Learning on heterogeneous graphs using high-order relations

Lee, See Hian; Ji, Feng; Tay, Wee Peng

2021

Lee, S. H., Ji, F. & Tay, W. P. (2021). Learning on heterogeneous graphs using high-order relations. 2021 International Conference on Acoustics, Speech and Signal Processing (ICASSP).

<https://hdl.handle.net/10356/147362>

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Downloaded on 26 Apr 2023 20:34:41 SGT

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

LEARNING ON HETEROGENEOUS GRAPHS USING HIGH-ORDER RELATIONS

See Hian Lee Feng Ji Wee Peng Tay

School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore

ABSTRACT

A heterogeneous graph consists of different vertices and edges types. Learning on heterogeneous graphs typically employs meta-paths to deal with the heterogeneity by reducing the graph to a homogeneous network, guide random walks or capture semantics. These methods are however sensitive to the choice of meta-paths, with suboptimal paths leading to poor performance. In this paper, we propose an approach for learning on heterogeneous graphs without using meta-paths. Specifically, we decompose a heterogeneous graph into different homogeneous relation-type graphs, which are then combined to create higher-order relation-type representations. These representations preserve the heterogeneity of edges and retain their edge directions while capturing the interaction of different vertex types multiple hops apart. This is then complemented with attention mechanisms to distinguish the importance of the relation-type based neighbors and the relation-types themselves. Experiments demonstrate that our model generally outperforms other state-of-the-art baselines in the vertex classification task on three commonly studied heterogeneous graph datasets.

Index Terms— Graph neural network, heterogeneous graph, high-order relations

1. INTRODUCTION

Learning representations that encode graph structural information is increasingly complex as we seek to model real-world graph data. In many applications, graph data are heterogeneous in nature. These heterogeneous graphs are innately ingrained with rich semantics and are more complex than their homogeneous counterparts [1, 2]. Thus, existing approaches [3, 4] developed for homogeneous graphs have to be adapted or re-formularized to cater to the disparity and to capture the semantics present in a heterogeneous graph.

A typical approach is to use meta-paths [5–7]. A meta-path [8] is a predefined sequence of vertex types. It is utilized to transform heterogeneous graphs into homogeneous graphs [5, 7], to guide random walks [9] or to capture the rich semantics [8]. This can be seen in techniques such as Heterogeneous graph Attention Network (HAN) [10], metapath2vec [9] and MEIRec [11]. Meta-paths are widely utilized as they express different semantics, which help to mine the

semantics information present [10].

For instance, in a bibliographic graph, a meta-path “APVPA”, which refers to the vertex sequence “author-paper-venue-paper-author”, conveys the idea that the authors published papers in the same venue. Meanwhile, “APA” indicates co-authorship relationship between two authors. However, selecting the optimal meta-paths for a specific task or dataset is a non-trivial task. It requires specialized domain knowledge regarding the relationships between different vertex types and edge types in a heterogeneous graph as well as the learning task at hand [7, 12]. Constructing meta-paths based on domain knowledge becomes inefficient for graphs with a large number of vertex and edge types as the number of possible meta-paths grows exponentially with the number of vertex and edge types [2].

In this paper, we propose an approach to learn on heterogeneous graphs that does not use meta-paths. Instead, it leverages on the heterogeneity and directions of edges to build high-order relations that span multiple hops. An edge type in the graph is a first order relation-type. The heterogeneous graph is first decomposed into subgraphs with homogeneous first order relation-types. The adjacency matrices of these subgraphs are then multiplied together to form (weighted) adjacency matrices that represent higher-order relation-types amongst the vertices. Learning, together with attention mechanisms, is then performed on the set of higher-order relation-types. We call our approach **leaRning hEterogeneous GrAphs wiTh High-order Relations** (REGATHER). The acronym REGATHER invokes the idea of decomposing a difficult problem into basic units that can be combined to facilitate learning.

