

# Graph-to-Tree Neural Networks for Learning Structured Input-Output Translation with Applications to Semantic Parsing and Math Word Problem

Shucheng Li<sup>\*†</sup>, Lingfei Wu<sup>\*‡</sup>, Shiwei Feng<sup>†</sup>, Fengyuan Xu<sup>†</sup>, Fangli Xu<sup>§</sup>, Sheng Zhong<sup>†</sup>

<sup>†</sup>National Key Lab for Novel Software Technology, Nanjing University, China

<sup>‡</sup>IBM Thomas J. Watson Research Center, Yorktown Heights, NY, USA

<sup>§</sup>Squirrel AI Learning, New Jersey, USA

<sup>†</sup>{shuchengli, fengyuan.xu, swfeng98, sheng.zhong}@smail.nju.edu.cn

<sup>‡</sup>lwu@email.wm.edu, <sup>§</sup>lili@yixue.us

## Abstract

The celebrated Seq2Seq technique and its numerous variants achieve excellent performance on many tasks such as neural machine translation, semantic parsing, and math word problem solving. However, these models either only consider input objects as sequences while ignoring the important structural information for encoding, or they simply treat output objects as sequence outputs instead of structural objects for decoding. In this paper, we present a novel Graph-to-Tree Neural Networks, namely Graph2Tree consisting of a graph encoder and a hierarchical tree decoder, that encodes an augmented graph-structured input and decodes a tree-structured output. In particular, we investigated our model for solving two problems, neural semantic parsing and math word problem. Our extensive experiments demonstrate that our Graph2Tree model outperforms or matches the performance of other state-of-the-art models on these tasks.

## 1 Introduction

Learning general functional dependency between arbitrary input and output spaces is one of the key challenges in machine learning. While many efforts in machine learning have mainly focused on designing flexible and powerful input representations for solving classification or regression problems, many applications require researchers to design novel models that can deal with complex structured inputs and outputs, such as graphs, trees, sequences, or sets. In this paper, we consider the general problem of learning a mapping between a graph input  $G \in \mathcal{G}$  and a tree output  $T \in \mathcal{T}$ , based on a training sample of structured input-output pairs  $(G_1, T_1), \dots, (G_n, T_n) \in \mathcal{G} \times \mathcal{T}$  drawn from some fixed but unknown probability distribution.

Such learning problems often arise in a variety of applications, ranging from semantic parsing, to

|     |  |
|-----|--|
| SP  | <b>Text Input:</b><br>what jobs are there for web developer who know 'c++' ?<br><br><b>Structured output:</b><br>answer( A , ( job ( A ) , title ( A , W ) , const ( W , 'Web Developer' ) , language ( A , C ) , const ( C , 'c++' ) ) )      |
| MWP | <b>Text input:</b><br>0.5 of the cows are grazing grass . 0.25 of the cows are sleeping and 9 cows are drinking water from the pond . find the total number of cows .<br><br><b>Structured output:</b><br>$((0.5 * x) + (0.25 * x)) + 9.0 = x$ |

Table 1: Examples of structured input and output of semantic parsing (SP) and math word problem (MWP). For inputs, we consider parsing tree augmented sequences to get structural information. For outputs, they are naturally a hierarchical structure with some structural meaning symbols like brackets.

math word problem, label sequence learning, and supervised grammar learning, to name just a few. As shown in Fig. 1, finding the parse tree of a sentence involves a structural dependency among the labels in the parse tree; generating a mathematical expression of a math word problem involves a hierarchical dependency between math logical operations and the numbers. Conventionally, there have been efforts in generalizing kernel methods to predict structured and inter-dependent variables in a supervised learning setting (Tsochantaridis et al., 2005; Altun et al., 2004; Joachims et al., 2009).

Recently, the celebrated Sequence-to-Sequence technique (Seq2Seq) and its numerous variants (Sutskever et al., 2014; Bahdanau et al., 2014; Luong et al., 2015) achieve excellent performance in neural machine translation. Encouraged by the success of Seq2Seq model, there is a surge of interests in applying Seq2Seq models to cope with other tasks such as developing neural semantic parser (Dong and Lapata, 2016) or solving math word problem (Ling et al., 2017). However, the two significant challenges making a Seq2Seq model ineffective in these tasks are that, i) for the natural

\* authors contributed equally to this research.

text description input, it often entails some hidden syntactic structure information such as dependency, constituency tree or even semantic structure information like AMR parsing tree; ii) for the meaningful representation output, it typically contains abundant information in a structured object like a parsing tree or a mathematical equation.

Inspired by these observations, in this work, we propose a Graph-to-Tree neural networks, namely Graph2Tree consisting of a graph encoder and a hierarchical tree decoder, which leverages the structural information of both source graphs and target trees. In particular, our Graph2Tree model learns the mapping from a structured object such as a graph to another structured object such as a tree. In addition, we also observe that the structured object translation typically follows a modular procedure, which translates the individual sub-graph in the source graph into the corresponding target one in target tree output, and then compose them to form the final target tree.

Therefore, we design a workflow to align with this procedure: our graph encoder first learns from an input graph that is constructed from the various inputs such as combining both a word sequence and the corresponding dependency or constituency tree, and then our tree decoder generates the tree object from the learned graph vector representations to explicitly capture the compositional structure of a tree. In particular, we present a novel Graph2tree model with a separated attention mechanism to jointly learn a final hidden vector of the corresponding graph nodes in order to align the generation process between a heterogeneous graph input and a hierarchical tree output.

To demonstrate the effectiveness of our model, we perform experiments on two important tasks – Semantic Parsing and Math Word Problem. First, we compare our approach against several neural network approaches on the Semantic Parsing task. Our experimental results show that our Graph2Tree model could outperform or match the performance of other state-of-the-art models on three standard benchmark datasets. Second, we further compare our approach with existing recently developed neural approaches on the math word problem and our results clearly show that our Graph2Tree model can achieve state-of-the-art performance compared to other baselines that use many task-specific techniques. We believe our Graph2Tree model is a solid attempt for learning structured input-output

translation.

## 2 Related Works

### 2.1 Graph Neural Networks

The graph representation learning recently attracted a lot of attention and interest from both academia and industry. One of the most important research lines is the semantic embedding learning of graph nodes or edges based upon the power of graph neural networks (GNNs) (Li et al., 2016; Kipf and Welling, 2017; Velickovic et al., 2017; Gilmer et al., 2017; Hamilton et al., 2017).

