Revisiting Attention Weights as Interpretations of Message-Passing Neural Networks

Yong-Min Shin [1] Siqing Li [2] Xin Cao [2] Won-Yong Shin [1]

Abstract

The self-attention mechanism has been adopted in several widely-used message-passing neural networks (MPNNs) (*e.g.*, GATs), which adaptively controls the amount of information that flows along the edges of the underlying graph. This usage of attention has made such models a baseline for studies on explainable AI (XAI) since interpretations via attention have been popularized in various domains (*e.g.,* natural language processing and computer vision). However, existing studies often use naïve calculations to derive attribution scores from attention, and do not take the precise and careful calculation of edge attribution into consideration. In our study, we aim to fill the gap between the widespread usage of attention-enabled MPNNs and their potential in largely under-explored *explainability*, a topic that has been actively investigated in other areas. To this end, as the first attempt, we formalize the problem of *edge attribution* from attention weights in GNNs. Then, we propose GATT, an edge attribution calculation method built upon the *computation tree*. Through comprehensive experiments, we demonstrate the effectiveness of our proposed method when evaluating attributions from GATs. Conversely, we empirically validate that simply averaging attention weights over graph attention layers is insufficient to interpret the GAT model's behavior. Code is publicly available at https://github.com/jordan7186/GAtt/tree/main.

1. Introduction

1.1. Background

In graph learning, graph neural networks (GNNs) (Wu et al., 2021) have been used as the *de facto* architecture, since they can effectively encode the graph structure along with the node (or edge) features. Among various GNNs (Kipf & Welling, 2017; Hamilton et al., 2017; Xu et al., 2019), several models have successfully incorporated the self-attention mechanism (Bahdanau et al., 2015; Vaswani et al., 2017) into the message passing neural networks (MPNNs) (Gilmer et al., 2017; Bronstein et al., 2021), *e.g.*, GAT (Velickovic et al., 2018). Such architectures have effectively allowed to adaptively weigh the incoming information during message passing.

Despite their widespread usage, GNNs, similarly as in many other neural networks, are considered as black-box models that lack interpretability. To shed light on the inner workings of such GNNs, a large number of studies have focused on developing new post-hoc explainability methods, *i.e.,* methods that are applied after the underlying GNN model completed its training (Li et al., 2022; Yuan et al., 2023). However, such approaches typically necessitate an optimization framework, making the explanation performance dependent on various hyperparameters (e.g., number of iterations and random seeds). Such post-hoc explanations are also known to be sub-optimal since they essentially perform a single-step projection to an information-controlled space (Miao et al., 2022). This motivates us to explore an exciting opportunity for GNNs that produces attention weights (dubbed AttGNNs), as it would be beneficial to introduce a new method that directly utilizes the attention weights for explanation without any additional optimization/learning frameworks and hyperparameters. Furthermore, attention itself is also viewed as a direct way of providing model interpretations without a separate post-hoc explanation method (Lee et al., 2017; Ghaeini et al., 2018; Hao et al., 2021). This viewpoint has already been extensively performed on transformers, the representative model architecture employing attention (Bahdanau et al., 2015; Xu et al., 2015; Vig, 2019; Dosovitskiy et al., 2021; Caron et al., 2021). Also, there has been a significant body of research debating the effectiveness of attention as explanation, particularly within the domain of

[1] Yonsei University, South Korea [2] University of New South Wales, Austrailia. Correspondence to: Won-Yong Shin <wy.shin@yonsei.ac.kr>.

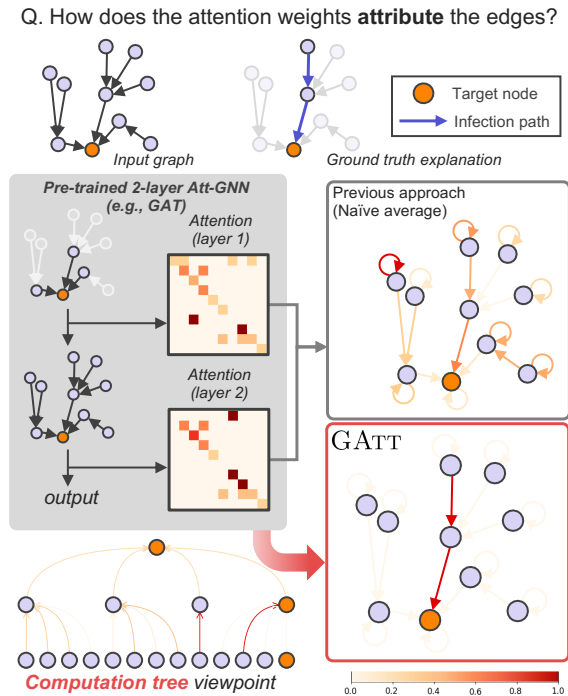Q. How does the attention weights **attribute** the edges?

Figure 1: A visualization of two different edge attribution methods given the attention weights in a GAT model, a representative GNN architecture with self-attention, trained on the Infection dataset. Contrast to simply averaging over graph attention layers, our proposed method, GATT, is built upon the computation tree, and reveals that the GAT model focuses mostly on the edges included in the ground truth explanation (blue edges in the input graph).

natural language processing (NLP) (Jain & Wallace, 2019; Wiegreffe & Pinter, 2019; Bibal et al., 2022).

1.2. Motivation

In contrast to other domains and architectures, the exploration of the attention weight's potential as explanation has been largely overlooked outside transformer architectures. While a few studies incorporated GATs as a baseline in their evaluations (Ying et al., 2019; Luo et al., 2020; Sánchez-Lengeling et al., 2020), they used attention as edge attribution (*i.e.,* highlighting relevant edges of the input graph) by simply averaging the attention weights over different layers, the precise computation of edge attribution from attention warrants a more intricate methodology.

Our study seeks to formulate *post-processing* for attention weights in Att-GNNs, enabling to extract higher-quality edge attributions that can better capture the essence of Att-GNNs.[1] While analogous investigations have been con-

---

[1]As mentioned, although graph transformers (Ying et al., 2021; Kreuzer et al., 2021; Chen et al., 2023) also heavily utilize attention, we focus on Att-GNNs in this study due to their widespread usage as well as their message passing-based architecture.

ducted in the context of transformers (Abnar & Zuidema, 2020; Chefer et al., 2021b;a), to the best of our knowledge, we are the first to address this issue within the domain of Att-GNNs (*i.e.,* MPNNs using attention).

1.3. Main Contributions

In this study, we tackle the intricate challenge of attributing edges using attention weights in Att-GNNs by dissecting the Att-GNNs feed-forward process, while taking into account the model's underlying *computation tree*. Drawing insights from this analysis, we assert that the edge attribution function to be designed should encompass two crucial dimensions: 1) proximity to the target node and 2) its position in the computation tree, thus aligning with the intrinsic feed-forward process.

