

Graph Transformer for Graph-to-Sequence Learning*

Deng Cai, Wai Lam

The Chinese University of Hong Kong
thisisjcykcd@gmail.com, wlam@se.cuhk.edu.hk

Abstract

The dominant graph-to-sequence transduction models employ graph neural networks for graph representation learning, where the structural information is reflected by the receptive field of neurons. Unlike graph neural networks that restrict the information exchange between immediate neighborhood, we propose a new model, known as Graph Transformer, that uses explicit relation encoding and allows direct communication between two distant nodes. It provides a more efficient way for global graph structure modeling. Experiments on the applications of text generation from Abstract Meaning Representation (AMR) and syntax-based neural machine translation show the superiority of our proposed model. Specifically, our model achieves 27.4 BLEU on LDC2015E86 and 29.7 BLEU on LDC2017T10 for AMR-to-text generation, outperforming the state-of-the-art results by up to 2.2 points. On the syntax-based translation tasks, our model establishes new single-model state-of-the-art BLEU scores, 21.3 for English-to-German and 14.1 for English-to-Czech, improving over the existing best results, including ensembles, by over 1 BLEU.

Introduction

Graphical structure plays an important role in natural language processing (NLP), they often serve as the central formalism for representing syntax, semantics, and knowledge. For example, most syntactic representations (e.g., dependency relation) are tree-based while most whole-sentence semantic representation frameworks (e.g., Abstract Meaning Representation (AMR) (Banarescu et al. 2013)) encode sentence meaning as directed acyclic graphs. A range of NLP applications can be framed as the process of graph-to-sequence learning. For instance, text generation may involve realizing a semantic graph into a surface form (Liu et al. 2015) and syntactic machine translation incorporates source-side syntax information for improving translation quality (Bastings et al. 2017). Fig. 1 gives an example of AMR-to-text generation.

*The work described in this paper is substantially supported by a grant from the Research Grant Council of the Hong Kong Special Administrative Region, China (Project Code: 14204418). Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

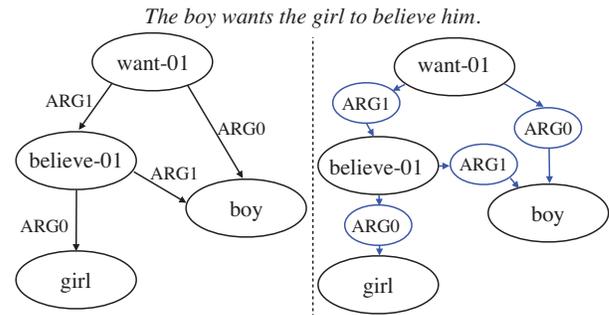


Figure 1: An AMR graph (left) for the reference sentence “The boy wants the girl to believe him.” and the corresponding Levi graph (right).

While early work uses statistical methods or neural models after the linearization of graphs, graph neural networks (GNNs) have been firmly established as the state-of-the-art approaches for this task (Damonte and Cohen 2019; Guo et al. 2019). GNNs typically compute the representation of each node iteratively based on those of its adjacent nodes. This inherently local propagation nature precludes efficient global communication, which becomes critical at larger graph sizes, as the distance between two nodes exceeds the number of stacked layers. For instance, for two nodes staying L hops away, at least L layers will be needed in order to capture their dependencies. Furthermore, even if two distant nodes are reachable, the information may also be disrupted in the long journey (Xu et al. 2018; Guo et al. 2019).

To address the above problems, we propose a new model, known as Graph Transformer, which relies entirely on the multi-head attention mechanism (Vaswani et al. 2017) to draw global dependencies.¹ Different to GNNs, the Graph Transformer allows modeling of dependencies between any two nodes without regard to their distance in the input graph. One undesirable consequence is that it essentially treats any

¹We note that the name *Graph Transformer* was used in a recent work (Koncel-Kedziorski et al. 2019). However, it merely focuses on the relations between directly connected nodes as other graph neural networks.

graph as a fully connected graph, greatly diluting the explicit graph structure. To maintain a graph structure-aware view, our proposed model introduces explicit relation encoding and incorporates it into the pairwise attention score computation as a dynamic parameter.

Our treatment of explicit relation encoding also brings other side advantages compared to GNN-based methods. Previous state-of-the-art GNN-based methods use Levi graph transformation (Beck, Haffari, and Cohn 2018; Guo et al. 2019), where two unlabeled edges are replacing one labeled edge that is present in the original graph. For example, in Fig. 1, the labeled edge $\text{want-01} \xrightarrow{\text{ARG1}} \text{believe-01}$ turns to be two unlabeled edges $\text{want-01} \rightarrow \text{ARG1}$ and $\text{ARG1} \rightarrow \text{believe-01}$. Since edge labels are represented as nodes, they end up sharing the same semantic space, which is not ideal as nodes and edges are typically different elements. In addition, the Levi graph transformation at least doubles the number of representation vectors, which will introduce more complexity for the decoder-side attention mechanism (Bahdanau, Cho, and Bengio 2015) and copy mechanism (Gu et al. 2016; See, Liu, and Manning 2017). Through explicit and separate relation encoding, our proposed Graph Transformer inherently avoids these problems.

