

# Conversational Question Answering over Knowledge Graphs with Transformer and Graph Attention Networks

Endri Kacupaj<sup>1</sup>, Joan Plepi<sup>2</sup>, Kuldeep Singh<sup>3</sup>, Harsh Thakkar<sup>4</sup>,  
Jens Lehmann<sup>1,5</sup>, and Maria Maleshkova<sup>1</sup>

<sup>1</sup>Smart Data Analytics Group, University of Bonn, Bonn, Germany

<sup>2</sup>Technische Universität Darmstadt, Darmstadt, Germany

<sup>3</sup>Zerotha Research and Cerence GmbH, Germany

<sup>4</sup>Zerotha Research and Osthus GmbH, Germany

<sup>5</sup>Fraunhofer IAIS, Dresden, Germany

{kacupaj, jens.lehmann, maleshkova}@cs.uni-bonn.de

joan.plepi@tu-darmstadt.de

kuldeep.singh1@cerence.com

harsh.thakkar@osthus.com

jens.lehmann@iais.fraunhofer.de

## Abstract

This paper addresses the task of (complex) conversational question answering over a knowledge graph. For this task, we propose LASAGNE (muLti-task semAntic parSing with trAnsformer and Graph atteNtion nEtworks). It is the first approach, which employs a transformer architecture extended with Graph Attention Networks for multi-task neural semantic parsing. LASAGNE uses a transformer model for generating the base logical forms, while the Graph Attention model is used to exploit correlations between (entity) types and predicates to produce node representations. LASAGNE also includes a novel entity recognition module which detects, links, and ranks all relevant entities in the question context. We evaluate LASAGNE on a standard dataset for complex sequential question answering, on which it outperforms existing baseline averages on all question types. Specifically, we show that LASAGNE improves the F1-score on eight out of ten question types; in some cases, the increase in F1-score is more than 20% compared to the state of the art.

## 1 Introduction

Since their inception in the late 2000s, publicly available Knowledge Graphs (e.g., DBpedia (Lehmann et al., 2015) and Yago (Suchanek et al., 2007)) have been widely used as a source of knowledge in several natural language processing (NLP) tasks such as entity linking, relation extraction, fact-checking, and question answering. Question answering (QA), in particular, is an essential task that maps a user natural language question

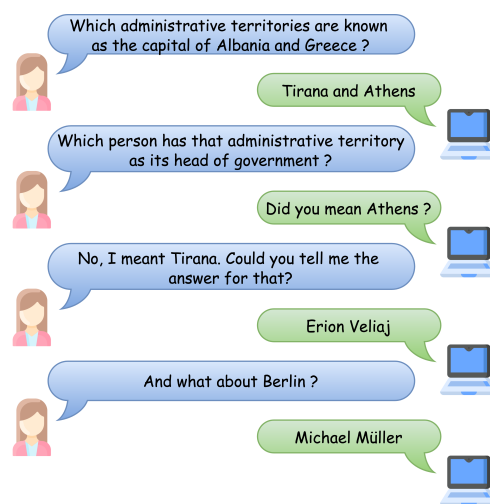


Figure 1: Conversational Question Answering task with examples similar to CSQA dataset (Saha et al., 2018).

to a query over a knowledge graph (KG) to retrieve the correct answer (Singh et al., 2018). With the increasing popularity of intelligent personal assistants (e.g., Alexa, Siri), the research focus has been shifted to conversational question answering that involves multi-turn dialogues, incorporating the phenomenon of anaphora and ellipses (Christmann et al., 2019; Shen et al., 2019)(c.f. Figure 1).

