aim to study whether important substructures can be detected. To this end, we feed the model with molecules of the dataset and collect the attention scores of each layer averaged by heads, as can be seen in Figure 2 for the molecules of Figure 1.

Patterns captured in the attention scores. While the scores in the first layer are close to the diffusion kernel, the following layers get sparser. Since the attention matrix is multiplied on the right by the values, the coefficients of the node aggregation step for the n -th node is given by the n -th *row* of the attention matrix. As a consequence, salient *columns* suggest the presence of important nodes. After verification, for many samples of Mutagenicity fed to our model, salient atoms indeed represent important groups in chemistry, many of them being known for causing mutagenicity. In compound 1a, the N atom of the nitro group (NO_2) is salient. 1a was correctly classified as a mutagen by our model and the nitro group is indeed known for its mutagenetic properties [8]. Note how information flows from O atoms to N in the first two layers and then, at the last layer, how every element of the carbon skeleton look at N, *i.e* the NO_2 group. We can apply the

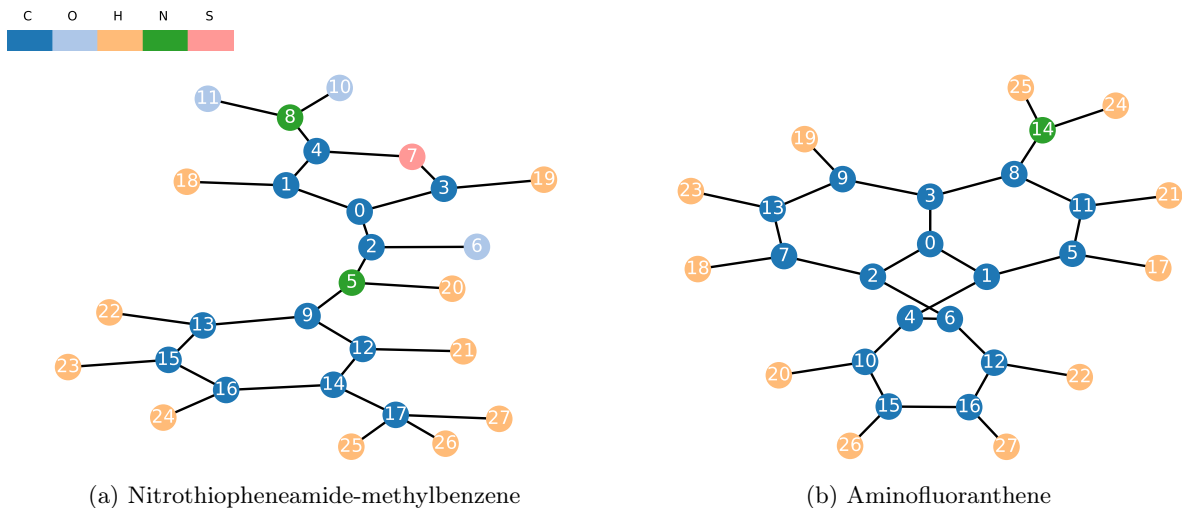


Figure 1: Examples of molecules from Mutagenicity correctly classified as mutagenetic by our model.




Figure 2: Attention scores averaged by heads for each layer of our trained model for the compounds in Figures 1a (*Top*) and 1b (*Bottom*). *Top Left*: diffusion kernel for 1a. *Top Right*: node 8 (N of NO₂) is salient. *Bottom Left*: diffusion kernel for 1b. *Bottom Right*: node 14 (N of NH₂) is salient.

same reasoning for the amino group (NH_2) in compound 1b [3]. We were also able to identify long-range intramolecular hydrogen bounds such as between H and Cl in other samples. We provide more examples of visualization in Appendix B.

6 Conclusion

In this work, we show that using a transformer to aggregate local substructures with appropriate position encoding, GraphiT, is a very effective strategy for graph representation, and that attention scores allow simple model interpretation. One of the reasons for the success of transformers lies in their scaling properties: in language modeling for example, it has been shown that for an equal large amount of data, the transformer’s loss follows a faster decreasing power law than Long-Short-Term-Memory networks when the size of the model increases [16]. We believe that an interesting future direction for this work would be to evaluate if GNNs and GraphiT also scale similarly in the context of large self-supervised pre-training, and can achieve a similar success as in natural language processing.