Experiments show that our model is able to achieve better performance for graph-to-sequence learning tasks for natural language processing. For the AMR-to-text generation task, our model surpasses the current state-of-the-art neural methods trained on LDC2015E86 and LDC2017T10 by 1.6 and 2.2 BLEU points, respectively. For the syntax-based neural machine translation task, our model is also consistently better than others, even including ensemble systems, showing the effectiveness of the model on a large training set. In addition, we give an in-depth study of the source of improvement gain and the internal workings of the proposed model.

Related Work

Early research efforts for graph-to-sequence learning use specialized grammar-based methods. Flanigan et al.(2016) split input graphs to trees and uses a tree-to-string transducer. Song et al.(2016) recast generation as a traveling salesman problem. Jones et al.(2012) leverage hyperedge replacement grammar and Song et al.(2017) use a synchronous node replacement grammar. More recent work employs more general approaches, such as phrase-based machine translation model (Pourdamghani, Knight, and Hermjakob 2016) and neural sequence-to-sequence methods (Konstas et al. 2017) after linearizing input graphs. Regarding AMR-to-text generation, Cao and Clark(2019) propose an interesting idea that factorizes text generation through syntax. One limitation of sequence-to-sequence models, however, is that they require serialization of input graphs, which inevitably incurs the obstacle of capturing graph structure information.

An emerging trend has been directly encoding the graph with different variants of graph neural networks, which in common stack multiple layers that restrict the update of node representation based on a first-order neighborhood but

use different information passing schemes. Some borrow the ideas from recurrent neural networks (RNNs), e.g, Beck, Haffari, and Cohn(2018) use gated graph neural network (Li et al. 2016) while Song et al.(2018) introduce LSTM-style information aggregation. Others apply convolutional neural networks (CNNs), e.g., Bastings et al.(2017);Damonte and Cohen(2019);Guo et al.(2019) utilize graph convolutional neural networks (Kipf and Welling 2017). Koncel-Kedziorski et al.(2019) update vertex information by attention over adjacent neighbors. Furthermore, Guo et al.(2019) allow the information exchange across different levels of layers. Damonte and Cohen(2019) systematically compare different encoders and show the advantages of graph encoder over tree and sequential ones. The contrast between our model and theirs is reminiscent of the contrast between the self-attention network (SAN) and CNN/RNN.

For sequence-to-sequence learning, the SAN-based Transformer model (Vaswani et al. 2017) has been the *de facto* approach for its empirical successes. However, it is unclear on the adaptation to graphical data and its performance. Our work is partially inspired by the introduction of relative position embedding (Shaw, Uszkoreit, and Vaswani 2018; Dai et al. 2019) in sequential data. However, the extension to graph is nontrivial since we need to model much more complicated relation instead of mere visual distance. To the best of our knowledge, the Graph Transformer is the first graph-to-sequence transduction model relying entirely on self-attention to compute representations.

Background of Self-Attention Network

The Transformer introduced by Vaswani et al.(2017) is a sequence-to-sequence neural architecture originally used for neural machine translation. It employs self-attention network (SAN) for implementing both the encoder and the decoder. The encoder consists of multiple identical blocks, of which the core is multi-head attention. The multi-head attention consists of H attention heads, and each of them learns a distinct attention function. Given a source vector $x \in \mathbb{R}^{d_x}$ and a set of context vectors $\{y_1, y_2, \dots, y_m\}$ with the same dimension d_x or in short $y_{1:m}$, for each attention head, x and $y_{1:m}$ are transformed into distinct query and value representations. The attention score is computed as the dot-product between them.

$$f(x, y_i) = (W_q x)^T W_k y_i$$

where $W_q, W_k \in \mathbb{R}^{d_z \times d_x}$ are trainable projection matrices. The attention scores are scaled and normalized by a softmax function to compute the final attention output $attn$.

$$a_i = \frac{\exp(f(x, y_i)/\sqrt{d_z})}{\sum_{j=1}^m \exp(f(x, y_j)/\sqrt{d_z})}$$

$$attn = \sum_{i=1}^m a_i W_v y_i$$

where $a \in \mathbb{R}^m$ is the attention vector (a distribution over all input $y_{1:m}$), $W_v \in \mathbb{R}^{d_z \times d_x}$ is a trainable projection matrix. Finally, the outputs of all attention heads are concatenated and projected to the original dimension of x , followed by