Conversational QA is often realised by using semantic parsing approaches, mapping an utterance to a logic form for extracting answers from a KG (Guo et al., 2018; Shen et al., 2019). The state of the art for semantic parsing approaches decomposes the semantic parsing process into two stages (Shen et al., 2019). First, a logical form is generated based on low-level features and then

the missing details are filled by considering both the question and the template. Other approaches (Dong and Lapata, 2016; Guo et al., 2018; Liang et al., 2017) first employ an entity linking model to identify entities in the question and subsequently use another model to map the question to a logical form. (Zhang et al., 2018; Shen et al., 2019) point out that the modular approaches suffer from the common issue of error propagation along the QA pipeline, resulting in accumulated errors. To mitigate these errors, Shen et al. (2019) proposed a multi-task framework, where a pointer-equipped semantic parsing model was designed to resolve coreference in conversations and empower joint learning with a type-aware entity detection model. Furthermore, the authors used simple classifiers to predict the required (entity) types and predicates for the generated logical forms. In this paper, we argue that Shen et al. (2019) model (the current SotA) has the following shortcomings: 1) the (entity) type and predicate classifiers share no common information, except for the supervision signal propagated to them. 2) Hence, due to missing common information, the model can produce ambiguous results, since the classifiers can predict entities and predicates that do not correlate with each other.

**Approach and Contributions:** We tackle the problem of conversational (complex) question answering over a large-scale knowledge graph. We propose LASAGNE (muLti-task semAntic parSing with trAnsformer and Graph atteNtion nEtworks) - a multi-task learning framework consisting of a transformer model extended with Graph Attention Networks (GATs) (Veličković et al., 2018) for multi-task neural semantic parsing. Our framework handles semantic parsing using the transformer (Vaswani et al., 2017) model similar to previous approaches. However, in LASAGNE we introduce the following two novel contributions: 1) the transformer model is supplemented with a Graph Attention Network to exploit the correlations between (entity) types and predicates due to its message-passing ability between the nodes. 2) We propose a novel entity recognition module that detects, links, filters, and permutes all relevant entities. (Shen et al., 2019) uses a pointer equipped decoder that learns and identifies the relevant entities for the logical form using only the encoder’s information. In contrast, we use both sources of information, i.e., the entity detection module and the encoder, to filter and permute the relevant en-

tities for a logical form. This avoids re-learning entity information in the current question context and relies on the entity detection module’s information. Our empirical results show that the proposed novel contributions lead to substantial performance improvements.

LASAGNE achieves the state of the art results in 8 out of 10 question types on the Complex Sequential Question Answering (CSQA) (Saha et al., 2018) dataset consisting of conversations over linked QA pairs. The dataset contains 200K dialogues with 1.6M turns, and over 12.8M entities from Wikidata<sup>1</sup>. Our implementation, the annotated dataset with the proposed grammar, and the results are publicly available to facilitate reproducibility and reuse<sup>2</sup>.

The structure of the paper is as follows: Section 2 summarises the related work. Section 3 presents the proposed LASANGE framework. Section 4 describes the experiments, including the experimental setup, the results, the ablation study and error analysis. We conclude in Section 5.

## 2 Related Work

We point to the survey by (Gao et al., 2018) that provides a holistic overview of neural approaches in conversational AI. In this paper, we stick to our closely related work, i.e., semantic parsing-based approaches in conversations. (Liang et al., 2017) introduce a neural symbolic machine (NSM) extended with a key-value memory network, where keys and values are the output of the sequence model in different encoding or decoding steps. The NSM model is trained using the REINFORCE algorithm with weak supervision and evaluated on the WebQuestionsSP dataset (Yih et al., 2016).

(Saha et al., 2018) propose a hybrid model of the HRED model (Serban et al., 2016) and the key-value memory network model (Miller et al., 2016). The model consists of three components. The first one is the Hierarchical Encoder, which computes a representation for each utterance. The next module is a higher-level encoder that computes a representation for the context. The second component is the Key-Value Memory Network. It stores each of the candidate tuples as a key-value pair where the key contains the concatenated embedding of the relation and the subject. In contrast, the value contains