Encouraged by the recent success in GNNs, various Sequence-to-Graph (Peng et al., 2018) or Graph-to-Sequence models (Xu et al., 2018a,b,c; Beck et al., 2018; Chen et al., 2020) have been proposed to handle the structured inputs, structured outputs or both of them, i.e. generating AMR graph generation from the text sequence. More recently, some researchers proposed the Tree-to-Tree (Chen et al., 2018b), Graph-to-Tree (Yin et al., 2019) and Graph-to-Graph (Guo et al., 2018) neural networks for targeted application scenarios.

However, these works are designed exclusively for specific downstream tasks like program translation or code edit. Compared to them, our proposed Graph2Tree neural network with novel design of graph encoder and tree decoder does not rely on any specific downstream task assumption. Additionally, our Graph2Tree is the first generic neural network translating graph inputs into tree outputs, which may have numerous applications in practice.

### 2.2 Neural Semantic Parsing

Semantic parsing is the task of translating natural language utterances into machine-interpretable meaning representations like logical forms or SQL queries. Recent years have witnessed a surge of interests in developing neural semantic parsers with sequence-to-sequence models. These parsers have achieved promising results (Jia and Liang, 2016; Dong and Lapata, 2016; Ling et al., 2016). Due to the fact that the meaning representations are usually structured objects (e.g. tree structures), many efforts have been devoted to develop structure-oriented decoders, including tree decoders (Dong and Lapata, 2016; Alvarez-Melis and Jaakkola, 2017), grammar constrained decoders (Yin and Neubig, 2017; Yin et al., 2018; Jie and Lu, 2018; Dong and Lapata, 2018), action sequences for semantic graph generation (Chen et al., 2018a), and

modular decoders based on abstract syntax trees (Rabinovich et al., 2017). However, those approaches could potentially be further improved because they only consider the word sequence information and ignore other rich syntactic information, such as dependency or constituency tree, available at the encoder side.

Researchers recently attempted to leverage of the power of GNNs in various NLP tasks, including the neural machine translation (Bastings et al., 2017; Beck et al., 2018), conversational machine reading comprehension (Chen et al., 2019b), and AMR-to-text (Song et al., 2018). Specifically in the semantic parsing field, a general Graph2Seq model (Xu et al., 2018b) is proposed to incorporate these dependency and constituency trees with the word sequence and then create a syntactic graph as the encoding input. However, this approach simply treats a logical form as a sequence, neglecting the abundant information in a structured object like tree in the decoder architecture. Therefore, we present the Graph2Tree model to utilize the structure information in both structured inputs and outputs.

### 2.3 Math Word Problems

The math word problem is the task of translating the short paragraph (typically consisting with multiple short sentences) into succinct mathematical equations. To solve a math word problem illustrated in Table 1, traditional approaches focus on generating numeric answer expressions by mapping verbs in problems text to categories (Hosseini et al., 2014) or by generating templates from problem texts (Kushman et al., 2014). However, these approaches either need additional hand-crafted annotations for problem texts or are limited to a set of predefined equation templates.

Inspired by the great success of Seq2Seq models in Neural Machine Translation, deep-learning based methods are intensively explored by researchers in the equation generation (Wang et al., 2017; Ling et al., 2017; Li et al., 2018, 2019; Zou and Lu, 2019; Xie and Sun, 2019). However, different forms of equations can be formed to solve the same math problem, which often makes models fail. To resolve the equation duplication issues, various equation normalization methods are proposed in (Wang et al., 2018a, 2019) to generate a unique expression tree with the cost of losing the understanding of problem-solving steps in equation expressions. In contrast, we propose to use a

Graph2Tree model to solve this task without any special mechanisms like equation normalization. To the best of our knowledge, this is the first work to use GNN to build a math word problem solver.

## 3 Problem Formulation and Structure Object Construction

### 3.1 Graph-to-Tree Translation Task

In this work, we consider the problem of translating a graph input to a tree output. In particular, we consider two important tasks - Semantic Parsing and Math Word Problem. Formally, we define both tasks as follows. The input side contains a set of text sequences, denoted as  $S = \{s_1, s_2, \dots, s_n\} \in \mathcal{S}$  where  $s_i$  is a text sequence consisting of a sequence of word embeddings  $s_i = \{w_1, w_2, \dots, w_{|s_i|}\} \in \mathcal{W}$ , where  $\mathcal{W}$  is a pre-trained word embedding space. We then construct a heterogeneous graph input  $G = (V, E) \in \mathcal{G}$ , where  $V = [V_1 V_2]$  contains all of the original word nodes  $V_1 \in \mathcal{V}_1$  as well as the relationship nodes  $V_2 \in \mathcal{V}_2$  from the relationships of a parsing tree (i.e. dependency or constituency tree), and  $E \in \mathcal{E}$  denotes if the two nodes are connected or not. The aim is to translate a set of heterogeneous graph inputs  $G = \{g_1, g_2, \dots, g_n\}$  into a set of tree outputs  $T = \{t_1, t_2, \dots, t_n\} \in \mathcal{T}$  where  $t_i$  is a logic form or math equation consisting of a sequence of tree node token  $t_i = \{y_1, y_2, \dots, y_{|t_i|}\} \in \mathcal{Y}$ .

### 3.2 Constructing Graph Inputs and Tree Outputs

To apply GNNs, the first step is to construct a graph input by combining the word sequence with their corresponding hidden structure information. How to construct such graphs is critical to incorporate the structured information and influences the final performance. Similarly, how to construct the tree outputs from logic form or math equations also play an important role in the final performance and model interpretability. In this section, we will introduce two methods for graph construction and one method for tree construction.




Figure 1: Dependency tree augmented text graph

**Combining Word Sequence with Dependency Parse Tree.** The dependency parse tree not only represents various grammatical relationships between pairs of text words, but also is shown to have an important role in transforming texts into logical forms (Reddy et al., 2016). Therefore, the first method integrates two types of features by adding dependency linkages between corresponding word pairs in word sequence. Concretely, we transform a dependency label into a node, which is linked respectively with two word nodes with dependency relationship. Figure 1 gives such an example of constructed heterogeneous graph from a text.