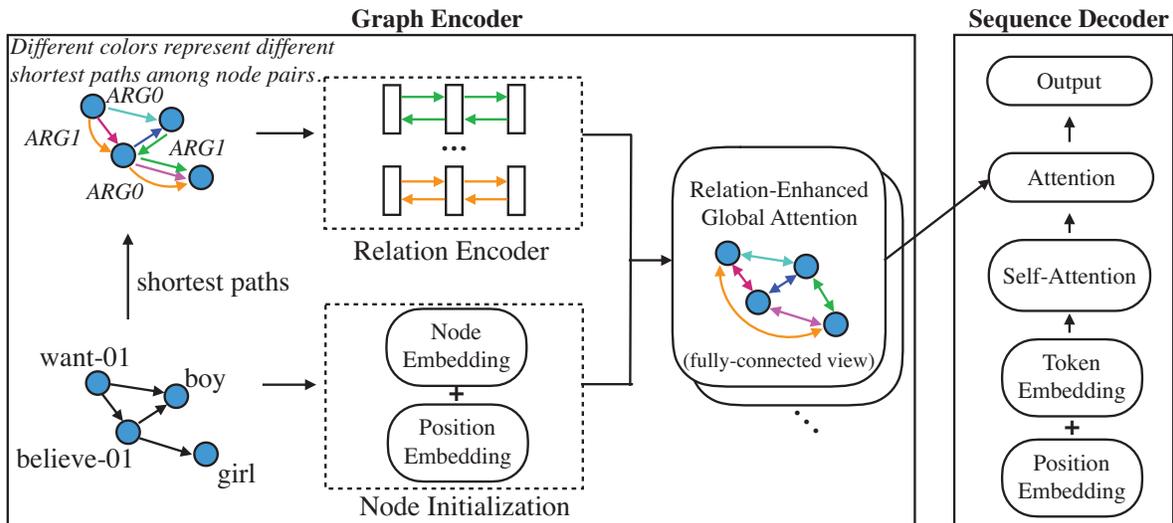


Figure 2: An overview of our propose model.

feed-forward layers, residual connection, and layer normalization.² For brevity, we will denote the whole procedure described above as a single function $\text{ATT}(x, y_{1:m})$.

For an input sequence $x_{1:n}$, the SAN-based encoder computes the vector representations iteratively by $x_i^L = \text{ATT}(x_i^L, x_{1:n}^{L-1})$, where L is the total number of blocks and $x_{1:n}^0$ are word embeddings. In this way, a representation is allowed to build a direct relationship with another long-distance representation. To feed the sequential order information, the deterministic or learned position embedding (Vaswani et al. 2017) is introduced to expose the position information to the model, i.e., x_i^0 becomes the sum of the corresponding word embedding and the position embedding for i .

The aforementioned treatment of SAN on sequential data can be drawn a close resemblance to graph neural networks by regarding the token sequence as an unlabeled fully-connected graph (each token as a node) and taking the multi-head attention mechanism as a specific message-passing scheme. Such view on the relationship between SAN and graph neural networks inspires our work.

Graph Transformer

Overview

For a graph with n nodes, previous graph neural networks compute the node representation v_i as a function of the input node i and all its first-order neighborhoods $N(i)$. The graph structure is implicitly reflected by the receptive field of each node representation. This local communication design, however, could be inefficient for long-distance information exchange. We introduce a new model, known as Graph Transformer, which provides an aggressively different paradigm that enables relation-aware global communication.

²We refer interesting readers to Vaswani et al.(2017) for more details.

The overall framework is shown in Fig. 2. The most important characteristic of the Graph Transformer is that it has a fully-connected view on arbitrary input graphs. A node is able to directly receive and send information to another node no matter whether they directly connected or not. These operations are achieved by our proposed extension to the original multi-head attention mechanism, the relation-enhanced global attention mechanism described below. Specifically, the relationship between any node pair is depicted as the shortest relation path between them. These pairwise relation paths are fed into a relation encoder for distributed relation encoding. We initialize node vectors as the sum of the node embedding and absolute position embeddings. Multiple blocks of global attention network are stacked to compute the final node representations. At each block, a node vector is updated based on all other node vectors and the corresponding relation encodings. The resulted node vectors at the last block are fed to the sequence decoder for sequence generation.

Graph Encoder

Our graph encoder is responsible for transforming an input graph into a set of corresponding node embeddings. To apply global attention on a graph, the central problem is how to maintain the topological structure of the graph while allowing fully-connected communication. To this end, we propose relation-enhanced global attention mechanism, which is an extension of the vanilla multi-head attention. Our idea is to incorporate explicit relation representation between two nodes into their representation learning. Recall that, in the standard multi-head attention, the attention score between the element x_i and the element x_j is simply the dot-product of their query vector and key vector respectively:

$$\begin{aligned} s_{ij} &= f(x_i, x_j) \\ &= x_i W_q^T W_k x_j \end{aligned} \quad (1)$$