<sup>1</sup><https://www.wikidata.org/>

<sup>2</sup><https://github.com/endrikacupaj/LASAGNE>

Action	Description
set $\rightarrow$ find( $e, p$ )	set of objects part of the triples with subject $e$ and predicate $p$
set $\rightarrow$ find_reverse( $e, p$ )	set of subjects part of the triples with object $e$ and predicate $p$
set $\rightarrow$ filter_type(set, tp)	filter the given set of entities based on the given type
set $\rightarrow$ filter_multi_types( $set_1, set_2$ )	filter the given set of entities based on the given set of types
dict $\rightarrow$ find_tuple_counts( $p, tp_1, tp_2$ )	extracts a dictionary, where keys are entities of $type_{e_1}$ and values are the number of objects of $type_{e_2}$ related with $p$
dict $\rightarrow$ find_reverse_tuple_counts( $p, tp_1, tp_2$ )	extracts a dictionary, where keys are entities of $type_{e_1}$ and values are the number of subjects of $type_{e_2}$ related with $p$
set $\rightarrow$ greater(dict, num)	set of those entities that have greater count than $num$
set $\rightarrow$ lesser(dict, num)	set of those entities that have lesser count than $num$
set $\rightarrow$ equal(dict, num)	set of those entities that have equal count with $num$
set $\rightarrow$ approx(dict, num)	set of those entities that have approximately same count with $num$
set $\rightarrow$ atmost(dict, num)	set of those entities that have at most same count with $num$
set $\rightarrow$ atleast(dict, num)	set of those entities that have at least same count with $num$
set $\rightarrow$ argmin(dict)	set of those entities that have the most count
set $\rightarrow$ argmax(dict)	set of those entities that have the least count
boolean $\rightarrow$ is_in(entity, set)	check if the entity is part of the set
number $\rightarrow$ count(set)	count the number of elements in the set
set $\rightarrow$ union( $set_1, set_2$ )	union of $set_1$ and $set_2$
set $\rightarrow$ intersection( $set_1, set_2$ )	intersection of $set_1$ and $set_2$
set $\rightarrow$ difference( $set_1, set_2$ )	difference of $set_1$ and $set_2$

Table 1: Predefined grammar with respective actions to generate logical forms.

the embedding of the object. The last component is the decoder used to create an end-to-end solution and produce multiple types of answers.

(Guo et al., 2018) present a model that converts an utterance in conversation to a logical form. The model follows a flexible grammar, in which the generation of a logical form is equivalent to predicting a sequence of actions. A dialogue memory management is proposed and integrated into the model, so that historical entities, predicates, and action sub-sequences can selectively be replicated. (Shen et al., 2019) proposed the first multi-task learning framework that learns type-aware entity detection and pointer-equipped logical form generation simultaneously. The multi-task learning framework takes advantage of the supervision from the subtasks.

### 3 LASAGNE

In a conversation, the input data consists of utterances  $u$  and their answers  $a$ , extracted from the knowledge graph. Our framework LASAGNE employs a multi-task semantic parsing approach. In particular, it maps the utterance  $u$  to a logical form  $z$ , depending on the conversation context. Figure 2 shows the architecture of LASAGNE.

#### 3.1 Grammar

For the semantic parsing task, we propose a grammar that can be used to capture the entire context of the input utterance with the minimum number of actions. Table 1 illustrates the complete grammar with all the defined actions. We considered the

work by (Guo et al., 2018) as a starting point for generating them, however, we have updated many of the semantic actions. For instance, for a couple of actions, we also define their reverse occurrence (e.g. *find*, *find\_reverse*)).

#### 3.2 Transformer

To translate the input conversation into a sequence of actions (logical form), we utilise a transformer model (Vaswani et al., 2017). Specifically, the transformer here aims to map a question  $q$ , that is a sequence  $x = \{x_1, \dots, x_n\}$ , to the answer label  $l$ , that can be also defined as a sequence  $y = \{y_1, \dots, y_m\}$ , by modelling the conditional probability  $p(y|x)$ .