Taking DBLP dataset as example, the first order relation-type subgraphs are shown in Fig. 1c)-j). We consider the schema edge directions and their reverse directions as independent first order relation-types. In many existing studies [3, 4, 13], it is common to consider directed graphs as undirected [14, 15]. However, this undesirably discards information captured in edge directions. By treating the reversed edge direction as an independent edge-type, we can learn different weights for an edge’s forward and backward directions.

2. RELATED WORK

Learning on heterogeneous graphs with meta-paths have shown good results as seen in [9–11]. Yet, given the limita-

Fig. 1. Structure of DBLP graph: a) Schema, b) An illustration of DBLP graph, c)-f) First order relation-type subgraphs, g)-j) Reversed first order subgraphs.

tions of manually designing meta-paths, methods to automate the discovery of relevant meta-paths have been developed such as using greedy strategies [16], enumerating meta-paths within a fixed length [17], using reinforcement learning [18] and stacking multiple graph transformer layers [7].

Learning on heterogeneous graphs based upon graph neural network (GNN) approaches that do not utilize meta-paths include Heterogeneous Graph Transformer (HGT) [12] and Heterogeneous Graph Neural Network (HetGNN) [1]. These methods differ from our approach in their ways to incorporate information from multi-hop neighbors. HGT uses cross-layer information transfer to obtain information of neighbors that are multi-hop apart and HetGNN applies random walk with restart (RWR) strategy to sample all types of neighbor nodes from each node which implicitly learns higher-order relations. Moreover, HetGNN focuses on integrating heterogeneous content information of each node which differs from REGATHER that learns without incorporating content information.

Performing random walks on heterogeneous graphs without the use of meta-paths for guidance is also seen in [19]. JUMP and STAY (JUST) [19] introduced a way to perform random walk on heterogeneous graphs without using meta-paths by using its jump and stay strategy to probabilistically balance *jumping* to a different vertex type and *staying* in the same vertex type during random walk. The whole process of JUST projects the vertices from different domains into a common vector space, generating vertex embeddings that are simply based upon structural information.

The attention mechanism is also evident in its usage in learning of graphs as seen in [4, 7, 10]. Heterogeneous graph Attention Network (HAN) proposed by [10] is a GNN designed for heterogeneous graphs employing hierarchical attention. Nevertheless, it uses a set of predefined meta-paths to transform a heterogeneous graph into a set of meta-path based homogeneous graphs. REGATHER employs dual-level attention like HAN but differs from HAN in three main ways.

Firstly, instead of transforming a heterogeneous graph into meta-path based graphs, it transforms the heterogeneous graph into a set of relation-type subgraphs. Secondly, the heterogeneity of edges and their directions are retained, viewing the reverse edge direction as an independent first order relation-type. This is because even if two vertices are only connected via one direction semantically, they are related in the reverse direction from an information theoretic viewpoint. Lastly, in order to encode information from higher-order relations, which are again typically performed using customized meta-paths in the literature [10], REGATHER uses higher-order relations, which are discussed in Section 3.1.

3. METHODOLOGY

In this section, we introduce the notion of relation-type subgraphs and present REGATHER’s architecture. Throughout, we consider a heterogeneous directed graph $G = (V, E)$ where V is the vertex set consisting of multiple vertices types, and E is the set of edges consisting of multiple directed edge types.

3.1. Relation-type subgraphs

Each edge-type is defined to be a first order relation. We decompose the graph G into a set of subgraphs $H_1, \dots, H_c$, where c is the total number of edge-types and each subgraph H_i has $|V|$ vertices and only one edge-type such that

$$\bigcup_{1 \leq i \leq c} H_i = G.$$

For each H_i , there is an associated adjacency or relation matrix M_i . As the graph G is directed, each matrix M_i may not be symmetric. To consider edges with reversed directions, we take the transpose of each of the relation matrices of $H_1, \dots, H_c$. Let S be the set of matrices $\{M_i : 1 \leq i \leq c\}$ and S^T be the set of transposed matrices $\{M_i^T : 1 \leq i \leq c\}$. We take their union as the first order relation matrices on G :