Suppose we have learned a vector representation for the relationship r_{ij} , which we will refer as relation encoding, between the node i and the node j . Following the idea of relative position embedding (Shaw, Uszkoreit, and Vaswani 2018; Dai et al. 2019), we propose to compute the attention score as follows:

$$[r_{i \rightarrow j}; r_{j \rightarrow i}] = W_r r_{ij} \quad (2)$$

where we split the relation encoding r_{ij} into the forward relation encoding $r_{i \rightarrow j}$ and the backward relation encoding $r_{j \rightarrow i}$. Then we compute the attention score based on both the node representations and their relation representation as shown below:

$$\begin{aligned} s_{ij} &= g(x_i, x_j, r_{ij}) \\ &= (x_i + r_{i \rightarrow j}) W_q^T W_k (x_j + r_{j \rightarrow i}) \\ &= \underbrace{x_i W_q^T W_k x_j}_{(a)} + \underbrace{x_i W_q^T W_k r_{j \rightarrow i}}_{(b)} \\ &\quad + \underbrace{r_{i \rightarrow j} W_q^T W_k x_j}_{(c)} + \underbrace{r_{i \rightarrow j} W_q^T W_k r_{j \rightarrow i}}_{(d)} \end{aligned} \quad (3)$$

Each term in Eq (3) corresponds to some intuitive meaning according to their formalization. The term (a) captures purely content-based addressing, which is the original term in vanilla attention mechanism. The term (b) represents a source-dependent relation bias. The term (c) governs a target-dependent relation bias. The term (d) encodes the universal relation bias. Our formalization provides a principled way to model the element-relation interactions. In comparison, it has broader coverage than Shaw, Uszkoreit, and Vaswani(2018) in terms of additional terms (c) and (d), and than Dai et al.(2019) in terms of the extra term (c) respectively. More importantly, previous methods only model the relative position in the context of sequential data, which merely adopts the immediate embeddings of the relative positions (e.g., $-1, +1$). To depict the relation between two nodes in a graph, we utilize a shortest-path based approach as described below.

Relation Encoder Conceptually, the relation encoding gives the model a global guidance about how information should be gathered and distributed, i.e., where to attend. For most graphical structures in NLP, the edge label conveys direct relationship between adjacent nodes (e.g., the semantic role played by concept-to-concept, and the dependency relation between two words). We extend this one-hop relation definition into multi-hop relation reasoning for characterizing the relationship between two arbitrary nodes. For example, in Fig 1, the shortest path from the concept `want-01` to `girl` is “`want-01` $\xrightarrow{ARG1}$ `believe-01` $\xrightarrow{ARG0}$ `girl`”, which conveys that `girl` is the object of the *wanted* action. Intuitively, the shortest path between two nodes gives the closest and arguably the most important relationship between them. Therefore, we propose to use the shortest paths (relation sequence) between two nodes to characterize their relationship.³ Following the sequential nature of the rela-

³For the case that there are multiple shortest paths, we randomly sample one during training and take the averaged representation

tion sequence, we employ recurrent neural networks with Gated Recurrent Unit (GRU) (Cho et al. 2014) for transforming relation sequence into a distributed representation. Formally, we represent the shortest relation path $sp_{i \rightarrow j} = [e(i, k_1), e(k_1, k_2), \dots, e(k_n, j)]$ between the node i and the node j , where $e(\cdot, \cdot)$ indicates the edge label and $k_{1:n}$ are the relay nodes. We employ bi-directional GRUs for sequence encoding:

$$\begin{aligned} \vec{s}_t &= \text{GRU}_f(\vec{s}_{t-1}, sp_t) \\ \overleftarrow{s}_t &= \text{GRU}_b(\overleftarrow{s}_{t+1}, sp_t) \end{aligned}$$

The last hidden states of the forward GRU network and the backward GRU networks are concatenated to form the final relation encoding $r_{ij} = [\vec{s}_n; \overleftarrow{s}_0]$.

Bidirectionality Though in theory, our architecture can deal with arbitrary input graphs, the most widely adopted graphs in the real problems are directed acyclic graphs (DAGs). This implies that the node embedding information will be propagated in one pre-specified direction. However, the reverse direction informs the equivalent information flow as well. To facilitate communication in both directions, we add reverse edges to the graph. The reverse edge connects the same two nodes as the original edge but in a different direction and with a reversed label. For example, we will draw a virtual edge `believe-01` $\xrightarrow{RARG1}$ `want-01` according to the original edge `want-01` $\xrightarrow{ARG1}$ `believe-01`. For convenience, we also introduce self-loop edges for each node. These extra edges have specific labels, hence their own parameters in the network. We also introduce an extra global node into every graph, who has a direct edge to all other nodes with the special label *global*. The final representation x_{global} of the global node serves as a whole graph representation.

Absolute Position Besides pairwise relationship, some absolute positional information can also be beneficial. For example, the root of an AMR graph serves as a rudimentary representation of the overall focus, making the minimum distance from the root node partially reflect the importance of the corresponding concept in the whole-sentence semantics. The sequence order of tokens in a dependency tree also provides complementary information to dependency relations. In order for the model to make use of the absolute positions of nodes, we add the positional embeddings to the input embeddings at the bottom of the encoder stacks. For example, `want-01` in Fig 1 is the root node of the AMR graph, so its index should be 0. Notice we denote the index of the global node as 0 as well.