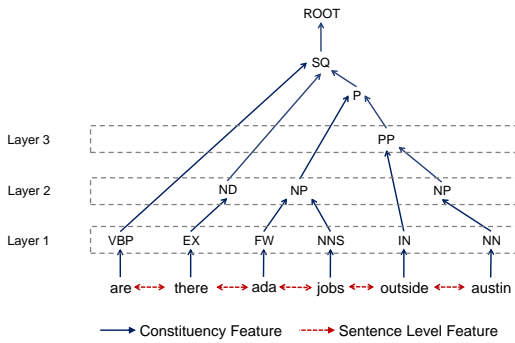


Figure 2: Constituency tree augmented text graph

**Combining Word Sequence with Constituency Tree.** The constituency tree contains the phrase structure information which is also critical to describe the word relationships and has shown to provide useful information for translation (Gū et al., 2018). Since the leaf nodes in the constituency tree are the word nodes in the text, this method merges these nodes with the identical ones in the bi-directional word sequence chain to create the syntactic graph. Figure 2 shows an example of constructed heterogeneous graph input.

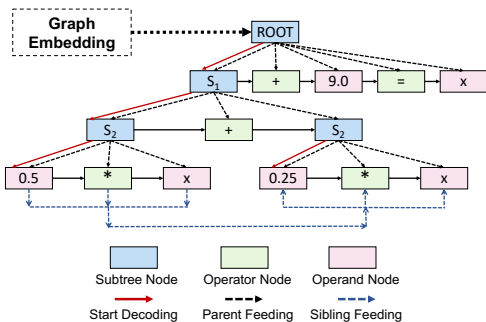


Figure 3: A sample tree output in our decoding process from expression  $((0.5 * x) + (0.25 * x)) + 9.0 = x$

**Constructing Tree Outputs.** To effectively learn the compositional nature of our structured outputs,

we need to firstly transform original outputs from logic forms or math equations to tree structured objects. Specifically, we follow the tree construction method in (Dong and Lapata, 2016), which is a top-down manner to generate tree-structured outputs. In original outputs containing structural meaning symbols like brackets, we first extract sub-tree structures and replace these sub-tree structures with sub-tree symbols. Then we grow branches from the generated sub-tree symbols until all hierarchical structures in the original sequence are processed. Figure 3 provides an example of constructed tree objects from mathematical expression.

## 4 Graph2Tree Neural Networks

We aim to learn a mapping that translates a heterogeneous graph-structured input  $G$  and its corresponding tree-structured outputs  $T$ . We illustrate the workflow of our proposed Graph2Tree model for semantic parsing in Figure 4, and present each component of the model as follows.

### 4.1 Graph Encoder

To effectively learn graph representations from our constructed heterogeneous text graph, we present a novel bidirectional graph node embeddings method - BiGraphSAGE. The proposed BiGraphSAGE extends the widely used GraphSAGE (Hamilton et al., 2017) by learning forward and backward node embeddings of a graph  $G$  in an interleaved fashion.

In particular, consider a word node  $v \in \mathcal{V}_1$  with pretrained word embedding  $\mathbf{w}_v$  like GloVe (Pennington et al., 2014) as  $v$ 's initial attributes. We then generate the contextualized node embeddings  $\mathbf{a}_v$  for all nodes  $v \in \mathcal{V}_1$  using Bi-directional Long Short Term Memory (BiLSTM) (Graves et al., 2013). For a relationship node  $v \in \mathcal{V}_2$ , we initialize  $\mathbf{a}_v$  with randomized embeddings. These feature vectors are used as initial node embeddings  $\mathbf{h}_v^0 = \mathbf{a}_v$ . Then each node embedding learns its vector representation by aggregating information from a node local neighborhood within  $K$  hops of the graph.

$$\mathbf{h}_{\mathcal{N}_+(v)}^k = \mathbf{M}_+^k(\{\mathbf{h}_{u^+}^{k-1}, \forall u \in \mathcal{N}_+(v)\}) \quad (1)$$

$$\mathbf{h}_{\mathcal{N}_-(v)}^k = \mathbf{M}_-^k(\{\mathbf{h}_{u^-}^{k-1}, \forall u \in \mathcal{N}_-(v)\}) \quad (2)$$

where  $k \in \{1, \dots, K\}$  is the iteration index and  $\mathcal{N}$  is the neighborhood function of node  $v$ .  $\mathbf{M}_+^k$  and  $\mathbf{M}_-^k$  are the forward and backward aggregator functions. Node  $v$ 's forward (backward) representation  $\mathbf{h}_{v^+}^k$  ( $\mathbf{h}_{v^-}^k$ ) aggregates the information of nodes in  $\mathcal{N}_+(v)$  ( $\mathcal{N}_-(v)$ ).

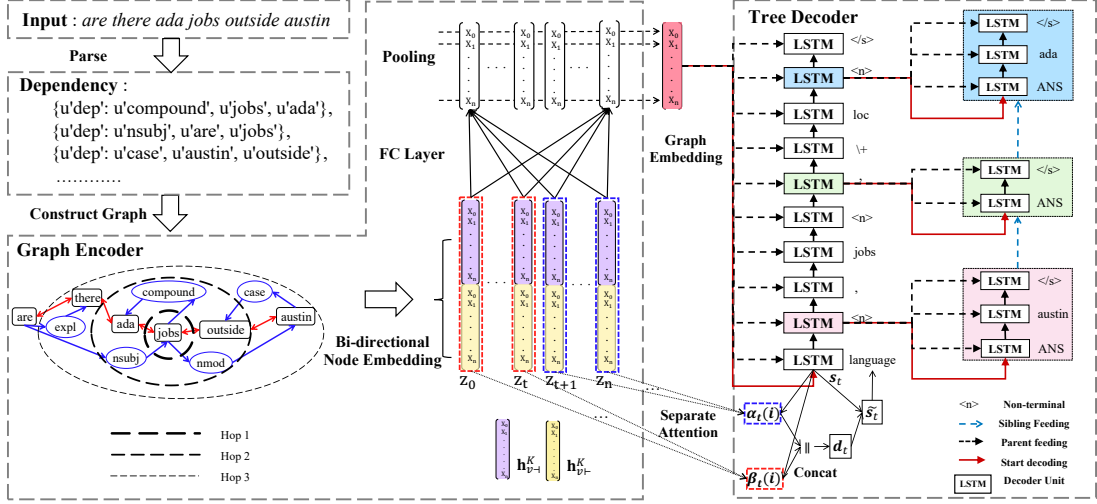


Figure 4: Overall architecture of our Graph2Tree model. We use semantic parsing task as an example.