##### 3.2.1 Input and Word Embedding

We have to incorporate the dialog history from previous interactions as an additional input to our model for handling coreference and ellipsis. To do so, we consider the following utterances for each turn: 1) the previous question, 2) the previous answer, and 3) the current question. Utterances are separated from one another by using a *[SEP]* token. At the end of the last utterance, we append a context token *[CTX]*, which is used as the semantic representation for the entire input question. In the next step, given an utterance  $q$  containing  $n$  words  $\{w_1, \dots, w_n\}$  we first tokenise the conversation context using WordPiece tokenization (Wu et al., 2016), and after that, we use the pre-trained model GloVe (Pennington et al., 2014) to embed the words into a vector representation space of di-

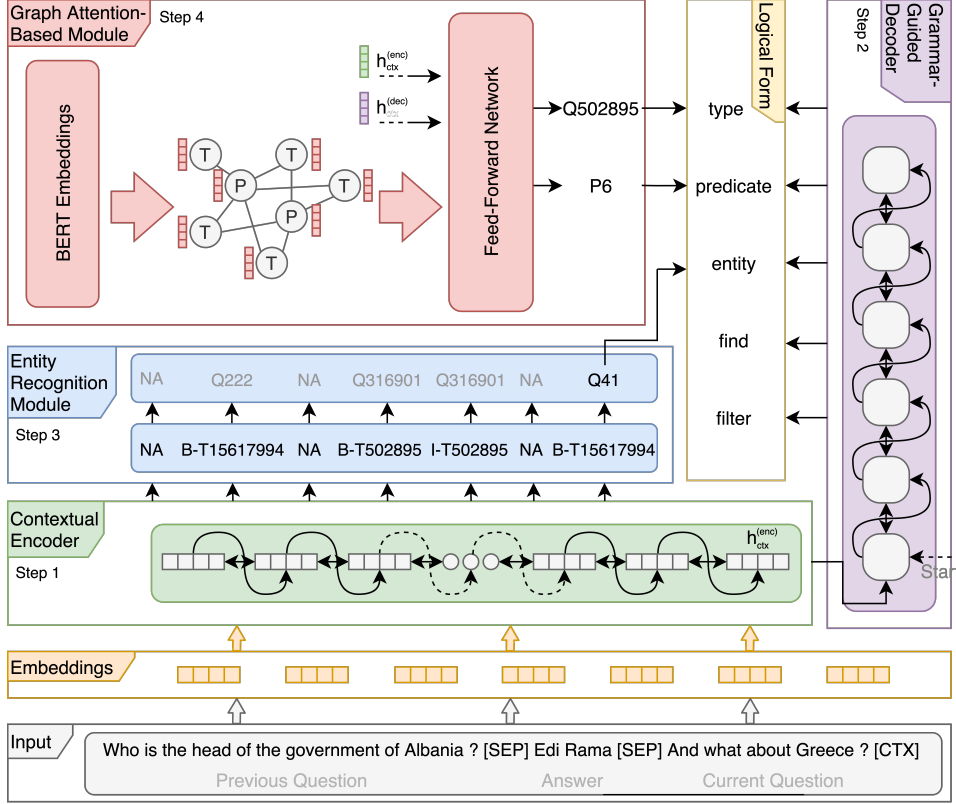


Figure 2: LASAGNE (Multi-task Semantic Parsing with Transformer and Graph Attention Networks) architecture. It consists of three modules: 1) A semantic parsing-based transformer model, containing a contextual encoder and a grammar guided decoder using the grammar defined in Table 1. 2) An entity recognition module, which identifies all the entities in the context, together with their types, linking them to the knowledge graph. It filters them based on the context and permutes them, in case of more than one required entity. Finally, 3) a graph attention-based module that uses a GAT network initialised with BERT embeddings to incorporate and exploit correlations between (entity) types and predicates. The resulting node embeddings, together with the context hidden state ( $h_{ctx}^{(enc)}$ ) and decoder hidden state ( $d_h$ ), are used to score the nodes and predict the corresponding type and predicate.

mension  $d$ <sup>3</sup>. Our word embedding model provides us with a sequence  $x = \{x_1, \dots, x_n\}$  where  $x_i$  is given by,  $x_i = GloVe(w_i)$  and  $x_i \in \mathbb{R}^d$ .