Sequence Decoder

Our sequence decoder basically follows the same spirit of the sequential Transformer decoder. The decoder yields the natural language sequence by calculating a sequence of hidden states sequentially. One distinct characteristic is that we use the global graph representation x_{global} for initializing the hidden states at each time step. The hidden state h_t at during testing.

Dataset	#train	#dev	#test	#edge types	#node types	avg #nodes	avg #edges	avg diameter
LDC2015E86	16,833	1,368	1,371	113	18735	17.34	17.53	6.98
LDC2017T10	36,521	1,368	1,371	116	24693	14.51	14.62	6.15
English-Czech	181,112	2,656	2,999	46	78017	23.18	22.18	8.36
English-German	226,822	2,169	2,999	46	87219	23.29	22.29	8.42

Table 1: Data statistics of all four datasets. #train/dev/test indicates the number of instances in each set, avg #nodes/edges/diameter represents the averaged value of nodes/edge/diameter size of a graph.

model component	hyper-parameter	value
char-level CNN	number of filters	256
	width of filters	3
	char embedding size	32
	final hidden size	128
Embeddings	node embedding size	300
	edge embedding size	200
	token embedding size	300
Multi-head attention	number of heads	8
	hidden state size	512
	feed-forward hidden size	1024

Table 2: Hyper-parameters settings.

each time step t is then updated by interleaving multiple rounds of attention over the output of the encoder (node embeddings) and attention over previously-generated tokens (token embeddings). Both are implemented by the multi-head attention mechanism. x_{global} is removed when performing the sequence-to-graph attention.

Copy mechanism To address the data sparsity issue in token prediction, we include a copy mechanism (Gu et al. 2016) in similar spirit to previous works. Concretely, a single-head attention is computed based on the decoder state h_t and the node representation $x_{1:n}$, where a_t^i denotes the attention weight of the node v_i in the current time step t . Our model can either directly copy the type name of a node or generate from a pre-defined vocabulary V . Formally, the prediction probability of a token y is given by:

$$P(y|h_t) = P(gen|h_t)gen(y|h_t) + P(copy|h_t) \sum_{i \in S(y)} a_t^i$$

where $S(y)$ is the set of nodes that have the same surface form as y . $P(gen|h_t)$ and $P(copy|h_t)$ are computed by a single layer neural network with softmax activation, and $gen(y|h_t) = \exp(w_y^T h_t) / \sum_{y' \in V} \exp(w_{y'}^T h_t)$, where w_y (for $y \in V$) denotes the model parameters. The copy mechanism facilitates the generation of dates, numbers, and named entities in both AMR-to-text generation and machine translation tasks in experiments.

Experiments

We assess the effectiveness of our models on two typical graph-to-sequence learning tasks, namely AMR-to-text generation and syntax-based machine translation (MT). Following previous work, the results are mainly evaluated by BLEU (Papineni et al. 2002) and CHRF++ (Popović 2017).

Specifically, we use case-insensitive scores for AMR and case sensitive BLEU scores for MT.

AMR-to-text Generation

Our first application is language generation from AMR, a semantic formalism that represents sentences as rooted DAGs (Banarescu et al. 2013). For this AMR-to-text generation task, we use two benchmarks, namely the LDC2015E86 dataset and the LDC2017T10 dataset. The first block of Table 1 shows the statistics of the two datasets. Similar to Konstas et al.(2017), we apply entity simplification and anonymization in the preprocessing steps and restore them in the postprocessing steps.

The graph encoder uses randomly initialized node embeddings as well as the output from a learnable CNN with character embeddings as input. The sequence decoder uses randomly initialized token embeddings and another char-level CNN. Model hyperparameters are chosen by a small set of experiments on the development set of LDC2017T10. The detailed settings are listed in Table 2. During testing, we use a beam size of 8 for generating graphs. To mitigate overfitting, we also apply dropout (Srivastava et al. 2014) with the drop rate of 0.2 between different layers. We use a special UNK token to replace the input node tag with a rate of 0.33. Parameter optimization is performed with the Adam optimizer (Kingma and Ba 2014) with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The same learning rate schedule of Vaswani et al.(2017) is adopted in our experiments.⁴ For computation efficiency, we gather all distinct shortest paths in a training/testing batch, and encode them into vector representations by the recurrent relation encoding procedure as described above.⁵

We run comparisons on systems without ensembling nor additional silver data. Specifically, the comparison methods can be grouped into three categories: (1) feature-based statistical methods (Song et al. 2016; Pourdamghani, Knight, and Hermjakob 2016; Song et al. 2017; Flanigan et al. 2016); (2) sequence-to-sequence neural models (Konstas et al. 2017; Cao and Clark 2019), which use linearized graphs as inputs; (3) recent works using different variants of graph neural networks for encoding graph structures directly (Song et al. 2018; Beck, Haffari, and Cohn 2018; Damonte and Cohen 2019; Guo et al. 2019). The results are shown in Table 3. For both datasets, our approach substantially outperforms

⁴Code available at <https://github.com/jcyk/gtos>.