$$S^1 = S \cup S^T.$$

Each first order relation-type captures only partial information contained in the heterogeneous graph. For example, consider the DBLP graph in Fig. 1. There are no edges of the form “author-author”. The fact that two authors may have co-authored a paper together, or attended the same conference previously, may contribute to our learning task. To capture such relationships, we consider higher-order relations by combining, e.g., “author-paper” and “paper-author” relation-type graphs. This is achieved by multiplication of the respective subgraph relation matrices together to obtain a (weighted) relation matrix. Nevertheless, it is possible to have some trivial resulting matrices given that some consecutive relation-types are unrealizable traversals. The trivial matrices are removed from the set. The k -th order relation-type is defined as

$$S^k = \{N_1 N_2 \dots N_k : N_i \in S^1, 1 \leq i \leq k\}.$$

We see that each meta-path of length k corresponds to an edge in one of the relation matrices in S^k . Our goal is to perform learning over

$$D = \bigcup_{k=1}^K S^k,$$

for some $K \geq 1$. As an additional step, we add self-loops to each of the matrices in D by adding an identity matrix I to obtain the set of relation matrices

$$\Phi = \{A + I : A \in D\}. \quad (1)$$

Using high-order relation-type subgraphs, we eliminate the need to predefine meta-paths. Instead, these are exhaustively included in Φ . In the following, we show how to use attention mechanisms to capture the importance of each relation-type in a data-driven fashion.

3.2. Attention for relation-type based neighbors

The relation matrices in Φ reflect the adjacency between vertices via the specific relation. Two vertices are considered adjacent if the entry in the matrix corresponding to the vertices is non-zero. The neighborhood of a vertex is the set of vertices adjacent to it. The different neighborhoods reflect the different aspects of a vertex. It is crucial to determine the importance of each of the adjacent vertices in a specific neighborhood because different adjacent vertices in a neighborhood can contribute differently to a vertex's representation. Hence, a single headed attention is applied to each of the relation matrices to determine the importance of a vertex's neighboring features specific to each relation matrix in Φ .

By using single headed graph attention, the number of parameters that have to be learned is reduced significantly. It is observed that reducing the attention heads to one did not negatively impact the performance of REGATHER. In some cases, it even improves REGATHER's performances. This is aligned with observations that multi-headedness plays an important role when dealing with specialized heads but most of the time, can be pruned for self-attention heads [20, 21].

The input to each of the single headed graph attention layer is an relation matrix $\phi \in \Phi$ and a set of vertex features, $\mathbf{h} = \{h_1, h_2, \dots, h_{|V|}\}$ where h_i is the feature vector of vertex i for $i = 1, \dots, |V|$. Given a vertex pair (i, j) , the unnormalized attention score between them can be learned as follows:

$$e_{ij}^\phi = \text{LeakyReLU}(\mathbf{a}_\phi^\top [\mathbf{W}_\phi h_i \parallel \mathbf{W}_\phi h_j]), \quad (2)$$

where $\parallel$ is the concatenation operation, $\mathbf{a}_\phi$ represents the learnable attention vector for the respective ϕ and $\mathbf{W}_\phi$ is the weight matrix for a linear transformation specific to the relation-type matrix ϕ and is shared by all the vertices. The attention vector $\mathbf{a}_\phi$ depends on the input features, $\mathbf{W}_\phi$ as well as the relation

matrix ϕ that are input into the layer of single headed graph attention. The attention scores are then normalized by applying the softmax function to obtain the weight coefficients

$$\alpha_{ij}^\phi = \text{softmax}_j(e_{ij}^\phi) = \frac{\exp(e_{ij}^\phi)}{\sum_{k \in N_i^\phi} \exp(e_{ik}^\phi)}, \quad (3)$$

where N_i^ϕ is the neighborhood of vertex i in the specific relation-type matrix ϕ . After the normalized weight coefficients are obtained, we can then employ these coefficients along with the corresponding neighbors' features to generate the output features. These output features are specific to ϕ for each of the vertices $i = 1, \dots, |V|$, and given as

$$z_i^\phi = \sigma \left(\sum_{j \in N_i^\phi} \alpha_{ij}^\phi \mathbf{W}_\phi h_j \right), \quad (4)$$

where σ denotes an activation function. For each $\phi \in \Phi$, the output is a set of relation-type specific vertex embeddings, $Z_\phi = \{z_1^\phi, \dots, z_{|V|}^\phi\}$.