### 3.2.2 Contextual Encoder

The word embeddings  $x$ , are forwarded as input to the contextual encoder, which uses the multi-head attention mechanism described by (Vaswani et al., 2017). The encoder here outputs the contextual embeddings  $h^{(enc)} = \{h_1^{(enc)}, \dots, h_n^{(enc)}\}$ , where  $h_i^{(enc)} \in \mathbb{R}^d$  and it can be defined as:

$$h^{(enc)} = \text{encoder}(x; \theta^{(enc)}), \quad (1)$$

where  $\theta^{(enc)}$  are the encoder’s trainable parameters.

### 3.2.3 Grammar-Guided Decoder

We use a grammar guided decoder for generating the logical forms. The decoder also employs the

multi-head attention mechanism. The decoder output is dependent on the encoder contextual embeddings  $h$ . The main task of the decoder is to generate each corresponding action, based on Table 1, alongside with the general semantic object from the knowledge graph (entity, type, predicate). In other words, the decoder will predict the main logical form without using or initialising any specific information from the knowledge graph. Here we define the decoder vocabulary as  $V^{(dec)} = \{find, find\_reverse, \dots, entity, type, predicate, value\}$ , where all the actions from Table 1 are included. On top of the decoder stack, we employ a linear layer alongside a softmax to calculate each token’s probability scores in the vocabulary. We define the decoder stack output as follows:

$$\begin{aligned} h^{(dec)} &= \text{decoder}(h^{(enc)}; \theta^{(dec)}), \\ p_t^{(dec)} &= \text{softmax}(\mathbf{W}^{(dec)} h_t^{(dec)}), \end{aligned} \quad (2)$$

<sup>3</sup>Across the model, we use the same dimension  $d$  for all the representations, unless it is explicitly noted.



where  $h_t^{(dec)}$  is the hidden state in time step  $t$ ,  $\theta^{(dec)}$  are the decoder trainable parameters,  $\mathbf{W}^{(dec)} \in \mathbb{R}^{|V^{(dec)}| \times d}$  are the linear layer weights, and  $p_t^{(dec)} \in \mathbb{R}^{|V^{(dec)}|}$  is the probability distribution over the decoder vocabulary in time step  $t$ . The  $|V^{(dec)}|$  denotes the decoder’s vocabulary size.

### 3.3 Entity Recognition Module

The entity recognition module is composed of two sub-modules, where each module is trained using a different objective.

#### 3.3.1 Entity Detection and Linking

**Entity Detection** It aims to detect and link the entities to the KG. The module is inspired by (Shen et al., 2019) and performs type-aware entity detection by using BIO sequence tagging jointly with entity type tagging. Specifically, the entity detection vocabulary is defined as  $V^{(ed)} = \{O, \{B, I\} \times \{TP_i\}_{i=1}^{N^{(tp)}}\}$ , where  $TP_i$  denotes the  $i$ -th entity type label,  $N^{(tp)}$  stands for the number of the distinct entity types in the knowledge graph and  $|V^{(ed)}| = 2 \times N^{(tp)} + 1$ . For performing the sequence tagging task we use an LSTM (Hochreiter and Schmidhuber, 1997) and the module is defined as:

$$\begin{aligned} h^{(l)} &= \text{LeakyReLU}(\text{LSTM}(h^{(enc)}; \theta^{(l)})), \\ p_t^{(ed)} &= \text{softmax}(\mathbf{W}^{(l)} h_t^{(l)}), \end{aligned} \quad (3)$$

where  $h^{(enc)}$  is the encoder hidden state,  $\theta^{(l)}$  are the LSTM layer trainable parameters,  $h_t^{(l)}$  is the LSTM hidden state for time step  $t$ ,  $\mathbf{W}^{(l)} \in \mathbb{R}^{|V^{(ed)}| \times d}$  are the linear layer weights and  $p_t^{(ed)}$  are the entity detection module prediction for time step  $t$ .  $|V^{(ed)}|$  denotes the entity detection vocabulary size.