⁵This strategy reduces the number of related sequences to encode from $O(mn^2)$ to a stable number when a large batch size m is used.

Model	LDC2015E86			LDC2017T10		
	BLEU	CHRFF++	METEOR	BLEU	CHRFF++	METEOR
Song et al.(2016) [†]	22.4	-	-	-	-	-
Flanigan et al.(2016) [†]	23.0	-	-	-	-	-
Pourdamghani, Knight, and Hermjakob(2016) [†]	26.9	-	-	-	-	-
Song et al.(2017) [†]	25.6	-	-	-	-	-
Konstas et al.(2017)	22.0	-	-	-	-	-
Cao and Clark(2019) [‡]	23.5	-	-	26.8	-	-
Song et al.(2018)	23.3	-	-	24.9	-	-
Beck, Haffari, and Cohn(2018)	-	-	-	23.3	50.4	-
Damonte and Cohen(2019)	24.4	-	23.6	24.5	-	24.1
Guo et al.(2019)	25.7	54.5*	31.5*	27.6	57.3	34.0*
Ours	27.4	56.4	32.9	29.8	59.4	35.1

Table 3: Main results on AMR-to-text generation. Numbers with * are from the contact from the authors. - denotes that the result is unknown because it is not provided in the corresponding paper.

Model	Type	English-German		English-Czech	
		BLEU	CHRFF++	BLEU	CHRFF++
Bastings et al.(2017)	Single	16.1	-	9.6	-
Beck, Haffari, and Cohn(2018)	Single	16.7	42.4	9.8	33.3
Guo et al.(2019)	Single	19.0	44.1	12.1	37.1
Beck, Haffari, and Cohn(2018)	Ensemble	19.6	45.1	11.7	35.9
Guo et al.(2019)	Ensemble	20.5	45.8	13.1	37.8
Ours	Single	21.3	47.9	14.1	41.1

Table 4: Main results on syntax-based machine translation.

all previous methods. On the LDC2015E86 dataset, our method achieves a BLEU score of 27.4, outperforming previous best-performing neural model (Guo et al. 2019) by a large margin of 2.6 BLEU points. Also, our model becomes the first neural model that surpasses the strong non-neural baseline established by Pourdamghani, Knight, and Hermjakob(2016). It is worth noting that those traditional methods marked with [†] train their language models on the external Gigaword corpus, thus they possess an additional advantage of extra data. On the LDC2017T10 dataset, our model establishes a new record BLEU score of 29.8, improving over the state-of-the-art sequence-to-sequence model (Cao and Clark 2019) by 3 points and the state-of-the-art GNN-based model (Guo et al. 2019) by 2.2 points. The results are even more remarkable since the model of Cao and Clark(2019) (marked with [‡]) uses constituency syntax from an external parser. Similar phenomena can be found on the additional metrics of CHRFF++ and METEOR (Denkowski and Lavie 2014). Those results suggest that current graph neural networks cannot make full use of the AMR graph structure, and our Graph Transformer provides a promising alternative.

Syntax-based Machine Translation

Our second evaluation is syntax-based machine translation, where the input is a source language dependency syntax tree and the output is a plain target language string. We employ the same data and settings from Bastings et al.(2017). Both the English-German and the English-Czech datasets from

the WMT16 translation task.⁶ The English sentences are parsed after tokenization to generate the dependency trees on the source side using SyntaxNet (Alberti et al. 2017).⁷ On the Czech and German sides, texts are tokenized using the Moses tokenizer.⁸ Byte-pair encodings (Sennrich, Haddow, and Birch 2016) with 8,000 merge operations are used to obtain subwords. The second block of Table 1 shows the statistics for both datasets. For model configuration, we just re-use the settings obtained in our AMR-to-text experiments.

Table 4 presents the results with comparison to existing methods. On the English-to-German translation task, our model achieves a BLEU score of 41.0, outperforming all of the previously published single models by a large margin of 2.3 BLEU score. On the English-to-Czech translation task, our model also outperforms the best previously reported single models by an impressive margin of 2 BLEU points. In fact, our single model already outperforms previous state-of-the-art models that use ensembling. The advantages of our method are also verified by the metric CHRFF++.

An important point about these experiments is that we did not tune the architecture: we simply employed the same model in all experiments, only adjusting the batch size for different dataset size. We speculate that even better results would be obtained by tuning the architecture to individual tasks. Nevertheless, we still obtained improved performance over previous works, underlining the generality of

⁶<http://www.statmt.org/wmt16/translation-task.html>.

⁷<https://github.com/tensorflow/models/tree/master/syntaxnet>

⁸<https://github.com/moses-smt/mosesdecoder>.