Conceptually, one can choose to keep these node embeddings for each direction independently, which ignores interactions between two intermediate node embeddings during the training. Therefore, we fuse two intermediate unidirectional node embeddings at each hop as follows,

$$\mathbf{h}_1 = \mathbf{h}_{\mathcal{N}_+(v)}^k, \mathbf{h}_2 = \mathbf{h}_{\mathcal{N}_-(v)}^k \quad (3)$$

$$\mathbf{h}_{\mathcal{N}(v)}^k = \mathbf{w}_g \odot \mathbf{h}_1 + (1 - \mathbf{w}_g) \odot \mathbf{h}_2, \quad (4)$$

$$\mathbf{w}_g = \sigma(\vec{W}_z [\mathbf{h}_1; \mathbf{h}_2; \mathbf{h}_1 \odot \mathbf{h}_2; \mathbf{h}_1 - \mathbf{h}_2]) \quad (5)$$

where  $\odot$  denotes component-wise multiplication,  $\sigma$  is a sigmoid function and  $\mathbf{w}_g$  is a gating vector.

The graph encoder learns node embeddings  $\mathbf{h}_v^k$  by repeating the following process  $K$  times:

$$\mathbf{h}_v^k = \sigma(\mathbf{W}^k \cdot \text{CONCAT}(\mathbf{h}_v^{k-1}, \mathbf{h}_{\mathcal{N}(v)}^k)) \quad (6)$$

where  $\mathbf{W}^k$  denotes weight matrices,  $\sigma$  is a non-linearity function,  $K$  is maximum number of hops.

The final bi-directional node embeddings  $\mathbf{z}_v$  is chosen to concatenate the two unidirectional node embeddings at the last hop,

$$\mathbf{z}_v = \text{CONCAT}(\mathbf{h}_{v+}^K, \mathbf{h}_{v-}^K) \quad (7)$$

$$\mathbf{g} = \text{MAXPOOL}(\text{FC}(\mathbf{z})). \quad (8)$$

After the bi-directional embeddings for all nodes  $\mathbf{z}$  are computed, we then feed the obtained node embeddings into a fully-connected neural network and apply the element-wise *max*-pooling operation on all node embeddings to compute the graph-level vector representation  $\mathbf{g}$ , where other alternative commutative operations such as *mean* or attention based weighted *sum* can be used as well.

## 4.2 Tree Decoder

We propose a new general tree decoder fully leveraging the outputs of our graph encoder, i.e. the bi-directional node embeddings and the graph embedding, and faithfully generating the tree-structured targets like logic forms or math equations.

Inspired by the thinking paradigm of human beings, our tree decoder at high level uses a divide-and-conquer strategy splitting the whole decoding task into sub ones. Figure 3 illustrates an example output of our tree decoder. In this example, we firstly initialize the root tree node *ROOT* with the graph embedding  $\mathbf{g}$ , and then apply a sub-decoder on the *ROOT* to generate a 1st-level coarse output containing a sub-tree node  $S_1$ . This  $S_1$  is further decoded with the similar sub-decoder to derive the 2nd-level coarse output. This procedure is repeated to generate the 3rd-level output in which there is no sub-tree nodes. In this way, we get the whole tree output in a top-down manner.

This whole procedure can be summarized as follows: 1) initialize the root tree node with the graph embedding from our encoder and perform the first level decoding with our LSTM based sub-decoder; 2) for each newly generated sub-tree node, a sub-decoder is applied to derive the next level coarse output; 3) repeat step 2 until there is no sub-tree nodes in the last level of tree structure.

### 4.2.1 Sub-Decoder Design

In each of our sub-decoder task, the conditional probability of the generated word at step  $t$  is calculated as follows:

$$p(y_t | \mathbf{y}_{<t}, \mathbf{x}) = f_{predict}(\mathbf{s}_t) \quad (9)$$

where  $\mathbf{x}$  denotes vectors of all input words,  $y_t$  is the predicted output word at  $t$ ,  $\mathbf{s}_t$  is the decoder hidden state at  $t$ , and  $f_{predict}$  is a non-linear function.

The key component of Eq. (9) is the computation of  $\mathbf{s}_t$ . Conceptually, this value is calculated as  $\mathbf{s}_t = f_{decoder}(y_{t-1}, \mathbf{s}_{t-1})$ , where  $f_{decoder}$  is usually a RNN unit. We propose two improvements on top of it, *parent feeding* and *sibling feeding*, to feed more information for decoding sub-tree nodes.

**Parent feeding.** For a sub-task in our tree decoding process, we aim to expand the sub-tree node in the parent layer. Therefore, it is reasonable to take the sub-tree node embedding  $\mathbf{s}_i$  into consideration. Therefore, we add the sub-tree node embedding as part of our input at every time-step, in order to capture the upper-layer information for decoding.

**Sibling feeding.** Besides the information from parent nodes, if two sub-tree nodes share the same parent node, then these two sub-tasks can also be related. Inspired by this observation, we employ the sibling feeding mechanism to feed the preceding sibling sentence embedding to the sub-task related to its closet neighbor sub-tree node. For example, imagine  $p_1$  is the parent node of  $c_1$ ,  $c_2$ , and we feed both embeddings of  $p_1$  and  $c_1$  when decoding  $c_2$ .

Therefore, our sub-decoder calculates the decoder hidden state  $\mathbf{s}_t$  as follows:

$$\mathbf{s}_t = f_{decoder}(y_{t-1}, \mathbf{s}_{t-1}; \mathbf{s}_{parent}; \mathbf{s}_{sibling}) \quad (10)$$

where  $\mathbf{s}_{parent}$  stands for sub-tree node embedding from parent layer and  $\mathbf{s}_{sibling}$  is the sentence embedding of the closest preceding sibling. By fully utilizing the information from parent nodes and sibling nodes, our tree decoder can effectively generate target hierarchical outputs.