Acknowledgments

GM, DC, MS, and JM were supported by the ERC grant number 714381 (SOLARIS project) and by ANR 3IA MIAI@Grenoble Alpes, (ANR-19-P3IA-0003). This work was granted access to the HPC resources of IDRIS under the allocation 2021-AD011012521 made by GENCI. GM thanks Robin Strudel for useful discussions on transformers.

References

- [1] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.
- [2] I. Beltagy, M. E. Peters, and A. Cohan. Longformer: The long-document transformer. *arXiv:2004.05150*, 2020.
- [3] D. L. Berry, G. M. Schoofs, and W. A. Vance. Mutagenicity of nitrofluoranthenes, 3-aminofluoranthene and 1-nitropyrene in Chinese hamster V79 cells. *Carcinogenesis*, 6(10):1403–1407, 1985.
- [4] K. Borgwardt, E. Ghisu, F. Llinares-López, L. O’Bray, and B. Rieck. Graph kernels: State-of-the-art and future challenges. *arXiv preprint arXiv:2011.03854*, 2020.
- [5] K. Borgwardt and H.-P. Kriegel. Shortest-path kernels on graphs. In *International conference on data mining (ICDM)*, 2005.
- [6] M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.
- [7] D. Chen, L. Jacob, and J. Mairal. Convolutional kernel networks for graph-structured data. In *International Conference on Machine Learning (ICML)*, 2020.
- [8] K.-T. Chung, C. A. Murdock, Y. Zhou, S. E. Stevens Jr., Y.-S. Li, C.-I. Wei, S. Y. Fernando, and M.-W. Chou. Effects of the nitro-group on the mutagenicity and toxicity of some benzamines. *Environmental and Molecular Mutagenesis*, 27(1):67–74, 1996.
- [9] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (ICLR)*, 2021.

- [10] S. S. Du, K. Hou, R. R. Salakhutdinov, B. Póczos, R. Wang, and K. Xu. Graph neural tangent kernel: Fusing graph neural networks with graph kernels. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [11] D. Duvenaud, D. Maclaurin, J. Aguilera-Iparraguirre, R. Gómez-Bombarelli, T. Hirzel, A. Aspuru-Guzik, and R. P. Adams. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.
- [12] V. P. Dwivedi and X. Bresson. A generalization of transformer networks to graphs. *AAAI Workshop on Deep Learning on Graphs: Methods and Applications*, 2021.
- [13] M. Fey and J. E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- [14] T. Gärtner, P. Flach, and S. Wrobel. On graph kernels: Hardness results and efficient alternatives. In *Learning theory and kernel machines*, pages 129–143. Springer, 2003.
- [15] Z. Harchaoui and F. Bach. Image classification with segmentation graph kernels. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007.
- [16] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei. Scaling laws for neural language models, 2020.
- [17] K. Kersting, N. M. Kriege, C. Morris, P. Mutzel, and M. Neumann. Benchmark data sets for graph kernels, 2016.
- [18] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- [19] N. Kitaev, Łukasz Kaiser, and A. Levskaya. Reformer: The efficient transformer. In *International Conference on Learning Representations (ICLR)*, 2020.
- [20] J. Klicpera, S. Weißenberger, and S. Günnemann. Diffusion improves graph learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [21] R. Kondor and J.-P. Vert. Diffusion kernels. In *Kernel Methods in Computational Biology*, pages 171–192. MIT Press, 2004.
- [22] Y. Li, X. Liang, Z. Hu, Y. Chen, and E. P. Xing. Graph transformer, 2019.
- [23] M. Låg, J. G. Omichinski, E. Dybing, S. D. Nelson, and E. J. SÅžderlund. Mutagenic activity of halogenated propanes and propenes: effect of bromine and chlorine positioning. *Chemico-Biological Interactions*, 93(1):73–84, 1994.
- [24] P. Mahé and J.-P. Vert. Graph kernels based on tree patterns for molecules. *Machine learning*, 75(1):3–35, 2009.
- [25] G. Mialon, D. Chen, A. d’Aspremont, and J. Mairal. A trainable optimal transport embedding for feature aggregation and its relationship to attention. In *International Conference on Learning Representations (ICLR)*, 2021.
- [26] A. Parikh, O. Täckström, D. Das, and J. Uszkoreit. A decomposable attention model for natural language inference. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016.
- [27] A. Rives, J. Meier, T. Sercu, S. Goyal, Z. Lin, J. Liu, D. Guo, M. Ott, C. L. Zitnick, J. Ma, and R. Fergus. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *bioRxiv*, 2019.