To this end, we introduce GATT, which *adds* attention weights across the computation tree while *adjusting* their influence by employing targeted multiplication factors for attention weights guiding towards the target node. When using a GAT model as a representative GNN architecture, we run extensive experiments by answering pivotal facets of interpretation—*faithfulness* and *explanation accuracy*—of the underlying GAT across diverse real-world and synthetic datasets. Empirical results demonstrate that the integration of GATT to process attention weights yields substantively enhanced explanations, excelling in both faithfulness and explanation accuracy. As an example, Figure 1 visualizes attribution scores from different edge attribution methods using the same GAT model. The attribution scores from GATT (see red box) imply that the GAT model places high emphasis on the edges included in the ground truth explanation. Such conclusion could not have been reached if we were to use simple averaging as the tool for interpretation, validating the superiority of our method. Finally, we also perform an ablation study where we introduce two variants of GATT, namely GATT$_{SIM}$ and GATT$_{AVG}$, where we each remove one critical design element of our method. Our analysis shows that although the variants still outperforms naïve layer-wise averaging, it deteriorates the quality of the attribution in all measures. In summary, we conclude that **Att-GNNs such as GAT are indeed highly explainable** when adopting the proper interpretation, *i.e.,* **proper adjustment of attention weights** by **taking the computation tree viewpoint**. Note that, although the principles of GATT is generally applicable to all Att-GNNs, we focus on the original GAT model in our study due to its widespread usage and significance in the literature. Our contributions are summarized as follows:

- **Key observations:** We make key observations and design principles that need to be considered in the edge attribution calculation by integrating the computation tree of the target node during its calculation.

- **Novel methodology:** We propose GATT, a new method to calculate edge attribution from attention weights in a Att-GNNs by integrating the computation tree of the given GNN model.

- **Extensive evaluations:** We extensively demonstrate that the GAT model is shown to be more faithful and accurate when using our proposed method compared to the simple alternative.

2. Related Work

In this section, we discuss relevant prior studies in two major themes. Specifically, we first address previous attempts to use attention weights as interpretations in other domains, and then we provide an overview of explainable AI (XAI).

2.1. Interpreting Models with Attention

Attention weights have been used as a useful tool to visualize the inner workings of the underlying model (Bahdanau et al., 2015; Xu et al., 2015; Vig, 2019; Dosovitskiy et al., 2021; Caron et al., 2021), and there have been multiple studies that systematically questioned the utility of attention as interpretation in NLP (Bibal et al., 2022). In (Jain & Wallace, 2019), attention was found to have low correlations to other importance measures. However, (Wiegreffe & Pinter, 2019) pointed out an unfair setting in the prior work, and argued that attention can still be an effective explanation. Furthermore, there have been studies that focused on post-processing attention in transformers for token attribution. (Abnar & Zuidema, 2020) proposed attention rollout and attention flow, and follow-up attempts (Chefer et al., 2021a;b) were built upon attention rollout by incorporating other attribution methods or attention gradients. Although our study lies in a similar objective, our focus is on graph data with the Att-GNN architecture, which has been largely under-explored in the literature.

2.2. Explainability in GNNs

The primary goal of XAI is to provide a comprehensible understanding of neural network models' decisions. In recent years, various studies have developed methods to explain GNN models. As one of the pioneering work, GNNExplainer (Ying et al., 2019) identified a subset of edges and node features around the target node that affect the underlying model's decision. PGExplainer (Luo et al., 2020) trained a separate parameterized mask predictor to generate edge masks that identify important edges. Although explanations of GNN models are still an active research area (Li et al., 2022; Yuan et al., 2023), most studies overlooked Att-GNN as an inherently explainable model. Several attempts (Ying et al., 2019; Luo et al., 2020; Sánchez-Lengeling et al., 2020) introduced GATs, a representitative model in Att-GNNs, as

a baseline by averaging attention over layers. In light of this, our study aims to scrutinize attention as a paramount candidate for explaining Att-GNNs.

3. Edge Attribution Calculation in Att-GNNs

In this section, we first describe our notations used in the paper. Then, we formalize the problem of calculating the edge attribution in Att-GNNs, and then propose GATT, an approach for incorporating the computation tree to compute edge attributions.

3.1. Notations

Let us denote a given (undirected) graph as a tuple of two sets $G = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the set of nodes and $\mathcal{E}$ is the set of edges. We denote the edge connecting two nodes $v_i, v_j \in \mathcal{V}$ as $e_{ij} \in \mathcal{E}$. For the undirected G, $e_{j,i} \in \mathcal{E}$ if $e_{i,j} \in \mathcal{E}$ and vice versa. We also denote the set of neighbors of node v_i as $\mathcal{N}_i$.

3.2. Problem Statement

We are given a graph $G = (\mathcal{V}, \mathcal{E})$, the Att-GNN model f with L layers and a target node $v_i \in \mathcal{V}$ of interest. The attention weights calculated from f are denoted as $\mathcal{A} = \{\mathbf{A}(l)\}_{l=1}^L$, where $\mathbf{A}(l) \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ and $[\mathbf{A}(l)]_{j,i} = \alpha_{i,j}^l$ is the attention weight of edge $e_{i,j}$ in the l-th layer ($l = 1$ being the input layer). The problem of edge attribution calculation is characterized by an edge attribution function $\Phi(v, \mathcal{A}, e_{i,j}) \triangleq \phi_{i,j}^v$ such that the *edge attribution score* $\phi_{i,j}^v$ accounts for the importance of edge $e_{i,j}$ to the model's calculation for node v (*i.e.,* faithfulness of f).

In our study, we design Φ using the *computation tree* in Att-GNNs alongside several observations and key design principles, which will be specified later. To design such a function Φ, we argue that the **computation tree of Att-GNNs** should be considered, and incorporating its several key properties. Note that most post-hoc instance-level explanation methods for GNNs (Ying et al., 2019; Luo et al., 2020) also have a similar objective in terms of calculating $\phi_{i,j}^v$, but they do not leverage the attention weights $\mathcal{A}$. Additionally, we mainly consider GATs as the representative Att-GNN model in our study, but our methodology can be generally applied to other Att-GNN models. We will focus on the original GAT model (Velickovic et al., 2018) as a default from now on.