### 4.3 Separate Attention Mechanism to Locate Source Sub-graph

Various attention mechanisms have been proposed (Bahdanau et al., 2014; Luong et al., 2015) to incorporate the hidden vectors of the inputs into account during the decoding processing. In particular, the context vector  $\mathbf{s}_t$  depends on a set of bidirectional node representations of the source graph  $(\mathbf{z}_1, \dots, \mathbf{z}_{|V|})$  to which the decoder locates the source sub-graph. Since our graph input is essentially a heterogeneous graph with two different input sources (word nodes with relationship nodes of a parsing tree), we propose to employ a separated

attention mechanism over the node representations corresponding to the different node types:

$$\alpha_{t(v)} = \frac{\exp(\text{score}(\mathbf{z}_v, \mathbf{s}_t))}{\exp(\sum_{k=1}^{V_1} \text{score}(\mathbf{z}_k, \mathbf{s}_t))}, \forall v \in \mathcal{V}_1 \quad (11)$$

$$\beta_{t(v)} = \frac{\exp(\text{score}(\mathbf{z}_v, \mathbf{s}_t))}{\exp(\sum_{k=1}^{V_2} \text{score}(\mathbf{z}_k, \mathbf{s}_t))}, \forall v \in \mathcal{V}_2 \quad (12)$$

where the  $\text{score}(\cdot)$  function estimates the similarity of  $\mathbf{z}_v$  and  $\mathbf{s}_t$ . Then, we compute the context vectors  $\mathbf{c}_{v1}$  and  $\mathbf{c}_{v2}$ , respectively.

$$\mathbf{c}_{v1} = \sum \alpha_{t(v)} \mathbf{z}_v, \forall v \in \mathcal{V}_1 \quad (13)$$

$$\mathbf{c}_{v2} = \sum \beta_{t(v)} \mathbf{z}_v, \forall v \in \mathcal{V}_2 \quad (14)$$

We concatenate the context vector  $\mathbf{c}_{v1}$ , context vector  $\mathbf{c}_{v2}$  and decoder hidden state  $\mathbf{s}_t$  to compute the final attention hidden state at this time step as:

$$\tilde{\mathbf{s}}_t = \tanh(W_c \cdot [\mathbf{c}_{v1}; \mathbf{c}_{v2}; \mathbf{s}_t] + b_c) \quad (15)$$

where  $W_c$  and  $b_c$  are learnable parameters. The final context vector  $\tilde{\mathbf{s}}_t$  is further used for decoding tree structured outputs. The output probability distribution over a vocabulary at the current time step is calculated by:

$$p(y_t | y_1, y_2, \dots, y_{t-1}, g) = \text{softmax}(W_v \tilde{\mathbf{s}}_t + b_v) \quad (16)$$

where  $W_v$  and  $b_v$  are learnable parameters. Our model is then jointly trained to maximize the conditional log-probability of the target tree given a heterogeneous graph input  $g$ .

## 5 Experiments

In this section, we evaluate the effectiveness and generality of Graph2Tree model on two important tasks – Semantic Parsing and Math Word Problem. The code and data for our Graph2Tree model are provided for research purpose<sup>1</sup>.

### 5.1 Experiments for Semantic Parsing

**Datasets.** We evaluate our Graph2Tree on three totally-different benchmark datasets, JOBS (Zettlemoyer and Collins, 2005), GEO (Zettlemoyer and Collins, 2005), and ATIS (Dahl et al., 1994), for the semantic parsing task. The first one JOBS is a set of 640 queries from a job listing database, the second one GEO is a set of 880 queries on a database of U.S. geography, and the last one ATIS is a dataset of 5410 queries from a flight booking system. We utilize the same train/dev/test split standard as used in previous works. We adopt the data preprocessing provided by (Dong and Lapata,

<sup>1</sup><https://github.com/IBM/Graph2Tree>

2016). Natural language utterances are in lower case and stemmed, and entity mentions are replaced by numbered markers. For the graph construction, we use the dependency parser and constituency parser from CoreNLP (Manning et al., 2014).

**Settings.** We use the Adam optimizer (Kingma and Ba, 2014) with a batch size of 20. For the JOBS and GEO datasets, our hyper-parameters are cross-validated on the training sets. For ATIS, we tune them on the development set. The learning rate is set to 0.001. In graph encoder, the BiRNN we use is a one-layer BiLSTM with a hidden size of 150, and the hop size in GNN is chosen from {2,3,4,5,6}. The decoder we employ is a one-layer LSTM with a hidden size of 300. The dropout rate is chosen from {0.1,0.3,0.5}.

**Baselines.** We compare our model against several state-of-the-art neural semantic parsers: i) Seq2Seq model with a Copy mechanism (Jia and Liang, 2016); ii) Seq2Seq and Seq2Tree models (Dong and Lapata, 2016); iii) Graph2Seq model (Xu et al., 2018a). We report the exact-match accuracy for each baseline on all three benchmarks.

| Methods                                | JOBS        | GEO         | ATIS        |
|--|-------------|-------------|-------------|
| Jia et al.(2016)                       | -           | 85.0        | 76.3        |
| Dong et al.(2016)-Seq2Seq              | 87.1        | 84.6        | 84.2        |
| Dong et al.(2016)-Seq2Tree             | 90.0        | 87.1        | <b>84.6</b> |
| Xu et al.(2018)-Graph2Seq <sup>2</sup> | 88.6        | 85.7        | 83.3        |
| <b>Graph2Tree</b>                      | <b>92.9</b> | <b>88.9</b> | <b>84.6</b> |

Table 2: Exact-match accuracy comparison on all three benchmarks *JOBS*, *GEO*, and *ATIS* for SP task

| Methods       | Translated logic form results  |
|---------------|--|
| Reference str | job (ANS, 'language (ANS, 'delphi'), title (ANS, 'developer'), loc (ANS, 'san antonio'), platform (ANS, 'windows') |
| Graph2tree    | job (ANS, 'language (ANS, 'delphi'), title (ANS, 'developer'), loc (ANS, 'san antonio'), platform (ANS, 'windows') |
| Graph2seq     | job (ANS, 'language (ANS, 'delphi'), title (ANS, 'developer'), platform (ANS, 'windows')                           |
| Seq2seq       | job (ANS, 'language (ANS, 'delphi'), title (ANS, 'developer'), loc (ANS, 'san antonio')                            |

Table 3: Case study of SP input: “*what jobs can a delphi developer find in san antonio on windows ?*”

**Results.** Table 2 shows that our proposed Graph2Tree outperforms or achieves comparable exact-match accuracy compared to other state-of-the-art baselines, highlighting the effectiveness of our proposed model by exploiting full utilization of structural information in both inputs and outputs.

<sup>2</sup>We run our own implementation of Graph2Seq on these-datasets using PyTorch.