- [28] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008.
- [29] A. W. Senior, R. Evans, J. Jumper, J. Kirkpatrick, L. Sifre, T. Green, et al. Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792):706–710, 2020.
- [30] P. Shaw, J. Uszkoreit, and A. Vaswani. Self-attention with relative position representations. In *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2018.
- [31] N. Shervashidze, P. Schweitzer, E. J. v. Leeuwen, K. Mehlhorn, and K. M. Borgwardt. Weisfeiler-Lehman graph kernels. *Journal of Machine Learning Research (JMLR)*, 12:2539–2561, 2011.
- [32] N. Shervashidze, S. Vishwanathan, T. Petri, K. Mehlhorn, and K. Borgwardt. Efficient graphlet kernels for large graph comparison. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2009.
- [33] A. J. Smola and R. Kondor. Kernels and regularization on graphs. In B. Schölkopf and M. K. Warmuth, editors, *Learning Theory and Kernel Machines*, pages 144–158. Springer Berlin Heidelberg, 2003.
- [34] Y. Tay, M. Dehghani, D. Bahri, and D. Metzler. Efficient transformers: A survey, 2020.
- [35] Y.-H. H. Tsai, S. Bai, M. Yamada, L.-P. Morency, and R. Salakhutdinov. Transformer dissection: A unified understanding of transformer’s attention via the lens of kernel. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2019.
- [36] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [37] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph attention networks. In *International Conference on Learning Representations (ICLR)*, 2018.
- [38] N. Verma, E. Boyer, and J. Verbeek. Feastnet: Feature-steered graph convolutions for 3d shape analysis. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [39] C. K. Williams and M. Seeger. Using the Nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2001.
- [40] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations (ICLR)*, 2019.
- [41] S. Yun, M. Jeong, R. Kim, J. Kang, and H. Kim. Graph transformer networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [42] J. Zhang, H. Zhang, C. Xia, and L. Sun. Graph-bert: Only attention is needed for learning graph representations, 2020.

Appendix

A Experimental Details

In this section, we provide implementation details and additional experimental results.

A.1 General Details.

Computing infrastructure. Computations were done on a GPU cluster equipped with Tesla V100-16G and Tesla V100-32G. We have monitored precisely the entire computational cost of this research project (including preliminary experiments, designing early and final models, evaluating baselines, running final experiments), which was approximately 20k GPU hours.

Position encoding and structure encoding parameters for all datasets. γ for p-step random walk kernels is fixed to 0.5 for both 2- and 3-step random walk kernels. β for the diffusion kernel is fixed to 1.0. Regarding the structure encoding in node features, the dimension of Laplacian positional encoding is set to 8 for ZINC as suggested by [12] and to 2 for graph classification datasets. The path size, the filter number and the bandwidth parameter of the unsupervised GCKN path features (used for structure encoding) are set to 8, 32 and 0.6 respectively for the ZINC dataset whereas the path size is set to 5 for graph classification datasets.

Other details. For all instances of GraphiT, the hidden dimensions of the feed-forward hidden layer in each layer of the encoder are fixed to twice of the dimensions of the attention.

A.2 Graph Classification Datasets

Here, we provide experimental details for MUTAG, PROTEINS, PTC, NCI1 and Mutagenicity datasets.

Datasets. These free datasets were collected by [17] for academic purpose. MUTAG consists in classifying molecules into mutagenetic or not. PROTEINS consists in classifying proteins into enzymes and non-enzymes. PTC consists in classifying molecules into carcinogenetic or not. NCI1 consists in classifying molecules into positive or negative to cell lung cancer. Mutagenicity consists in classifying compounds into mutagenetic or not.

Training splits. For MUTAG, PROTEINS, PTC and NCI1, our outer splits correspond to the splits used in [40]. Our inner splits that divide their train splits into our train and validation splits are provided in our code. The error bars are computed as the standard deviation of the test accuracies across the 10 outer folds. On the other side, as the purpose of using Mutagenicity is model interpretation, we use simple train, validation and test splits respectively composed of 80%, 10%, and 10% of the whole dataset.

Hyperparameter choices. For smaller datasets (MUTAG, PTC, PROTEINS), we use the search grids in Table 3 for GNNs, the search grids in Table 4 for GCKN and the search grids in Table 5 for both GraphiT and the transformer model of [12]. For bigger datasets (NCI1), we use the search grid in Table 3 for GNNs, the search grids in Table 4 for GCKN and the search grids in Table 6 for both GraphiT and the transformer model of [12]. The best models are selected based on the validation scores that are computed as the average of the validation scores over the last 50 epochs. Then, the selected model is retrained on train plus validation sets and the average of the test scores over the last 50 epochs is reported. For Mutagenicity, we train a GraphiT

Table 3: Parameter grid for GNNs trained on MUTAG, PROTEINS, PTC, NCI1.