3.3. Attention to fuse different relation-types

The new vertex embeddings $\{Z_\phi : \phi \in \Phi\}$ reflect the different aspects of the vertices under different relation-types ϕ and have to be fused. A second attention layer accounts for the different contribution of each relation-type in describing the vertices, giving more weight to the more crucial relation-type and less weights to the less salient relation-type.

For each $\phi \in \Phi$, the relation-type specific vertex embeddings are transformed via a linear transformation consisting of a weight matrix $\mathbf{F}$ and a bias vector $\mathbf{b}$, then followed by a tanh function. The unnormalized attention score for each of the relation-type specific embeddings is then learnt as follows:

$$w_\phi = \frac{1}{|V|} \sum_{i \in V} \mathbf{q}^\top \tanh(\mathbf{F} z_i^\phi + \mathbf{b}), \quad (5)$$

where $\mathbf{q}$ refers to the learnable relation-type specific attention vector. The second attention layer employs averaging instead of concatenation which means that the importance of each of the relation-type w_ϕ where $\phi \in \Phi$, is generated as a weighted average of all the vertex embeddings within a group of relation-type specific embeddings. After w_ϕ is obtained, the softmax function is then applied. The normalized values are the weights of the respective relation-type denoted as β_ϕ :

$$\beta_\phi = \frac{\exp(w_\phi)}{\sum_{\phi' \in \Phi} \exp(w_{\phi'})}. \quad (6)$$

The different relation-type specific vertex embeddings are then fused to generate the final vertex embedding via linear combination using the learned weights as follows:

$$Z = \sum_{\phi \in \Phi} \beta_\phi Z_\phi. \quad (7)$$

The vertex embedding Z can then be applied to various tasks, optimising the model via backpropagation. Suppose that $|\Phi| = P$, the architecture of REGATHER is as shown in Fig. 2.

Fig. 2. Overall architecture of REGATHER.

4. EXPERIMENTS

We perform experiments on three heterogeneous graph datasets, similar to those utilized in [19] and evaluate REGATHER on vertex classification tasks. Our objective is to predict the labels of a target vertex type. The statistics of the datasets utilised are as shown in Table 1. For the loss function, cross-entropy loss is employed. Only the vertices belonging to the target type is considered when measuring cross entropy L

$$L = - \sum_{l \in y_L} Y(l) \ln(\sigma(Z(l))),$$

where σ denotes an activation function, y_L is the set of target vertex type’s indices with labels, $Y(l)$ and $Z(l)$ are the labels and learned embeddings of the labeled vertex l , respectively.

Table 1. Statistics of the heterogeneous graph datasets utilized.

Dataset	Relations (A-B)	Feature	#Label	#Vertex	#Edge	#Heterogeneous Edge	#Homogeneous Edge
DBLP	Paper-Paper	334	4	15,649	51,377	44,379	6,998
	Author-Paper						
	Venue-Paper						
Movie	Movie-Movie	334	5	21,345	89,038	34,354	54,684
	Actor-Actor						
	Director-Movie						
Foursquare	Composer-Movie	128	10	29,771	83,407	77,712	5,695
	User-User						
	User-Check_in						

We compare our model with JUST [19], GAT [4] and HAN [10], which are state-of-the-art methods to demonstrate the effectiveness of our proposed model. For GAT, the heterogeneity in vertex types and edge types are ignored. JUST is evaluated using the same settings as in [19]. For the observed vertex feature vectors, we use the JUST pre-trained embeddings for all methods to enable fair comparison as this ensures that each method’s performance is not due to engineered features that might have incorporated additional information. Nevertheless, we note that engineered features like those used in [10] can also be used with REGATHER.