3.3. From Attention to Attribution

We first visualize the computation tree in a GAT, which will lead to several observations to guide GATT, an edge attribution calculation method given the attention weights in the GAT model.
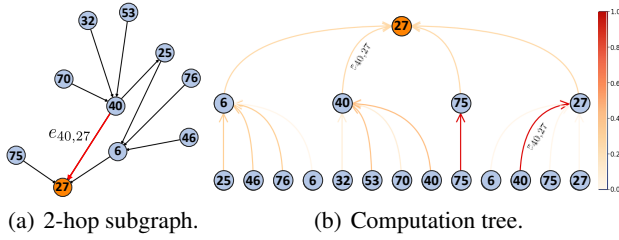
(a) 2-hop subgraph. (b) Computation tree.

Figure 2: A visualization for a 2-layer GAT on node 27 on the infection dataset. Figure 2(a) shows the local 2-hop subgraph with the edge $e_{40,27}$ marked as red. Figure 2(b) shows the computation tree in the GAT, where the information flows from leaf nodes to node 27 at the root. The edges are colored by the attention weights from the model, while highlighting the two occurrences of edge $e_{40,27}$.

3.3.1. VISUALIZING THE COMPUTATION TREE

To provide an illustrative example, we train a 2-layer GAT model (Velickovic et al., 2018) with a single attention head on the synthetic infection benchmark dataset (Faber et al., 2021). Figure 2(a) shows the 2-hop subgraph from target node 27, which contains all nodes and edges that the GAT model utilizes from node 27's point of view. The computation tree in the GAT is commonly expressed as a *rooted subtree* (Sato et al., 2021), which is shown in Figure 2(b) for node 27. In the figure, the information flows from leaf nodes at depth 2 to the root node 27 at depth 0, which exhibits a different structure from that of the subgraph in Figure 2(a). Since each graph attention layer calculates attention weights for each edge in $\mathcal{E}$, we color-code those weights in the computation tree (note that self-attention always includes self-loops regardless of the original graph structure).

3.3.2. OBSERVATIONS AND DESIGN PRINCIPLES

We make several observations from the computation tree:

(O1) Identical edges can appear multiple times in the computation tree. For example, edge $e_{40,27}$ in Figure 2(a) appears twice in Figure 2(b).

(O2) Nodes do not appear uniformly in the computation tree. Specifically, nodes that are k-hops away from the target node do not exist in depth k' for $0 < k' < k$ (*e.g.,* node 70 appears only at depth 2 while node 40 appears three times).

(O3) The graph attention layer always includes self-loops during its feed-forward process.

Based on the above observations, we would like to state two design principles that must be reflected when building the edge attribution function Φ.

(P1) Proximity effect: Edges within closer proximity to the target node tend to highly impact the model's prediction compared with faraway edges, since they **tend to appear more frequently** in the computation tree.

(P2) Contribution adjustment: The contribution of an edge in the computation tree should be **adjusted** by its position. (*i.e.*, depth in the tree)

By these standards, we look again at the attention weights in Figure 2(b). We first see that edges close to the target node such as $e_{40,27}$ appear twice, whereas distant edges such as $e_{70,40}$ appear only once (**P1**). Additionally, for the attention weights from the last graph attention layer (*i.e.,* edges connecting nodes at depth 1 to the root node), each edge tends to have roughly the value of 0.25 for the attention weights. In consequence, the information flowing from the first graph attention layer (*i.e.,* edges connecting leaf nodes to nodes at depth 1) will be diminished by 0.25 as it reaches the root node (**P2**).

3.3.3. PROPOSED METHOD

To design the edge attribution function Φ, we start by formally defining the computation tree alongside *flow* and *attention flow*.

Definition 3.1 (Computation tree). The computation tree for an L-layer GAT in our study is defined to as a rooted subtree of height L with the root (target) node. For each node in the tree at depth d, the neighboring nodes are at depth $d + 1$ with edges directed towards node v.

According to Definition 3.1, we define the concept of flows.

Definition 3.2 (Flow in a computation tree). Given a computation tree as a rooted subtree of height L with the root (target) node v, we define a flow $\lambda_{i,j,v}^l$ as the list of edges that sequentially appear in a path of length l starting from a given edge $e_{i,j}$ within the computation tree and ending with some edge $e_{*,v}$.[2] We indicate the k-th position within the flow as $\lambda_{i,j,v}^l(k)$ for $k \in [1, L]$. We denote the set of all flows in the computation tree with node v at its root that starts from edge $e_{i,j}$ with length $m \in [1, L]$ as $\Lambda_v^m(e_{i,j})$.

From Definition 3.2, it follows that $\lambda_{i,j,v}^l(1) = e_{i,j}$ and $\lambda_{i,j,v}^l(l) = e_{*,v}$ for all flows in $\Lambda_v(e_{i,j})$.

Definition 3.3 (Attention flow in a computation tree). Given a flow $\lambda_{v_0,v_1,w}^m = [e_{v_0,v_1}, \cdots, e_{*,w}]$ of length $m \leq L$ for an L-layer GAT model, we define an attention flow $\alpha[\lambda_{v_0,v_1,w}^m]$ as the corresponding attention weights assigned to each edge

---

[2]As stated in (**O1**), nodes/edges are not unique in the computation tree. Nonetheless, we will use the node indices from the original graph and avoid differentiating them in the computation tree as long as it does not cause any confusion.

by the associated graph attention layers:

$$\alpha[\lambda_{v_0,v_1,w}^m] = [\alpha_{v_0,v_1}^{L-m+1}, \cdots, \alpha_{*,w}^L]. \qquad (1)$$

Then, it follows that $\alpha[\lambda_{v_0,v_1,w}^m](i) = \alpha_{v_{i-1},v_i}^{L-m+i}$.

Example 1. In Figure 2(b), $\Lambda_{27}(e_{40,27})$ includes two flows, *i.e.,* $\lambda_{40,27,27}^1 = [e_{40,27}]$ and $\lambda_{40,27,27}^2 = [e_{40,27}, e_{27,27}]$, along with the corresponding attention flows $\alpha[\lambda_{40,27,27}^1] = [0.25]$ and $\alpha[\lambda_{40,27,27}^2] = [0.9, 0.25]$, respectively.

Finally, we are ready to present GATT.