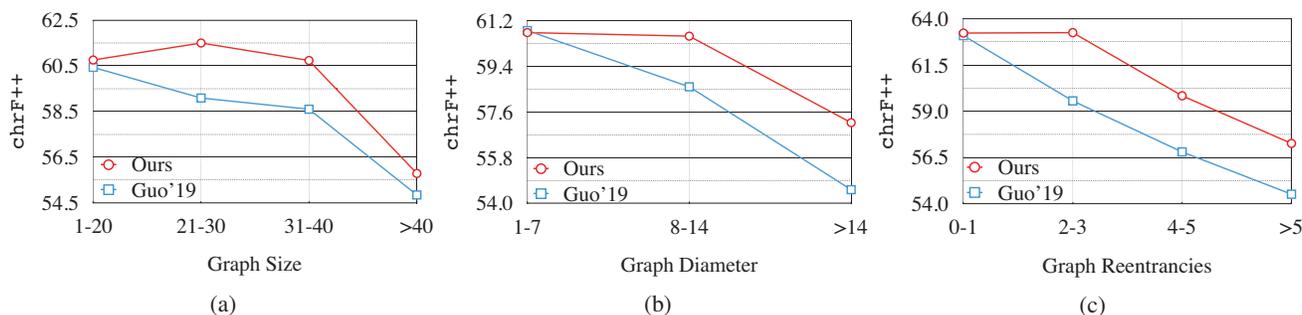


Figure 3: CHRf++ scores with respect to (a) the graph size, (b) the graph diameter, and (c) the the number of reentrancies.

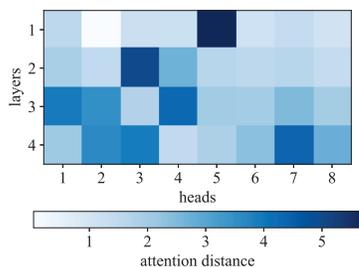


Figure 4: The average distance for maximum attention for each head.

our model.

More Analysis

The overall scores show a great advantage of the Graph Transformer over existing methods, including the state-of-the-art GNN-based models. However, they do not shed light into how this is achieved. In order to further reveal the source of performance gain, we perform a series of analysis based on different characteristics of graphs. For those analyses, we use sentence-level CHRf++ scores, and take the macro average of them when needed. All experiments are conducted with the test set of LDC2017T10.

Graph Size To assess the model’s performance for different sizes of graphs, we group graphs into four classes and show the curves of CHRf++ scores in Figure 3a. The results are presented with the contrast with the state-of-the-art GNN-based model of Guo et al.(2019), denoted as Guo’19. As seen, the performance of both models decreases as the graph size increases. It is expected since a larger graph often contains more complex structure and the interactions between graph elements are more difficult to capture. The gap between ours and Guo’19 becomes larger for relatively larger graphs while for small graphs, both models give similar performance. This result demonstrates that our model has better ability for dealing with complicated graphs. As for extremely large graphs, the performance of both models have a clear drop, yet ours is still slightly better.

Graph Diameter We then study the impact of graph diameter.⁹ Graphs with large diameters have interactions between two nodes that appear distant from each other. We conjecture that it will cause severe difficulties for GNN-based models because they solely rely on local communication. Figure 3b confirms our hypothesis, as the curve of the GNN-based model shows a clear slope. In contrast, our model has more stable performance, and the gap between the two curves also illustrates the superiority of our model on featuring long-distance dependencies.

Number of Reentrancies We study the ability for handling the reentrancies, where the same node has multiple parent nodes (or the same concept participates in multiple relations for AMR). The recent work (Damonte and Cohen 2019) has been identified reentrancies as one of the most difficult aspects of AMR structure. We bin the number of reentrancies occurred in a graph into four classes and plot Fig. 3c. It can be observed that the gap between the GNN-based model and the Graph transformer becomes noticeably wide when more than one reentrancies start to happen. Since then, our model is consistently better than the GNN-based model, maintaining a margin of over 1 CHRf++ score.

How Far Does Attention Look At The Graph Transformer shows a strong capacity for processing complex and large graphs. We attribute the success to the global communication design, as it provides opportunities for direct communication in long distance. A natural and interesting question is how well the model makes use of this property. To answer this question, following Voita et al.(2019), we turn to study the attention distribution of each attention head. Specifically, we record the specific distance of its maximum attention weight is assigned to. Fig. 4 shows the averaged the attention distance after we run on the development set of LDC2017T10. We can observe that nearly half of the attention heads have an average attention distance larger than 2. The number of these distance heads generally increases as layers go deeper. Interestingly, the longest-reaching head (layer1-head5) and the shortest-sighted head (layer1-head2) coexist in the very first layer, while the former has an average distance over 5.

⁹The diameter of a graph is defined as the length of the longest shortest path between two nodes.

Conclusions

In this paper, we presented the Graph Transformer, the first graph-to-sequence learning based entirely on automatic attention. Different from previous recurrent models that require linearization of input graph and previous graph neural network models that restrict the message passing in the first-order neighborhood, our model enables global node-to-node communication. With the Graph Transformer, we achieve the new state-of-the-art on two typical graph-to-sequence generation tasks with four benchmark datasets.