**Case study.** Next we analyze the different decoding results of all models for an example case in Table 3. The challenge in semantic parsing is the high-order neighborhood estimation of the noun key word “jobs” to its attribute words “windows” and “san antonio”. It is hard for the traditional sequence encoder to encode high-order neighborhood (long-range dependency). For instance, there are 10 hops between the word “jobs” and “windows” according to the sequential dependency, while there are only two hops if we introduce the syntactic dependency information. Therefore, syntactic graph with graph encoder is an effective way to learn a high-quality representation for decoding. This partially explains why our Graph2tree model outperforms Seq2Seq and Seq2Tree models.

| Methods                      | JOBS | GEO  |
|------------------------------|------|------|
| Full model                   | 92.9 | 88.9 |
| w/o const tree               | 90.0 | 86.8 |
| w/ original GraphSage        | 90.7 | 88.2 |
| w/ only parent feeding       | 91.4 | 87.9 |
| w/ only sibling feeding      | 89.2 | 84.3 |
| w/o parent & sibling feeding | 88.6 | 83.9 |
| w/o separated attention      | 83.6 | 77.1 |
| w/ uniform attention         | 90.7 | 87.1 |
| w/o bilstm                   | 89.3 | 86.4 |

Table 4: Ablation study of Graph2Tree on the semantic parsing (*JOBS* and *GEO*). We employ exact match accuracy as evaluation metric.

**Ablation study.** Table 4 presents the ablation study on our Graph2Tree using a constituency tree based graph (on SP datasets *JOBS* and *GEO*). This is done with test sets (*JOBS* and *GEO* have no dev set). Firstly, We observe that the syntactic information in constituency tree, which is helpful for describing word relationships, is critical to our overall performance. And we found that our bidirectional GraphSAGE, encoding from both forward and backward nodes according to edge direction, is proved to enhance the final performance. Furthermore, parent feeding and sibling feeding mechanism, which can enrich both the paternal and fraternal information in decoding, also play important roles in the whole model. In addition, designed for different types of nodes in the input graph, the separate attention mechanism is proved useful in our model. Last but not least, it is also necessary to use Bi-LSTM in encoder to learn the contextualized word embeddings from the word sequences.

## 5.2 Experiments for Math Word Problems

**Datasets.** We here evaluate our Graph2Tree model on two benchmark datasets, MAWPS (Koncel-Kedziorski et al., 2016) and MATHQA (Amini et al., 2019), for the Math Word Problems automatically solving task. The MAWPS dataset is a Math Word Problem dataset in English and contains 2373 pairs after harvesting equations with single unknown variable. The other MATHQA dataset is a recently proposed large-scale Math Word Problem dataset with 37k English pairs, where each math expression is corresponding to an annotated formula for better interpretability. This dataset is more difficult for covering complex multivariate problems.

**Baselines.** We compare our Graph2Tree model against several state-of-the-art methods. We report the solution accuracy for each baseline in test set. On MAWPS, our baselines are: i) Retrieval, Classification, and Seq2Seq (Robaidek et al., 2018); ii) Seq2Tree (Dong and Lapata, 2016); iii) Graph2Seq (Xu et al., 2018a); iv) MathDQN (Wang et al., 2018b); v) T-RNN (Wang et al., 2019); vi) Group-Att (Li et al., 2019). On MATHQA, our baselines are: i) Sequence-to-program (Amini et al., 2019); ii) TP-N2F (Chen et al., 2019a); iii) Seq2Seq, Seq2Tree and Graph2Seq.

| Methods        |                            | MAWPS       |
|----------------|----------------------------|-------------|
| Oracle         |                            | 84.8        |
| Retrieval      | Jaccard                    | 45.6        |
|                | Cosine                     | 38.8        |
| Classification | BiLSTM                     | 62.8        |
|                | Self-attention             | 60.4        |
| Seq2seq        | LSTM                       | 25.6        |
|                | CNN                        | 44.0        |
| Seq2Tree       |                            | 65.2        |
| Graph2Seq      |                            | 70.4        |
| MathDQN        |                            | 60.25       |
| T-RNN          | Full model                 | 66.8        |
|                | W/o equation normalization | 63.9        |
|                | W/o self-attention         | 66.3        |
| Group-Att      |                            | 76.1        |
| Graph2Tree     | with constituency graph    | <b>78.8</b> |
|                | with dependency graph      | 76.8        |

Table 5: Solution accuracy comparison on MAWPS

**Results.** As shown in Table 5, our Graph2Tree model consistently outperforms other state-of-the-art baselines by a large margin up to 10 points absolute accuracy except Group-Att baseline. To the best of our knowledge, we make the first attempt to employ the graph neural network for solving Math Word Problems, and our Graph2Tree model with constituency graph achieves the best performance

| Methods      |                         | MATHQA       |
|--------------|-------------------------|--------------|
| Seq2Prog     |                         | 51.9         |
| Seq2Prog+Cat |                         | 54.2         |
| TP-N2F       |                         | 55.95        |
| Seq2seq      |                         | 58.36        |
| Seq2Tree     |                         | 64.15        |
| Graph2Seq    |                         | 65.36        |
| Graph2Tree   | with constituency graph | <b>69.65</b> |
|              | with dependency graph   | 65.66        |

Table 6: Solution accuracy comparison on MATHQA

so far on this MAWPS benchmark. We have observed similar conclusions on a more challenging and larger dataset – MATHQA. This highlights the importance of having our Graph2Tree neural networks that can leverage the structured information from both inputs and outputs for automatic solving of math problems.

It is worth noting that our hierarchical tree decoder directly generates original mathematical expressions, which faithfully reflect reasoning steps when building math equations. However, state-of-the-art math word problem solvers like Group-Att (Li et al., 2019) or T-RNN (Wang et al., 2019) have achieved high performance by utilizing Equation Normalization (EN) proposed by (Wang et al., 2019) to keep structures of output equations unified. This method can improve solution accuracy because it reduces the difficulty of equation generation. On the other hand, the normalized equations completely lose the semantic meaning of operands and operators, making them difficult to reason rationales how answer math equations are built.

**Attention visualization.** For better understanding of our separated attention, we give a visualization sample from MAWPS. As shown in Figure 5(a), we give an augmented graph input and equation tree, where  $\langle N \rangle$  is sub-tree node and  $1, 2$  are indexed markers for original numbers. Specifically, Figure 5(b) and 5(c) illustrates alignments with word nodes and compositional nodes in graph input respectively. For example, in Figure 5(c), the equation part “ $2 * 1$ ” is matched with “a bee has 2 legs” in the original natural language sentence which is actually semantically connected with “NP” and “VP” in the constituency tree.