Definition 3.4 (GATT). Given a target node v, an edge $e_{i,j}$ of interest, the set of flows, $\Lambda_v(e_{i,j})$, and the attention flows for all flows in $\Lambda_v(e_{i,j})$, we define the edge attribution of $e_{i,j}$ in L-layer GATs as

$$\phi_{i,j}^v = \sum_{m'=1}^L \sum_{\lambda_{i,j,v}^{m'} \in \Lambda_v^{m'}(e_{i,j})} C(\alpha[\lambda_{i,j,v}^{m'}])\alpha[\lambda_{i,j,v}^{m'}](1), \quad (2)$$

where $C(\alpha[\lambda_{i,j,v}^m]) = \prod_{2 \leq k \leq m} \alpha[\lambda_{i,j,v}^m](k)$ (or 1 if $m = 1$). Equation (2) can be interpreted as follows. We first find all occurrences of the target edge $e_{i,j}$ in the computation tree, and then re-weight its attention score (*i.e.,* $\alpha[\lambda_{i,j,v}^m](1)$) by the product of all attention weights that appear *after* $e_{i,j}$ (*i.e.,* $\alpha[\lambda_{i,j,v}^m](k)$ for $k \geq 2$) in the flow before the summation over all relevant flows. Next, let us turn to addressing how our design principles (**P1)** and (**P2)** are met. First, (**P1)** holds as we *add* the contributions from each flow rather than taking the average, as this will offset the total number of occurrences of $e_{i,j}$ in the computation tree. Next, (**P2)** is fulfilled by the adjustment factor from $C(\alpha[\lambda_{i,j,v}^m])$, since its value is dependent on the position of $\lambda_{i,j,v}^m(1)$. Essentially, $C(\alpha[\lambda_{i,j,v}^m])$ takes the chain of calculation from an edge to the target node into account. We provide an insightful example below.

Example 2. Let us recall $\lambda_{40,27,27}^2 = [e_{40,27}, e_{27,27}]$ and its attention flow $\alpha[\lambda_{40,27,27}^2] = [0.9, 0.25]$ on node 27 from Example 1. At face value, the contribution of edge $e_{40,27}$ within the flow $\lambda_{40,27,27}^2$ should be 0.9. However, this is inappropriate since the information will eventually get muted significantly by $\alpha_{27,27}^2 = 0.25$; thus, we need to consider $C(\alpha[\lambda_{40,27,27}^2])$ before calculating the final edge attribution. From Definition 3.4, the edge attribution $\phi_{40,27}^{27}$ from the attention weights is calculated as:

$$\phi_{40,27}^{27} = 1 \times 0.25 + 0.25 \times 0.9 = 0.475. \qquad (3)$$

3.4. Practical Calculation of GATT

Although GATT is defined by Equation (2), directly using this to compute the edge attribution $\phi_{i,j}^v$ is not desirable since it involves constructing the computation tree in the form of a rooted subtree for each node v, as well as computing over all relevant attention flows, resulting in high

redundancy during computation and not being proper for batch computation. To overcome these challenges, we introduce a *matrix-based computation* method that is much preferred in practice.

We first define

$$\mathbf{C}_L(k) = \begin{cases} \mathbf{I}, & \text{if } k = 0, \\ \mathbf{A}(L)\mathbf{A}(L-1)\cdots \mathbf{A}(L-k+1), & \text{otherwise.} \end{cases}$$

Then, we would like to establish the following proposition.

Proposition 3.5. *For a given set of attention weights $\mathcal{A} = \{\mathbf{A}(l)\}_{l=1}^L$ for an L-layer GAT with $L \geq 1$, GATT in Definition 3.4 is equivalent to*

$$\phi_{i,j}^v = \sum_{m=1}^L [\mathbf{C}_L(L-m)]_{v,j}[\mathbf{A}(m)]_{j,i}. \qquad (4)$$

Proposition 3.5 signifies that GATT sums the attention scores, weighted by the sum of the products of attention weights $[\mathbf{A}(m)]_{j,i}$ along the paths from node j to node v, over all graph attention layers.

3.5. Complexity Analysis

We first analyse the computational complexity of GATT. According to Equation (4), the bottleneck for calculating $\phi_{i,j}^v$ is to acquire $\prod_{k=m+1}^L \mathbf{A}(k)$. However, this matrix can be pre-computed and does not require re-calculation after its initial acquirement. Since we only count the number of multiplications in the summation, the computation complexity is $O(L)$, which is extremely efficient. Next, according to Equation (4), the memory complexity requires analyzing $\mathbf{C}_L(L-m)$ and $\mathbf{A}(m)$. For an L-layer GAT, while storing all attention weights in $\mathbf{A}(m)$ requires $O(L|\mathcal{E}|)$, $\mathbf{C}_L(L-m)$ requires at most $O(L||T^{L-1}||_0)$, where T denotes the adjacency matrix, and $||\cdot||_0$ is the 0-norm. In conclusion, the total memory complexity is $O(L||T^{L-1}||_0 + L|\mathcal{E}|)$.

4. Can Attention Interpret Att-GNNs?

In this section, we carry out empirical studies to validate the effectiveness of GATT, the proposed method interpreting a given GAT model.

4.1. Is Attention Faithful to the Model?

We focus primarily on one of the most important properties in evaluating an explanation method: *faithfulness*, which indicates how closely it truly reflects the underlying model's inner workings (Jacovi & Goldberg, 2020; Chrysostomou & Aletras, 2021; Liu et al., 2022; Li et al., 2022). Measuring the faithfulness involves 1) manipulating the input (such as masking a part of the input) according to the attribution suggested by the explanation method and 2) observing the

Table 1: Experimental results with respect to the faithfulness for GATT, AVGATT, and random attribution on four real-world datasets. Results for 2-layer and 3-layer GATs are shown for each case. The best performer is highlighted as **bold**.