The input to JUST is the heterogeneous graph and it outputs embeddings that are used as the initial features for all the other methods under comparison. The parameters in JUST like the embeddings’ dimension d , the SkipGram model’s window size k , the memorized domains m and the initial stay parameter γ , are optimized for the different datasets. For REGATHER, GAT and HAN, we evaluated all of them using the same setting as in [10]. The differences are that REGATHER uses single head attention and the highest order of relation-type K in REGATHER is set to 3. In brief, we trained the models for a maximum of 200 epochs using Adam with a learning rate of 0.005 and early stopping with a window size of 100. The relation-type specific attention vector q in HAN and REGATHER is set to a dimension of 128.

The experimental results are summarised in Table 2. We report the average Macro-F1 and Micro-F1 scores (with standard deviation) from 10 repeated trials. Our results demonstrate that our method considerably improves JUST pre-trained embeddings and in most cases, yield better results in vertex classification tasks compared to other baselines.

Table 2. Quantitative results (%) of vertex classification task.

Datasets	Metrics	Training Size	JUST	GAT	HAN	REGATHER
DBLP	Macro-F1	20%	84.41 ± 0.68	88.25±1.15	87.15±0.90	89.04±0.51
		40%	84.52 ± 0.90	89.22±1.03	88.51±1.72	90.87±0.95
		60%	85.06 ± 1.07	89.40±1.70	89.21±2.02	92.25±2.00
		80%	85.28 ± 1.63	89.54±2.61	90.13±1.26	92.66±1.70
	Micro-F1	20%	84.41 ± 0.67	88.23±1.15	87.15±0.89	89.04±0.52
		40%	84.53 ± 0.90	89.20±1.05	88.53±1.72	90.85±0.97
		60%	85.09 ± 1.07	89.41±1.71	89.25±2.04	92.25±2.01
		80%	85.29 ± 1.67	89.53±2.61	90.16±1.35	92.62±1.73
Movie	Macro-F1	20%	35.88 ± 0.86	40.64±1.29	39.36±2.82	46.87±4.04
		40%	39.12 ± 0.86	45.87±2.07	49.62±2.43	54.68±3.03
		60%	40.18 ± 0.64	47.07±4.09	53.22±2.71	57.94±2.72
		80%	41.07 ± 1.18	48.68±2.49	54.40±3.03	58.44±3.82
	Micro-F1	20%	38.19 ± 0.67	49.67±0.72	48.35±1.33	55.07±2.05
		40%	42.74 ± 0.67	54.08±1.33	56.88±2.05	60.93±2.25
		60%	44.11 ± 0.64	56.63±2.83	60.09±2.13	65.29±1.99
		80%	45.11 ± 0.94	58.81±2.38	62.84±3.07	66.50±3.08
Foursquare	Macro-F1	20%	32.53 ± 2.14	46.74±1.46	56.11±2.42	49.47±2.88
		40%	38.09 ± 1.73	59.49±2.78	64.33±2.25	61.13±3.79
		60%	41.65 ± 2.99	67.76±3.16	66.73±4.82	68.55±5.34
		80%	46.49 ± 3.41	68.24±4.75	66.39±5.48	69.19±6.13
	Micro-F1	20%	40.17 ± 1.73	52.12±1.28	61.48±1.44	56.29±1.56
		40%	44.13 ± 1.60	63.94±1.91	68.80±1.78	66.45±2.46
		60%	47.08 ± 3.17	71.36±2.73	70.29±4.17	72.77±4.12
		80%	51.76 ± 2.28	72.00±3.56	71.04±2.96	73.68±4.34

5. CONCLUSION

In this paper, we explored learning on heterogeneous graphs without using meta-paths as optimal meta-path selection is problematic especially when the dataset is complex and has many vertex types. Our method REGATHER uses relation-type subgraphs to learn representations and to capture interactions multiple hops apart. We also use dual-level attention to learn the importance of neighboring vertices and different relation-types. Experimental results demonstrated REGATHER’s effectiveness in improving JUST pre-embeddings compared to GAT and HAN across three datasets.