Parameter	Grid
Layers	[1, 2, 3]
Hidden dimension	[64, 128, 256]
Global pooling	[sum, max, mean]
Learning rate	[0.1, 0.01, 0.001]
Weight decay	[0.1, 0.01, 0.001]

Table 4: Parameter grid for GCKN trained on MUTAG, PROTEINS, PTC, NCI1.

Parameter	Grid
Path size	[3, 5, 7]
Hidden dimension	[16, 32, 64]
Sigma	[0.5]
Global pooling	[sum, max, mean]
Learning rate	[0.001]
Weight decay	[0.01, 0.001, 0.0001]

model with 3 layers, 4 heads, 64 hidden dimensions. Initial learning rate was set to 1e-3, weight-decay was fixed to 1e-4 and structural information was encoded only via relative position encoding with the diffusion kernel.

Optimization. We use the cross-entropy loss and Adam optimizer with batch size 128 for GNNs and 32 for both GraphiT and the transformer model of [12]. For transformers, we do not observe significant improvement using warm-up strategies on these classification tasks. Thus, we simply follow a scheduler that halves the learning rate after every 50 epochs, as for GNNs. All models are trained for 300 epochs.

A.3 Graph Regression Dataset

Here, we present experimental details and additional experimental results for ZINC dataset.

Dataset. ZINC is a free dataset consisting of 250k compounds and the task is to predict the constrained solubility of the compounds, formulated as a graph property regression problem. This problem is crucial for designing generative models for molecules. Each molecular graph in ZINC has the type of heavy atoms as node features (represented by a binary vector using one-hot encoding) and the type of bonds between atoms as edge features. In order to focus on the exploitation of the topological structures of the graphs, we omitted the edge features in our experiments. They could possibly be incorporated into GraphiT through the approach of [12] or considering different kernels on graph for each edge bond type, which is left for future work.

Table 5: Parameter grid for transformers trained on MUTAG, PROTEINS, PTC.

Parameter	Grid
Layers	[1, 2, 3]
Hidden dimension	[32, 64, 128]
Heads	[1, 4, 8]
Learning rate	[0.001]
Weight decay	[0.01, 0.001, 0.0001]

Table 6: Parameter grid for transformers trained on NCI1.

Parameter	Grid
Layers	[2, 3, 4]
Hidden dimension	[64, 128, 256]
Heads	[1, 4, 8]
Learning rate	[0.001]
Weight decay	[0.01, 0.001, 0.0001]

Training splits. Following [12], we use a subset of the ZINC dataset composed of respectively 10k, 1k and 1k graphs for train, validation and test split. The error bars are computed as the standard deviation of test accuracies across 4 different runs.

Hyperparameter choice. In order to make fair comparisons with the most relevant work [12], we use the same number of layers, number of heads and hidden dimensions, namely 10, 8 and 64. The number of model parameters for our transformers is only 2/3 of that of [12] as we use a symmetric variant for the attention score function in (5). Regarding the GNNs, we use the values reported in [12] for GIN, GAT and GCN. In addition, we use the same grid to train the MF model [11], *i.e.*, a learning rate starting at 0.001, two numbers of layers (4 and 16) and the hidden dimension equal to 145. Regarding the GCKN-subtree model, we use the same hidden dimension as GNNs, that is 145. We fix the bandwidth parameter to 0.5 and the path size is fixed to 10. We select the global pooling form max, mean, sum and weight decay from 0.001, 0.0001 and 1e-5, similar to the search grid used in [12].

Optimization. Following [12], we use the L1 loss and the Adam optimization method with batch size of 128 for training. Regarding the scheduling of the learning rate, we observe that a standard warm-up strategy used in [36] leads to more stable convergence for larger models (hidden dimension equal to 128). We therefore adopt this strategy throughout the experiments with a warm-up of 2000 iteration steps.