Dataset	Measure	Metric	2-layer GAT			3-layer GAT		
			GATT	AVGATT	Random	GATT	AVGATT	Random
Cora	Δ_{PC}	$\rho_{Pearson}$ (↑)	**0.8468**	0.1764	-0.0056	**0.8642**	0.0967	0.0045
		$\tau_{Kendall}$ (↑)	**0.7051**	-0.1826	0.0082	**0.6512**	-0.0537	-0.0025
		$\rho_{Spearman}$ (↑)	**0.6516**	-0.1240	0.0061	**0.5679**	-0.0379	-0.0018
	Δ_{NE}	$\rho_{Pearson}$ (↑)	**0.7112**	0.1526	-0.0076	**0.7690**	0.0859	0.0040
		$\tau_{Kendall}$ (↑)	**0.7948**	-0.2463	0.0060	**0.7616**	-0.0820	0.0007
		$\rho_{Spearman}$ (↑)	**0.7371**	-0.1736	0.0044	**0.6737**	-0.0580	0.0005
	Δ_P	AUROC (↑)	**0.9755**	0.7251	0.4389	**0.9875**	0.7075	0.5235
Citeseer	Δ_{PC}	$\rho_{Pearson}$ (↑)	**0.8516**	0.3096	0.0012	**0.8711**	0.2110	-0.0073
		$\tau_{Kendall}$ (↑)	**0.7584**	-0.0106	0.0041	**0.6456**	-0.0130	0.0021
		$\rho_{Spearman}$ (↑)	**0.8321**	-0.0187	0.0057	**0.7318**	-0.0191	0.0031
	Δ_{NE}	$\rho_{Pearson}$ (↑)	**0.7653**	0.2780	0.0021	**0.8291**	0.2006	-0.0058
		$\tau_{Kendall}$ (↑)	**0.8469**	-0.0312	-0.0003	**0.7235**	-0.0263	0.0023
		$\rho_{Spearman}$ (↑)	**0.9206**	-0.0517	-0.0004	**0.8204**	-0.0376	0.0032
	Δ_P	AUROC (↑)	**0.9846**	0.9213	0.3695	**0.9920**	0.8979	0.4039
Pubmed	Δ_{PC}	$\rho_{Pearson}$ (↑)	**0.8812**	0.1648	-0.0064	**0.8489**	0.0592	0.0009
		$\tau_{Kendall}$ (↑)	**0.6268**	-0.0797	0.0002	**0.5349**	-0.0964	-0.0003
		$\rho_{Spearman}$ (↑)	**0.6746**	-0.1097	-0.0003	**0.5946**	-0.1348	-0.0004
	Δ_{NE}	$\rho_{Pearson}$ (↑)	**0.8201**	0.1477	-0.0068	**0.8612**	0.0600	0.0015
		$\tau_{Kendall}$ (↑)	**0.7031**	-0.0823	0.0025	**0.5378**	-0.1138	-0.0004
		$\rho_{Spearman}$ (↑)	**0.7568**	-0.1133	0.0033	**0.6187**	-0.1628	-0.0006
	Δ_P	AUROC (↑)	**0.9915**	0.8834	0.3974	**0.9993**	0.8932	0.5172
OGB-Arxiv	Δ_{PC}	$\rho_{Pearson}$ (↑)	**0.7790**	0.0794	0.0007	**0.7721**	0.0465	-0.0004
		$\tau_{Kendall}$ (↑)	**0.2047**	0.0128	0.0009	**0.1327**	-0.0158	-0.0041
		$\rho_{Spearman}$ (↑)	**0.2590**	0.0187	0.0013	**0.1778**	-0.0246	-0.0061
	Δ_{NE}	$\rho_{Pearson}$ (↑)	**0.8287**	0.0804	0.0016	**0.8282**	0.0478	-0.0017
		$\tau_{Kendall}$ (↑)	**0.2619**	0.0053	-0.0010	**0.1557**	-0.0086	-0.0038
		$\rho_{Spearman}$ (↑)	**0.3275**	-0.0066	-0.0015	**0.2106**	-0.0142	-0.0056
	Δ_P	AUROC (↑)	**0.9908**	0.8470	0.4962	**0.9985**	0.8331	0.5004

change of the model's response. We specify our experiment settings below.

4.1.1. ATTENTION ERASURE

To quantitatively analyze the importance of edge $e_{i,j}$, we mask the attention coefficients (i.e.,, attention weights before softmax) corresponding to edge $e_{i,j}$ with zeros in the computation tree, which is compared with the original response of the GAT model (Tomsett et al., 2020). Due to the fact that removing $e_{i,j}$ instead of erasing attention weights causes the whole attention distribution to change, we refrain from edge masking in our study. Moreover, we do not mask the attention weights after softmax, which never happens in a normal feed-forward procedure of GATs and makes the attention distribution not properly normalized. In other words, we only remove attention coefficients from one edge at a time to prevent the model from out-of-distribution computation, a common pitfall for perturbation-based approaches (Hooker et al., 2019; Hase et al., 2021).

4.1.2. DATASETS

In this experiment, we use four citation datasets, Cora, Citeseer, Pubmed (Yang et al., 2016), and a large-scale dataset, OGB-Arxiv (Hu et al., 2020), for node classification. They are real-world benchmark datasets where nodes correspond to papers and edges represent citations between papers. The nodes are labeled according to their topics.

4.1.3. BASELINE METHODS

We compare the proposed GATT against another attention-based explanation method (Ying et al., 2019; Luo et al., 2020; Sánchez-Lengeling et al., 2020), named as AVGATT, which attributes each edge as the average of the attention weights over different layers and attention heads. As another baseline, we additionally include random attribution, called 'Random', by randomly assigning scores between $[0, 1]$ to each edge.

4.1.4. MEASUREMENT OF GAT

Denoting the output probability vector of the GAT for node v as $\mathbf{p}_v$ and the output probability vector after the attention erasure for $e_{i,j}$ as $\mathbf{p}_{v \setminus e_{i,j}}$, we measure the model's behavior from three points of view: 1) *decline in prediction confidence* Δ_{PC} (Guo et al., 2017) defined as the decrease of the probability for the predicted label (*i.e.,* $\Delta_{\mathrm{PC}} = \mathbf{p}_v[i] - \mathbf{p}_{v \setminus e_{i,j}}[i]$, where $i = \arg\max_i \mathbf{p}_v[i]$), 2) *change in negative entropy* Δ_{NE} (Moon et al., 2020) defined as the increase of 'smoothness' of the probability vector (*i.e.,* $\Delta_{\mathrm{NE}} = -\sum \mathbf{p}_{v \setminus e_{i,j}} \log \mathbf{p}_{v \setminus e_{i,j}} + \sum \mathbf{p}_v \log \mathbf{p}_v$), which also reflects the model's confidence, and 3) *change in prediction* Δ_{P}, which measures the difference between $\arg\max_i \mathbf{p}_v[i]$ and $\arg\max_i \mathbf{p}_{v \setminus e_{i,j}}[i]$, where $\mathbf{p}_v[i]$ is the i-th entry of $\mathbf{p}_v$.

4.1.5. QUANTITATIVE ANALYSIS WITH ERASURE

We investigate the relationship between the attribution scores acquired from attention erasure and edge attribution calculation methods. In each dataset, we randomly select 100 nodes as target nodes v and calculate GATT values for all edges (i, j) that affect the target node (*i.e.,* $\phi_{i,j}^v$). We also perform attention erasure for the same edges (i, j) and measure Δ_{PC}, Δ_{NE}, and Δ_{P} to observe the correlation between GATT values. Specifically, for Δ_{PC} and Δ_{NE}, we adopt three metrics: Pearson's rho ρ_{Pearson}, Kendall's tau τ_{Kendall}, and Spearman's rho ρ_{Spearman}. For Δ_{P}, we use the area under receiver operating characteristic (AUROC), basically assessing the quality of an attribution score as a predictor of whether the prediction of the target node will change after attention erasure.

Table 1 summarizes the experimental results with respect to the faithfulness on the four real-world citation datasets, using 2-layer and 3-layer trained GATs with a single attention head for each dataset. The results strongly indicate using GATT reflects the behavior of the GAT model much better than AVGATT, resulting in a drastic increase in explainability performance compared to AVGATT and random attribution. Although AVGATT shows modest performance in Δ_{P}, it performs poorly in terms of changes in confidence (*i.e.,* Δ_{PC} and Δ_{NE}), sometimes performing worse than random attribution. This is because AVGATT does not account for the proximity effect and contribution adjustment, and rather naïvely averages the attention weights over different layers with no context of the computation tree.