**Ablation study.** Similarly, we also perform the ablation study for math word problem (MAWPS), as shown in Table 7. This is done with dev set. Attention mechanism, constituency structure and other components in our model play significant roles for Graph2tree to achieve high performance in MWP solving, which is consistent with our ob-



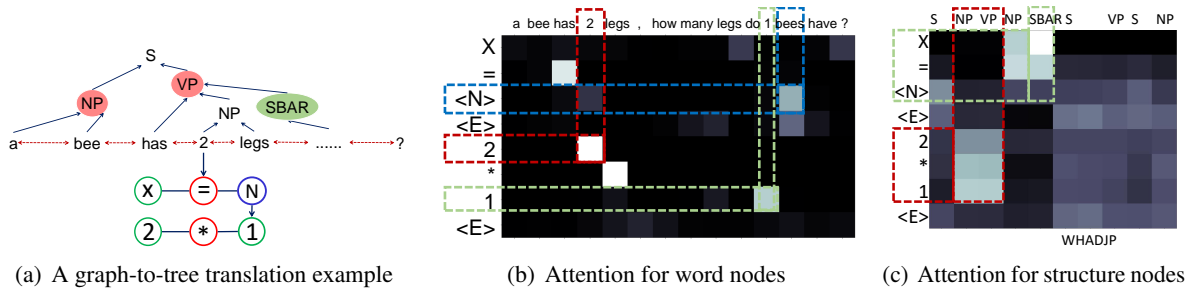


Figure 5: Effect visualization of our separated attentions on both word and structure nodes in a graph.

| Methods                      | MAWAPS |
|------------------------------|--------|
| Full model                   | 78.8   |
| w/o const tree               | 75.6   |
| w/ original GraphSage        | 76.4   |
| w/ only parent feeding       | 75.6   |
| w/ only sibling feeding      | 72.4   |
| w/o parent & sibling feeding | 67.6   |
| w/o separated attention      | 67.6   |
| w/ uniform attention         | 71.6   |
| w/o bilstm                   | 72.8   |

Table 7: Ablation study of Graph2Tree on the math word problem (MAWAPS). We employ solution accuracy as evaluation metric. The *Methods* settings is same as Table 4.

ervation in the semantic parsing task. However, it is worth noting that, according to the experiment, the sibling mechanism is obviously more important to the MWP task than the semantic parsing task, which is in line with our expectations. In the MWP task, the result of decoding, math expressions, is relatively simple compared to semantic parsing. And in math expressions, the order between leaf nodes (numbers), which directly affects the correctness of expressions, is very important. The sibling mechanism plays exactly such a role. One potential interesting extension is that, if we can connect leaf nodes in the input graph and employ edge weights to dynamically represent the order between the nodes, it may achieve a similar or even better effect than the sibling mechanism.

## 6 Conclusion and Future Work

We presented a novel Graph2Tree model consisting of a graph encoder and a hierarchical tree decoder, for learning the translation between structured inputs and structured outputs. Studies on two tasks - Semantic Parsing and Math Word Problem demonstrated our model consistently outperformed

or matched the performance of the state-of-the-art. Our Graph2Tree model is generic and agnostic to the downstream tasks and thus one of the future works is to adapt it to the other NLP applications.

## References

- Yasemin Altun, Thomas Hofmann, and Alexander J Smola. 2004. Gaussian process classification for segmenting and annotating sequences. In *ICML*, page 4. ACM.
- David Alvarez-Melis and Tommi S Jaakkola. 2017. Tree-structured decoding with doubly-recurrent neural networks. *ICLR*.
- Aida Amini, Saadia Gabriel, Shanchuan Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh Hajishirzi. 2019. [MathQA: Towards interpretable math word problem solving with operation-based formalisms](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2357–2367, Minneapolis, Minnesota. Association for Computational Linguistics.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. [Neural machine translation by jointly learning to align and translate](#). *arXiv e-prints*, abs/1409.0473.
- Joost Bastings, Ivan Titov, Wilker Aziz, Diego Marcheggiani, and Khalil Sima'an. 2017. [Graph convolutional encoders for syntax-aware neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1957–1967, Copenhagen, Denmark. Association for Computational Linguistics.
- Daniel Beck, Gholamreza Haffari, and Trevor Cohn. 2018. [Graph-to-sequence learning using gated graph neural networks](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 273–283, Melbourne, Australia. Association for Computational Linguistics.

- Bo Chen, Le Sun, and Xianpei Han. 2018a. [Sequence-to-action: End-to-end semantic graph generation for semantic parsing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 766–777, Melbourne, Australia. Association for Computational Linguistics.
- Kezhen Chen, Qiuyuan Huang, Hamid Palangi, Paul Smolensky, Kenneth D Forbus, and Jianfeng Gao. 2019a. Natural-to formal-language generation using tensor product representations. *arXiv preprint arXiv:1910.02339*.
- Xinyun Chen, Chang Liu, and Dawn Song. 2018b. Tree-to-tree neural networks for program translation. In *NIPS*, pages 2547–2557.
- Yu Chen, Lingfei Wu, and Mohammed J Zaki. 2019b. Graphflow: Exploiting conversation flow with graph neural networks for conversational machine comprehension. *arXiv preprint arXiv:1908.00059*.
- Yu Chen, Lingfei Wu, and Mohammed J Zaki. 2020. Reinforcement learning based graph-to-sequence model for natural question generation. *ICLR*.
- Deborah A. Dahl, Madeleine Bates, Michael Brown, William Fisher, Kate Hunicke-Smith, David Pallett, Christine Pao, Alexander Rudnicky, and Elizabeth Shriberg. 1994. [Expanding the scope of the atis task: The atis-3 corpus](#). In *HUMAN LANGUAGE TECHNOLOGY: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*.
- Li Dong and Mirella Lapata. 2016. [Language to logical form with neural attention](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 33–43, Berlin, Germany. Association for Computational Linguistics.
- Li Dong and Mirella Lapata. 2018. [Coarse-to-fine decoding for neural semantic parsing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 731–742, Melbourne, Australia. Association for Computational Linguistics.
- Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. 2017. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1263–1272. JMLR. org.
- Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. 2013. Speech recognition with deep recurrent neural networks. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 6645–6649. IEEE.
- Jetic Gū, Hassan S. Shavarani, and Anoop Sarkar. 2018. [Top-down tree structured decoding with syntactic connections for neural machine translation and parsing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 401–413, Brussels, Belgium. Association for Computational Linguistics.
- Xiaojie Guo, Lingfei Wu, and Liang Zhao. 2018. Deep graph translation. *arXiv preprint arXiv:1805.09980*.
- Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pages 1024–1034.
- Mohammad Javad Hosseini, Hannaneh Hajishirzi, Oren Etzioni, and Nate Kushman. 2014. [Learning to solve arithmetic word problems with verb categorization](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 523–533, Doha, Qatar. Association for Computational Linguistics.
- Robin Jia and Percy Liang. 2016. [Data recombination for neural semantic parsing](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12–22, Berlin, Germany. Association for Computational Linguistics.
- Zhanming Jie and Wei Lu. 2018. [Dependency-based hybrid trees for semantic parsing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2431–2441, Brussels, Belgium. Association for Computational Linguistics.
- Thorsten Joachims, Thomas Hofmann, Yisong Yue, and Chun-Nam Yu. 2009. Predicting structured objects with support vector machines. *Communications of the ACM*, 52(11):97.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*.
- Rik Koncel-Kedziorski, Subhro Roy, Aida Amini, Nate Kushman, and Hannaneh Hajishirzi. 2016. [MAWPS: A math word problem repository](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1152–1157, San Diego, California. Association for Computational Linguistics.
- Nate Kushman, Yoav Artzi, Luke Zettlemoyer, and Regina Barzilay. 2014. [Learning to automatically solve algebra word problems](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 271–281, Baltimore, Maryland. Association for Computational Linguistics.