References

- Alberti, C.; Andor, D.; Bogatyy, I.; Collins, M.; Gillick, D.; Kong, L.; Koo, T.; Ma, J.; Omerick, M.; Petrov, S.; et al. 2017. Syntaxnet models for the conll 2017 shared task. *arXiv preprint arXiv:1703.04929*.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Banarescu, L.; Bonial, C.; Cai, S.; Georgescu, M.; Griffitt, K.; Hermjakob, U.; Knight, K.; Koehn, P.; Palmer, M.; and Schneider, N. 2013. Abstract meaning representation for sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, 178–186.
- Bastings, J.; Titov, I.; Aziz, W.; Marcheggiani, D.; and Sima'an, K. 2017. Graph convolutional encoders for syntax-aware neural machine translation. In *EMNLP*, 1957–1967.
- Beck, D.; Haffari, G.; and Cohn, T. 2018. Graph-to-sequence learning using gated graph neural networks. In *ACL*, 273–283.
- Cao, K., and Clark, S. 2019. Factorising AMR generation through syntax. In *NAACL*, 2157–2163.
- Cho, K.; Van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Dai, Z.; Yang, Z.; Yang, Y.; Carbonell, J.; Le, Q.; and Salakhutdinov, R. 2019. Transformer-XL: Attentive language models beyond a fixed-length context. In *ACL*, 2978–2988.
- Damonte, M., and Cohen, S. B. 2019. Structural neural encoders for AMR-to-text generation. In *NAACL*, 3649–3658.
- Denkowski, M., and Lavie, A. 2014. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the EACL 2014 Workshop on Statistical Machine Translation*.
- Flanigan, J.; Dyer, C.; Smith, N. A.; and Carbonell, J. 2016. Generation from abstract meaning representation using tree transducers. In *NAACL*, 731–739.
- Gu, J.; Lu, Z.; Li, H.; and Li, V. O. 2016. Incorporating copying mechanism in sequence-to-sequence learning. In *ACL*, 1631–1640.
- Guo, Z.; Zhang, Y.; Teng, Z.; and Lu, W. 2019. Densely connected graph convolutional networks for graph-to-sequence learning. *Transactions of the Association for Computational Linguistics* 7:297–312.
- Jones, B.; Andreas, J.; Bauer, D.; Hermann, K. M.; and Knight, K. 2012. Semantics-based machine translation with hyperedge replacement grammars. In *COLING*, 1359–1376.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kipf, T. N., and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*.
- Koncel-Kedziorski, R.; Bekal, D.; Luan, Y.; Lapata, M.; and Hajishirzi, H. 2019. Text Generation from Knowledge Graphs with Graph Transformers. In *NAACL*, 2284–2293.
- Konstas, I.; Iyer, S.; Yatskar, M.; Choi, Y.; and Zettlemoyer, L. 2017. Neural AMR: Sequence-to-sequence models for parsing and generation. In *ACL*, 146–157.
- Li, Y.; Tarlow, D.; Brockschmidt, M.; and Zemel, R. 2016. Gated graph sequence neural networks. In *ICLR*.
- Liu, F.; Flanigan, J.; Thomson, S.; Sadeh, N.; and Smith, N. A. 2015. Toward abstractive summarization using semantic representations. In *NAACL*, 1077–1086.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 311–318.
- Popović, M. 2017. chr++: words helping character n-grams. In *Proceedings of the second conference on machine translation*, 612–618.
- Pourdamghani, N.; Knight, K.; and Hermjakob, U. 2016. Generating english from abstract meaning representations. In *INLG*, 21–25.
- See, A.; Liu, P. J.; and Manning, C. D. 2017. Get to the point: Summarization with pointer-generator networks. In *ACL*, 1073–1083.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016. Neural machine translation of rare words with subword units. In *ACL*, 1715–1725.
- Shaw, P.; Uszkoreit, J.; and Vaswani, A. 2018. Self-attention with relative position representations. In *NAACL*, 464–468.
- Song, L.; Zhang, Y.; Peng, X.; Wang, Z.; and Gildea, D. 2016. AMR-to-text generation as a traveling salesman problem. In *EMNLP*, 2084–2089.
- Song, L.; Peng, X.; Zhang, Y.; Wang, Z.; and Gildea, D. 2017. AMR-to-text generation with synchronous node replacement grammar. In *ACL*, 7–13.
- Song, L.; Zhang, Y.; Wang, Z.; and Gildea, D. 2018. A graph-to-sequence model for AMR-to-text generation. In *ACL*, 1616–1626.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research* 15(1):1929–1958.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *NIPS*, 5998–6008.
- Voita, E.; Talbot, D.; Moiseev, F.; Sennrich, R.; and Titov, I. 2019. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In *ACL*, 5797–5808.
- Xu, K.; Wu, L.; Wang, Z.; Feng, Y.; Witbrock, M.; and Sheinin, V. 2018. Graph2seq: Graph to sequence learning with attention-based neural networks. *arXiv preprint arXiv:1804.00823*.