Additionally, we consider the effect of multi-head attention, which allows each attention head to potentially capture different patterns in GATs. To this end, we carry out an experiment while increasing the number of attention heads from 1 to 8. Table 2 summarizes the performance of edge attribution calculation methods according to different attention head configurations. We find that the trend in performance is largely consistent with our findings in Table 1 even with

Table 2: Experimental results with respect to the faithfulness by increasing the number of attention heads from 1 to 8 for 2 and 3-layer GAT models. The table shows the case for the Cora dataset, measuring Δ_{PC} with ρ_{Spearman}.

Model	Method	Number of heads			
		1	2	4	8
2-layer	GATT	0.8477	0.8625	0.8468	0.8496
	AVGATT	0.1768	0.1807	0.1697	0.1728
	Random	-0.0079	0.0031	0.0021	0.0073
3-layer	GATT	0.8624	0.8857	0.8674	0.7048
	AVGATT	0.0966	0.0965	0.0994	0.0857
	Random	0.0092	-0.0011	0.0001	-0.0033

an increased number of attention heads. However, we find no distinguishable pattern with respect to the number of heads in both 2-layer and 3-layer GATs for the Cora dataset; nonetheless, this is not necessarily the case for other two datasets.

4.2. Does Attention Reveal Accurate Explanations?

We evaluate the edge attributions of the GAT model in comparison with ground truth explanations. Since only the synthetic datasets are equipped with proper ground truth explanations, we only use these datasets during evaluation.

4.2.1. DATASETS

We use the BA-shapes and Infection synthetic benchmark datasets, which also includes ground truth explanations. *BA-shapes* (Ying et al., 2019) attaches 80 house-shaped motifs to a base graph made from the Barabási-Albert model with 300 nodes, where the edges included in the motif is set as the ground-truth explanation. *Infection benchmark* (Faber et al., 2021) generates a backbone graph from the Erdös-Rényi model. Then, a small portion of the nodes are assigned as 'infected', and the ground-truth explanation is the path from an infected node to the target node. We expect that edge attributions should highlight such ground truth explanations for the underlying GAT with high performance.

4.2.2. BASELINE METHODS

In this experiment, we mainly compare the performance among attention-based edge attribution calculation methods (*i.e.,* GATT and AVGATT) including random attribution. Additionally, we consider seven popular *post-hoc* explanation methods: Saliency (Simonyan et al., 2014), Guided Backpropagation (Springenberg et al., 2015), Integrated Gradient (Sundararajan et al., 2017), GNNExplainer (Ying et al., 2019), PGExplainer (Luo et al., 2020), GraphMask (Schlichtkrull et al., 2021), and FastDnX (Pereira et al., 2023). We emphasize that post-hoc explanation methods are treated as a complementary category of inherent

Table 3: Experimental results on the explanation accuracy for the synthetic datasets using 3-layer GATs, measured in terms of the AUROC. The results for directly using attention weights as explanation are colored in gray. The best and runner up performances are marked as **bold** and <u>underline</u>, respectively.

Expl. type	Method	BA-Shapes	Infection
Attention	GATT	<u>0.9591</u>	**0.9976**
	AVGATT	0.7977	0.8786
-	Random	0.4975	0.4811
Post-hoc	Saliency	0.9563	0.8237
	Guided Backprop	0.6231	0.8949
	Integrated Gradient	0.6231	<u>0.9472</u>
	GNNExplainer	0.8912	0.9272
	PGExplainer	0.8289	0.7173
	GraphMask	0.5316	0.6859
	FastDnX	**0.9917**	0.6574

explanations, thus belonging to a **different category** (Du et al., 2020). However, we include them for a more comprehensive comparison.

4.2.3. MEASUREMENT

As in prior studies (Ying et al., 2019; Luo et al., 2020), we treat evaluation as binary classification of edges, aiming to predict whether each edge belongs to ground truth explanations by using the attribution scores as probability values. In this context, we adopt the AUROC as our metric.

4.2.4. EXPERIMENTAL RESULTS

Table 3 summarizes the results on the explanation accuracy for two synthetic datasets with ground truth explanations. For both datasets, we observe that GATT is superior to AVGATT. Even compared to the representative post-hoc explanation methods, GATT performs surprisingly better. This indicates that the attention weights can inherently capture the GAT model's behavior without external post-hoc explanation methods.

4.3. Ablation Study

The definition of GATT by Definition 3.4 satisfies the two design principles (*i.e.*, proximity effect (**P1**) and contribution adjustment (**P2**)). We now perform an ablation study to validate the effectiveness of the contribution adjustment principle. To this end, we devise two variants GATT$_{\text{SIM}}$ and GATT$_{\text{AVG}}$, simply adding all attention weights uniformly corresponding to the target edge in the computation tree and replacing the weighted summation in Equation (2) with averaging to remove the effects of the proximity effect, respectively. Specifically, GATT$_{\text{SIM}}$ and GATT$_{\text{AVG}}$ are defined

Table 4: Properties of different edge attribution calculation methods.

Method	GATT	GATT$_{\text{SIM}}$	GATT$_{\text{AVG}}$	AVGATT
(**P1**) Proximity effect	✓	✓	✗	✗
(**P2**) Contribution adjustment	✓	✗	✓	✗

Table 5: Properties of different edge attribution calculation methods.

Dataset	Model	GATT	GATT$_{\text{SIM}}$	GATT$_{\text{AVG}}$	AVGATT
Cora	2-layer	**0.7793**	0.7565	0.7650	0.0461
	3-layer	**0.7598**	0.6603	0.6940	0.0954
Citeseer	2-layer	**0.8514**	0.8295	0.7317	0.1995
	3-layer	**0.8019**	0.7210	0.7890	0.1734
Pubmed	2-layer	**0.7792**	0.7413	0.6796	0.1158
	3-layer	**0.7136**	0.5727	0.7051	0.0721

as,

$$\phi_{i,j}^v = \sum_{m'=1}^{L} \sum_{\lambda_{i,j,v}^{m'} \in \Lambda_v^{m'}(e_{i,j})} \alpha[\lambda_{i,j,v}^{m'}](1), \tag{5}$$

$$\phi_{i,j}^v = \frac{1}{|\Lambda_v(e_{i,j})|} \sum_{m'=1}^{L} \sum_{\lambda_{i,j,v}^{m'} \in \Lambda_v^{m'}(e_{i,j})} C(\alpha[\lambda_{i,j,v}^{m'}])\alpha[\lambda_{i,j,v}^{m'}](1), \tag{6}$$

respectively. The properties of different edge attribution calculation methods are summarized in Table 4.