6. REFERENCES

- [1] Chuxu Zhang, Dongjin Song, Chao Huang, Ananthram Swami, and Nitesh V Chawla, “Heterogeneous graph neural network,” in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019, pp. 793–803.
- [2] Xiaohuan Cao, Yuyan Zheng, Chuan Shi, Jingzhi Li, and Bin Wu, “Meta-path-based link prediction in schema-rich heterogeneous information network,” *International Journal of Data Science and Analytics*, vol. 3, no. 4, pp. 285–296, 2017.
- [3] Thomas N Kipf and Max Welling, “Semi-supervised classification with graph convolutional networks,” in *International Conference on Learning Representations*, 2016.
- [4] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio, “Graph attention networks,” in *International Conference on Learning Representations*, 2017.
- [5] Huiting Hong, Hantao Guo, Yucheng Lin, Xiaoqing Yang, Zang Li, and Jieping Ye, “An attention-based graph neural network for heterogeneous structural learning,” *arXiv preprint arXiv:1912.10832*, 2019.
- [6] Chuan Shi, Yitong Li, Jiawei Zhang, Yizhou Sun, and S Yu Philip, “A survey of heterogeneous information network analysis,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 29, no. 1, pp. 17–37, 2016.
- [7] Seongjun Yun, Minbyul Jeong, Raehyun Kim, Jaewoo Kang, and Hyunwoo J Kim, “Graph transformer networks,” in *Advances in Neural Information Processing Systems*, 2019, pp. 11960–11970.
- [8] Yizhou Sun, Jiawei Han, Xifeng Yan, Philip S Yu, and Tianyi Wu, “Pathsim: Meta path-based top-k similarity search in heterogeneous information networks,” in *Proceedings of the International Conference on Very Large Data Based*, 2011, pp. 992–1003.
- [9] Yuxiao Dong, Nitesh V Chawla, and Ananthram Swami, “Metapath2vec: Scalable representation learning for heterogeneous networks,” in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2017, pp. 135–144.
- [10] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu, “Heterogeneous graph attention network,” in *Proceedings of the International World Wide Web Conference*, 2019, pp. 2022–2032.
- [11] Shaohua Fan, Junxiong Zhu, Xiaotian Han, Chuan Shi, Linmei Hu, Biyu Ma, and Yongliang Li, “Metapath-guided heterogeneous graph neural network for intent recommendation,” in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019, pp. 2478–2486.
- [12] Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun, “Heterogeneous graph transformer,” in *Proceedings of the International World Wide Web Conference*, 2020, pp. 2704–2710.
- [13] Jie Chen, Tengfei Ma, and Cao Xiao, “FastGCN: Fast learning with graph convolutional networks via importance sampling,” *arXiv preprint arXiv:1801.10247*, 2018.
- [14] Liyu Gong and Qiang Cheng, “Exploiting edge features for graph neural networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 9211–9219.
- [15] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip, “A comprehensive survey on graph neural networks,” *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [16] Changping Meng, Reynold Cheng, Silviu Maniu, Pierre Senellart, and Wangda Zhang, “Discovering meta-paths in large heterogeneous information networks,” in *Proceedings of the International Conference on World Wide Web*, 2015, pp. 754–764.
- [17] Ni Lao and William W Cohen, “Relational retrieval using a combination of path-constrained random walks,” *Machine Learning*, pp. 53–67, 2010.
- [18] Carl Yang, Mengxiong Liu, Frank He, Xikun Zhang, Jian Peng, and Jiawei Han, “Similarity modeling on heterogeneous networks via automatic path discovery,” in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2018, pp. 37–54.
- [19] Rana Hussein, Dingqi Yang, and Philippe Cudré-Mauroux, “Are meta-paths necessary? Revisiting heterogeneous graph embeddings,” in *Proceedings of the ACM International Conference on Information and Knowledge Management*, 2018, pp. 437–446.
- [20] Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov, “Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned,” *arXiv preprint arXiv:1905.09418*, 2019.
- [21] Paul Michel, Omer Levy, and Graham Neubig, “Are sixteen heads really better than one?,” in *Advances in Neural Information Processing Systems*, 2019, pp. 14014–14024.