Additional results. Besides the relatively small models presented in Table 1, we also show in Table 7 the performance of larger models with 128 hidden dimensions. Increasing the number of hidden dimensions generally results in boosting performance, especially for transformer variants combined with Laplacian positional encoding in node features. While sparse local positional encoding is shown to be more useful compared to the positional encoding with a diffusion kernel, we show here that a variant of diffusion kernel positional encoding outperforms all other sparse local positional encoding schemes. Since the skip-connection in transformers already assigns some weight to each node itself and the diagonal of our kernels on graphs also have important weight, we considered setting these diagonal to zero in order to remove the effect of the attention scores on self-loop. This modification leads to considerable performance improvement on longer range relative positional encoding schemes, especially for the transformer with diffusion kernel positional encoding combined with GCKN in node features, which results in best performance.

B Additional visualization of Mutagenicity compounds

In this section, we use attention scores as a visualization mechanism for detecting substructures that may possibly induce mutagenicity. We apply our model to samples from the Mutagenicity data set [17] and analyze molecules that were correctly classified as mutagenic.

1,2-Dibromo-3-Chloropropane (DBCP). DBCP in Figure 3 was used as a soil fumigant in various countries. It was banned from use in the United States in 1979 after evidences that DBCP causes diseases, male sterility, or blindness, which can be instances of mutagenicity. In Figure 4, attention focuses on the

Table 7: Mean absolute error for regression problem ZINC. The results are computed from 4 different runs following [12].

Relative PE in attention score	Structural encoding in node features		
	None	LapPE	GCKN (p=8,d=32)
[12]	0.3585±0.0144	0.3233±0.0126	-
Transformers with d=64			
T	0.6964±0.0067	0.5065±0.0025	0.2741±0.0112
T + Adj PE	0.2432±0.0046	0.2020±0.0112	0.2106±0.0104
T + 2-step RW kernel	0.2429±0.0096	0.2272±0.0303	0.2133±0.0161
T + 3-step RW kernel	0.2435±0.0111	0.2099±0.0027	0.2028±0.0113
T + diffusion	0.2548±0.0102	0.2209±0.0038	0.2180±0.0055
Setting diagonal to zero, d=64			
T	0.7041±0.0044	0.5123±0.0232	0.2735±0.0046
T + 2-step	0.2427±0.0053	0.2108±0.0072	0.2176±0.0430
T + 3-step	0.2451±0.0043	0.2054±0.0072	0.1986±0.0091
T + diffusion	0.2468±0.0061	0.2027±0.0084	0.1967±0.0023
Larger models with d=128			
T	0.7044±0.0061	0.4965±0.0338	0.2776±0.0084
T + Adj PE	0.2310±0.0072	0.1911±0.0094	0.2055±0.0062
T + 2-step RW kernel	0.2759±0.0735	0.2005±0.0064	0.2136±0.0062
T + 3-step RW kernel	0.2501±0.0328	0.2044±0.0058	0.2128±0.0069
T + diffusion	0.2371±0.0040	0.2116±0.0103	0.2238±0.0068
Setting diagonal to zero, d=128			
T	0.7044±0.0061	0.4964±0.0340	0.2721±0.0099
T + 2-step	0.2348±0.0010	0.2012±0.0038	0.2031±0.0083
T + 3-step	0.2402±0.0056	0.2031±0.0076	0.2019±0.0084
T + diffusion	0.2351±0.0121	0.1985±0.0032	0.2019±0.0018




Figure 3: 1,2-Dibromo-3-Chloropropane.




Figure 4: Attention scores averaged by heads for each layer of our trained model for the compound in Figure 3. *Left*: diffusion kernel for 3. *Right*: node 3 and 5 (Br) are salient.

carbon skeleton and on the two Bromine (Br) atoms, the latter being indeed known for being associated with mutagenicity [23].

Nitrobenzene-nitroimidazothiazole. This compound is shown in Figure 5. As for compound 1a in Section 5.3, our model puts emphasis on two nitro groups which are indeed known for inducing mutagenicity.

Triethylenemelamine. Triethylenemelamine in Figure 7 is a compound exhibiting mutagenic properties, and is used to induce cancer in experimental animal models. Our model focuses on the three nitrogen atoms of the aziridine groups which are themselves mutagenic compounds.




Figure 5: Nitrobenzene-nitroimidazothiazole.




Figure 6: Attention scores averaged by heads for each layer of our trained model for the compound in Figure 5. *Left*: diffusion kernel for 5. *Right*: node 5 and 17 (N) are salient.




Figure 7: Triethylenemelamine.




Figure 8: Attention scores averaged by heads for each layer of our trained model for the compound in Figure 7. *Left*: diffusion kernel for 7. *Right*: node 1, 8, and 10 (N) are salient.