Evaluation of faithfulness. We compare the performance among GATT, GATT$_{\text{SIM}}$, GATT$_{\text{AVG}}$, and AVGATT by running experiments with respect to the faithfulness on the Cora, Citeseer, and Pubmed datasets. Table 5 summarizes the performance of each method by averaging over all 7 measures and metrics (*i.e.*, 3 correlation metrics for Δ_{PC} and Δ_{NE} each and AUROC for Δ_{p}). Comparing GATT and its variants, we observe that GATT consistently outperforms both GATT$_{\text{SIM}}$ and GATT$_{\text{AVG}}$ for all cases. In particular, the performance degradation of the variants is generally more severe for 3-layer GATs. This is because the effect of the contribution adjustment (**P2**) and the cardinality of $\Lambda_v(e_{i,j})$ are more significant in 3-layer GATs since the length of each attention flow are longer as well as the number of flows to consider is much higher compared to 2-layer GATs.

Case study. We conduct case studies on the BA-Shapes and Infection datasets for 2-layer GATs. In Figure 3, each of two cases shows the target node (orange node) and the ground truth explanations (blue edges). We aim to observe how much the attribution scores from different methods focus on the ground truth explanation edges. Indeed, for both datasets, GATT focuses primarily on the ground truth explanations, while the attribution scores from GATT$_{\text{SIM}}$ and AVGATT tend to be more spread throughout the whole graph.

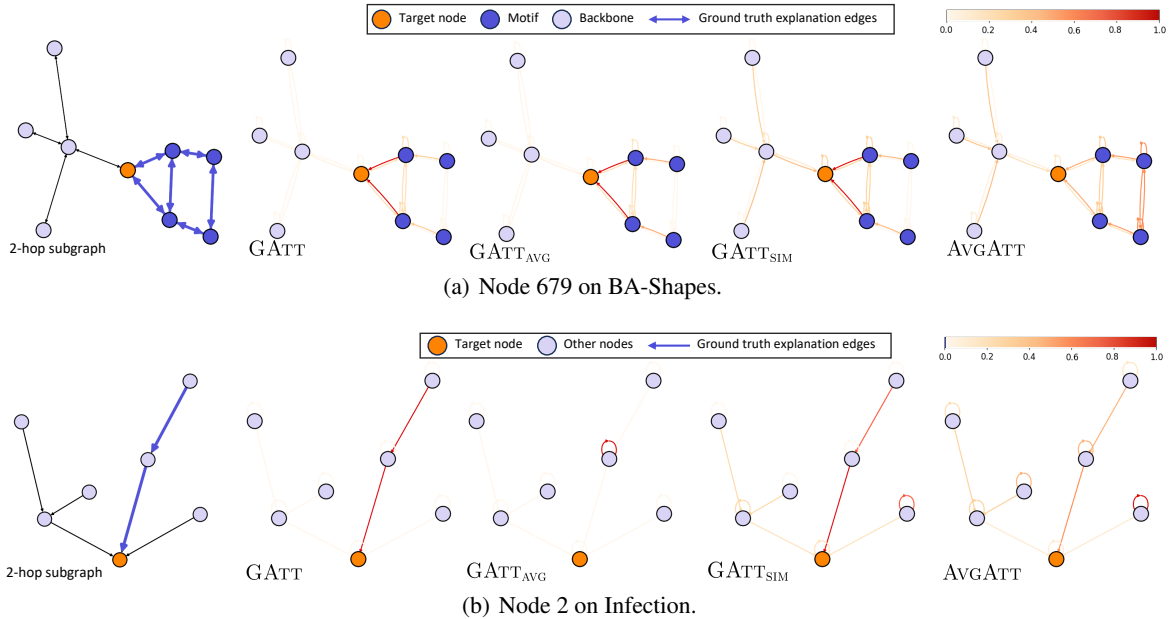(a) Node 679 on BA-Shapes.

(b) Node 2 on Infection.

Figure 3: Case study on the BA-Shapes and Infection datasets for a 2-layer GAT.

This indicates that the attention weights in GATs indeed recognize the ground truth explanations. For GATT$_{\text{AVG}}$, it is likely that the attributions are not much different from GATT in BA-Shapes. However, GATT$_{\text{AVG}}$ in the Infection dataset attributes its attribution scores to a single self-loop edge that does not belong to the ground truth explanations, failing to provide adequate explanations. Interestingly, on BA-Shapes, GATT tends to strongly emphasize local edges even within the house-shaped motifs, which coincides with the pitfall addressed in (Faber et al., 2021).

5. Conclusion and Future Work

In this study, we have investigated the largely under-explored problem of interpreting Att-GNNs by utilizing the attention weights. Although Att-GNNs were not regarded as a candidate as an inherently explainable model, our empirical evaluations have demonstrated the affirmative results when our proposed method, GATT, is used to calculate edge attribution scores. Even though GATT is generally applicable, this work does not include a systematic analysis on how different designs of attention weights will interact with GATT, which we leave as future work.

References

Abnar, S. and Zuidema, W. H. Quantifying attention flow in transformers. In *ACL*, pp. 4190–4197, Online, Jul. 2020.

Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. In *ICLR*, San Diego, CA, May 2015.

Bibal, A., Cardon, R., Alfter, D., Wilkens, R., Wang, X.,

François, T., and Watrin, P. Is attention explanation? an introduction to the debate. In *ACL*, pp. 3889–3900, Dublin, Ireland, May 2022.

Bronstein, M. M., Bruna, J., Cohen, T., and Velickovic, P. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges, 2021.

Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., and Joulin, A. Emerging properties in self-supervised vision transformers. In *ICCV*, pp. 9630–9640, Montreal, Canada, Oct. 2021.

Chefer, H., Gur, S., and Wolf, L. Transformer interpretability beyond attention visualization. In *CVPR*, pp. 782–791, virtual, Jun. 2021a.

Chefer, H., Gur, S., and Wolf, L. Generic attention-model explainability for interpreting bi-modal and encoder-decoder transformers. In *ICCV*, pp. 387–396, Montreal, Canada, Oct. 2021b.

Chen, J., Gao, K., Li, G., and He, K. NAGphormer: A tokenized graph transformer for node classification in large graphs. In *ICLR*, May 2023.

Chrysostomou, G. and Aletras, N. Improving the faithfulness of attention-based explanations with task-specific information for text classification. In *ACL/IJCNLP*, pp. 477–488, Virtual Event, Aug. 2021.

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, Virtual event, May 2021.