- Jian Li, Zhaopeng Tu, Baosong Yang, Michael R. Lyu, and Tong Zhang. 2018. [Multi-head attention with disagreement regularization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2897–2903, Brussels, Belgium. Association for Computational Linguistics.
- Jierui Li, Lei Wang, Jipeng Zhang, Yan Wang, Bing Tian Dai, and Dongxiang Zhang. 2019. [Modeling intra-relation in math word problems with different functional multi-head attentions](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6162–6167, Florence, Italy. Association for Computational Linguistics.
- Yujia Li, Richard Zemel, Marc Brockschmidt, and Daniel Tarlow. 2016. [Gated graph sequence neural networks](#). In *Proceedings of ICLR’16*.
- Wang Ling, Phil Blunsom, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kočiský, Fumin Wang, and Andrew Senior. 2016. [Latent predictor networks for code generation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 599–609, Berlin, Germany. Association for Computational Linguistics.
- Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. 2017. [Program induction by rationale generation: Learning to solve and explain algebraic word problems](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 158–167, Vancouver, Canada. Association for Computational Linguistics.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. [Effective approaches to attention-based neural machine translation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal. Association for Computational Linguistics.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. [The Stanford CoreNLP natural language processing toolkit](#). In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60, Baltimore, Maryland. Association for Computational Linguistics.
- Xiaochang Peng, Daniel Gildea, and Giorgio Satta. 2018. [Amr parsing with cache transition systems](#). In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Maxim Rabinovich, Mitchell Stern, and Dan Klein. 2017. [Abstract syntax networks for code generation and semantic parsing](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1139–1149, Vancouver, Canada. Association for Computational Linguistics.
- Siva Reddy, Oscar Täckström, Michael Collins, Tom Kwiatkowski, Dipanjan Das, Mark Steedman, and Mirella Lapata. 2016. [Transforming dependency structures to logical forms for semantic parsing](#). *Transactions of the Association for Computational Linguistics*, 4:127–140.
- Benjamin Robaidek, Rik Koncel-Kedziorski, and Hananeh Hajishirzi. 2018. [Data-driven methods for solving algebra word problems](#). *arXiv preprint arXiv:1804.10718*.
- Lin Feng Song, Yue Zhang, Zhiguo Wang, and Daniel Gildea. 2018. [A graph-to-sequence model for AMR-to-text generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1616–1626, Melbourne, Australia. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In *NIPS*, pages 3104–3112.
- Ioannis Tsochantaridis, Thorsten Joachims, Thomas Hofmann, and Yasemin Altun. 2005. [Large margin methods for structured and interdependent output variables](#). *Journal of machine learning research*, 6(Sep):1453–1484.
- Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2017. [Graph attention networks](#). *ArXiv*, abs/1710.10903.
- Lei Wang, Yan Wang, Deng Cai, Dongxiang Zhang, and Xiaojiang Liu. 2018a. [Translating a math word problem to a expression tree](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1064–1069, Brussels, Belgium. Association for Computational Linguistics.
- Lei Wang, Dongxiang Zhang, Lianli Gao, Jingkuan Song, Long Guo, and Heng Tao Shen. 2018b. [Mathdqn: Solving arithmetic word problems via deep reinforcement learning](#). In *AAAI*.
- Lei Wang, Dongxiang Zhang, Jipeng Zhang, Xing Xu, Lianli Gao, Bing Tian Dai, and Heng Tao Shen. 2019. [Template-based math word problem solvers with recursive neural networks](#). In *AAAI*.

- Mingxuan Wang, Zhengdong Lu, Jie Zhou, and Qun Liu. 2017. [Deep neural machine translation with linear associative unit](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 136–145, Vancouver, Canada. Association for Computational Linguistics.
- Zhipeng Xie and Shichao Sun. 2019. [A goal-driven tree-structured neural model for math word problems](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5299–5305. International Joint Conferences on Artificial Intelligence Organization.
- Kun Xu, Lingfei Wu, Zhiguo Wang, and Vadim Sheinin. 2018a. [Graph2seq: Graph to sequence learning with attention-based neural networks](#). *arXiv preprint arXiv:1804.00823*.
- Kun Xu, Lingfei Wu, Zhiguo Wang, Mo Yu, Liwei Chen, and Vadim Sheinin. 2018b. [Exploiting rich syntactic information for semantic parsing with graph-to-sequence model](#). *arXiv preprint arXiv:1808.07624*.
- Kun Xu, Lingfei Wu, Zhiguo Wang, Mo Yu, Liwei Chen, and Vadim Sheinin. 2018c. [Sql-to-text generation with graph-to-sequence model](#). *arXiv preprint arXiv:1809.05255*.
- Pengcheng Yin and Graham Neubig. 2017. [A syntactic neural model for general-purpose code generation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 440–450, Vancouver, Canada. Association for Computational Linguistics.
- Pengcheng Yin, Graham Neubig, Miltiadis Allamanis, Marc Brockschmidt, and Alexander L. Gaunt. 2019. [Learning to represent edits](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.
- Pengcheng Yin, Chunting Zhou, Junxian He, and Graham Neubig. 2018. [StructVAE: Tree-structured latent variable models for semi-supervised semantic parsing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 754–765, Melbourne, Australia. Association for Computational Linguistics.
- Luke S. Zettlemoyer and Michael Collins. 2005. [Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars](#). In *Proceedings of the 21st Conference on Uncertainty in AI*, pages 658–666.
- Yanyan Zou and Wei Lu. 2019. [Text2Math: End-to-end parsing text into math expressions](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5327–5337, Hong Kong, China. Association for Computational Linguistics.