Du, M., Liu, N., and Hu, X. Techniques for interpretable machine learning. *Commun. ACM*, 63(1):68–77, 2020.

Faber, L., Moghaddam, A. K., and Wattenhofer, R. When comparing to ground truth is wrong: On evaluating GNN explanation methods. In *KDD*, pp. 332–341, Virtual Event, Aug. 2021.

Ghaeini, R., Fern, X. Z., and Tadepalli, P. Interpreting recurrent and attention-based neural models: a case study on natural language inference. In *EMNLP*, pp. 4952–4957, Brussels, Belgium, Nov. 2018.

Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. In *ICML*, pp. 1263–1272, Sydney, NSW, Aug. 2017.

Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. In *ICML*, pp. 1321–1330, Sydney, Australia, Aug. 2017.

Hamilton, W. L., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. In *NeurIPS*, pp. 1024–1034, Long Beach, CA, Dec. 2017.

Hao, Y., Dong, L., Wei, F., and Xu, K. Self-attention attribution: Interpreting information interactions inside transformer. In *AAAI*, pp. 12963–12971, Virtual event, Feb. 2021.

Hase, P., Xie, H., and Bansal, M. The out-of-distribution problem in explainability and search methods for feature importance explanations. In *NeurIPS*, pp. 3650–3666, Virtual event, Dec. 2021.

Hooker, S., Erhan, D., Kindermans, P., and Kim, B. A benchmark for interpretability methods in deep neural networks. In *NeurIPS*, pp. 9734–9745, Vancouver, Canada, Dec. 2019.

Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. Open graph benchmark: Datasets for machine learning on graphs. In *NeurIPS*, Virtual event, Dec. 2020.

Jacovi, A. and Goldberg, Y. Towards faithfully interpretable NLP systems: How should we define and evaluate faithfulness? In *ACL*, pp. 4198–4205, Online, Jul. 2020.

Jain, S. and Wallace, B. C. Attention is not explanation. In *NAACL-HLT*, pp. 3543–3556, Minneapolis, MN, Jun. 2019.

Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *ICLR*, Toulon, France, Apr. 2017.

Kreuzer, D., Beaini, D., Hamilton, W. L., Létourneau, V., and Tossou, P. Rethinking graph transformers with spectral attention. In *NeurIPS*, pp. 21618–21629, Virtual event, Dec. 2021.

Lee, J., Shin, J., and Kim, J. Interactive visualization and manipulation of attention-based neural machine translation. In *EMNLP*, pp. 121–126, Copenhagen, Denmark, Sept. 2017.

Li, P., Yang, Y., Pagnucco, M., and Song, Y. Explainability in graph neural networks: An experimental survey, 2022.

Liu, Y., Li, H., Guo, Y., Kong, C., Li, J., and Wang, S. Rethinking attention-model explainability through faithfulness violation test. In *ICML*, pp. 13807–13824, Baltimore, MD, Jul. 2022.

Luo, D., Cheng, W., Xu, D., Yu, W., Zong, B., Chen, H., and Zhang, X. Parameterized explainer for graph neural network. In *NeurIPS*, Virtual event, Dec. 2020.

Miao, S., Liu, M., and Li, P. Interpretable and generalizable graph learning via stochastic attention mechanism. In *ICML*, Baltimore, MD, Jul. 2022.

Moon, J., Kim, J., Shin, Y., and Hwang, S. Confidence-aware learning for deep neural networks. In *ICML*, pp. 7034–7044, Vienna, Austria, Jul. 2020.

Pereira, T. A., Nascimento, E., Resck, L. E., Mesquita, D., and Souza, A. H. Distill n' explain: explaining graph neural networks using simple surrogates. In *AISTATS*, Palau de Congressos, Spain, Apr. 2023.

Sánchez-Lengeling, B., Wei, J. N., Lee, B. K., Reif, E., Wang, P., Qian, W. W., McCloskey, K., Colwell, L. J., and Wiltschko, A. B. Evaluating attribution for graph neural networks. In *NeurIPS*, virtual, Dec. 2020.

Sato, R., Yamada, M., and Kashima, H. Random features strengthen graph neural networks. In *SDM*, pp. 333–341, Virtual event, Apr.–May 2021.

Schlichtkrull, M. S., Cao, N. D., and Titov, I. Interpreting graph neural networks for NLP with differentiable edge masking. In *ICLR*, Virtual Event, May 2021.

Simonyan, K., Vedaldi, A., and Zisserman, A. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *ICLR*, Banff, Canada, Apr. 2014.

Springenberg, J. T., Dosovitskiy, A., Brox, T., and Riedmiller, M. A. Striving for simplicity: The all convolutional net. In *ICLR*, San Diego, CA, May 2015.

Sundararajan, M., Taly, A., and Yan, Q. Axiomatic attribution for deep networks. In *ICML*, Sydney, Australia, Aug. 2017.

Tomsett, R., Harborne, D., Chakraborty, S., Gurram, P., and Preece2, A. Sanity checks for saliency metrics. In *AAAI*, New York, NY, Feb. 2020.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. In *NIPS*, pp. 5998–6008, Long Beach, CA, Dec. 2017.

Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. Graph attention networks. In *ICLR*, Vancouver, Canada, Apr. 2018.

Vig, J. A multiscale visualization of attention in the transformer model. In *ACL*, pp. 37–42, Florence, Italy, Jul.–Aug. 2019.

Wiegreffe, S. and Pinter, Y. Attention is not not explanation. In *EMNLP-IJCNLP*, pp. 11–20, Hong Kong, China, Nov. 2019.

Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., and Yu, P. S. A comprehensive survey on graph neural networks. *IEEE Trans. Neural Networks Learn. Syst.*, 32(1):4–24, 2021.

Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A. C., Salakhutdinov, R., Zemel, R. S., and Bengio, Y. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, pp. 2048–2057, Lille, France, Jul. 2015.

Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? In *ICLR*, New Orleans, LA, May 2019.

Yang, Z., Cohen, W. W., and Salakhutdinov, R. Revisiting semi-supervised learning with graph embeddings. In *ICML*, New York City, NY, Jun. 2016.

Ying, C., Cai, T., Luo, S., Zheng, S., Ke, G., He, D., Shen, Y., and Liu, T. Do transformers really perform badly for graph representation? In *NeurIPS*, pp. 28877–28888, virtual, Dec. 2021.

Ying, Z., Bourgeois, D., You, J., Zitnik, M., and Leskovec, J. GNNExplainer: Generating explanations for graph neural networks. In *NeurIPS*, pp. 9240–9251, Vancouver, Canada, Dec. 2019.

Yuan, H., Yu, H., Gui, S., and Ji, S. Explainability in graph neural networks: A taxonomic survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(5):5782–5799, 2023.