**Entity Linking** Once the entity BIO labels and their types are recognised, the next steps for the entity linking are: 1) the BIO labels are used to locate the entity spans from the input utterances. 2) An inverted index built for the knowledge graph entities is used to retrieve candidates for each predicted entity span. Finally, 3) the candidate lists are filtered using the predicted (entity) types. From the filtered candidates, the first entity is considered as correct.

#### 3.3.2 Filtering and Permutation

After finding all the input utterances’ entities, we perform two additional tasks in order to use entities in the generated logical form. First, we filter the relevant entities, and then we need to permute the

entities in the order required for the logical form. The module receives as an input the concatenation of the hidden states of the encoder  $h^{(enc)}$  and the hidden states of the LSTM  $h^{(l)}$  from the entity detection model. The module here learns to assign index tags to each input token. We define the module vocabulary as  $V^{(ef)} = \{0, 1, \dots, m\}$  where 0 is the index assigned to the context entities that are not considered. The remaining values are indices that permute our entities based on the logical form. Here,  $m$  is the total number of indices based on the maximum number of entities from all logical forms. Overall, our filtering and permutation module is modelled using a feed-forward network with two linear layers separated with a Leaky ReLU activation function and appended with a softmax. Formally we define the module as:

$$\begin{aligned} h^{(ef)} &= \text{LeakyReLU}(\mathbf{W}^{(ef_1)}[h^{(enc)}; h^{(l)}]), \\ p_t^{(ef)} &= \text{softmax}(\mathbf{W}^{(ef_2)} h_t^{(ef)}), \end{aligned} \quad (4)$$

where  $\mathbf{W}^{(ef_1)} \in \mathbb{R}^{d \times 2d}$  are the weights of the first linear layer and  $h_t^{(ef)}$  is the hidden state of the module in time step  $t$ .  $\mathbf{W}^{(ef_2)} \in \mathbb{R}^{|V^{(ef)}| \times d}$  are the weights of the second linear layer,  $|V^{(ef)}|$  is the size of the vocabulary and  $p_t^{(ef)}$  denotes the probability distribution over the tag indices for the time step  $t$ .

### 3.4 Graph Attention-Based Module

A knowledge graph (KG) can be denoted as a set of triples  $\mathcal{K} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$  where  $\mathcal{E}$  and  $\mathcal{R}$  are the set of entities and relations respectively. To build the (local) graph, we consider the relations and the types of entities that are linked with these relations in the knowledge graph  $\mathcal{K}$ . We define a graph  $\mathcal{G} = \{\mathcal{T} \cup \mathcal{R}, \mathcal{L}\}$  where  $\mathcal{T}$  is the set of types,  $\mathcal{R}$  is the set of relations and  $\mathcal{L}$  is a set of links  $(tp_1, r)$  and  $(r, tp_2)$  such that  $\exists (e_1, r, e_2) \in \mathcal{K}$  where  $e_1$  is of type  $tp_1$  and  $e_2$  is of type  $tp_2$ .

To propagate information in the graph and to project prior KG information into the embedding space, we use the Graph Attention Networks (GATs) (Veličković et al., 2018).

We initialise each node embedding  $h^{(g)} = \{h_1^{(g)}, \dots, h_n^{(g)}\}$  using pretrained BERT embeddings, and  $n = |\mathcal{T} \cup \mathcal{R}|$ . A GAT layer uses a parameter weight matrix, and self-attention, to produce a transformation of input representations  $\bar{h}^{(g)} = \{\bar{h}_1^{(g)}, \dots, \bar{h}_n^{(g)}\}$ , where  $\bar{h}_i^{(g)} \in \mathbb{R}^d$  as

shown below:

$$\bar{h}^{(g)} = g(h^{(g)}; \theta^{(g)}), \quad (5)$$

and  $\theta^{(g)}$  are the trainable parameters. We model the task of predicting the correct type or predicate in the logical form as a classification task over the nodes in graph  $\mathcal{G}$ , given the current conversational context and decoder hidden state. For each time step  $t$  in the decoder, we calculate the probability distribution  $p_t^{(g)}$  over the graph nodes as:

$$p_t^{(g)} = \text{softmax}(\bar{h}^{(g)T} h_t^{(c)}), \quad (6)$$

where  $\bar{h}^{(g)} \in \mathbb{R}^{d \times n}$  and  $h_t^{(c)}$  is a linear projection of the concatenation of the context representation and the decoder hidden state, given as follows,

$$h_t^{(c)} = \text{LeakyReLU}(\mathbf{W}^{(g)}[h_{ctx}^{(enc)}; h_t^{(dec)}]), \quad (7)$$

and  $\mathbf{W}^{(g)} \in \mathbb{R}^{d \times 2d}$ .

### 3.5 Learning

The framework consists of four trainable modules, grammar guided decoder, entity detection, filtering and permutation, and the GAT-based module for types and predicates. Every module consists of a loss function that contributes to the overall performance of the framework, as shown in Section 4.3. To account for multi-tasking, we perform a weighted average of all the single losses:

$$L = \lambda_1 L^{dec} + \lambda_2 L^{ed} + \lambda_3 L^{ef} + \lambda_4 L^g, \quad (8)$$

where  $\lambda_1, \lambda_2, \lambda_3, \lambda_4$  are the relative weights, which are learned during training by taking into account the difference in magnitude between losses by incorporating the log standard deviation (Armitage et al., 2020; Cipolla et al., 2018).  $L^{dec}, L^{ed}, L^{ef}$ , and  $L^g$  are the respective negative log-likelihood losses of the grammar guided decoder, entity detection, filtering and permutation, and GAT-based modules. These losses are defined as follows:

$$\begin{aligned} L^{dec} &= - \sum_{k=1}^m \text{logp}(y_k^{(dec)} | x), \\ L^{ed} &= - \sum_{j=1}^n \text{logp}(y_j^{(ed)} | x), \\ L^{ef} &= - \sum_{i=1}^n \text{logp}(y_i^{(ef)} | x), \\ L^g &= - \sum_{k=1}^m I_{(y_k^{(dec)} \in \{type, pred\})} \text{logp}(y_k^{(g)} | x), \end{aligned} \quad (9)$$

<sup>4</sup>For more details about GAT please refer to the appendix.

where  $n$  and  $m$  are the length of the input utterance  $x$  and the gold logical form, respectively.  $y_k^{(dec)} \in V^{(dec)}$  are the gold labels for the decoder,  $y_j^{(ed)} \in V^{(ed)}$  are the gold labels for entity detection,  $y_i^{(ef)} \in V^{(ef)}$  are the gold labels for filtering and permutation, and  $y_k^{(g)} \in \{\mathcal{T} \cup \mathcal{R}\}$  are the gold labels for the GAT-based module. The model benefits from multiple supervision signals from each module, and this improves the performance in the given task.

## 4 Experiments

### 4.1 Experimental Setup

**Datasets** We use the Complex Sequential Question Answering (CSQA) dataset<sup>5</sup> (Saha et al., 2018). CSQA was built on the large-scale knowledge graph Wikidata. Wikidata consists of 21.2M triples with over 12.8M entities, 3,054 entity types, and 567 predicates. The CSQA dataset consists of around 200K dialogues where each partition – train, valid, test contains 153K, 16K, 28K dialogues, respectively. The questions involve complex reasoning to determine the correct answers.

**Model Configurations** We incorporate a semi-automated preprocessing step to annotate the CSQA dataset with gold logical forms. For each question type and subtype in the dataset, we create a general template with a pattern sequence that the actions should follow. Thereafter, we follow a set of rules to create the specific gold logical form that extracts the gold sequence of actions based on the type of question for each question. The actions used for this process are the ones in Table 1. For all the modules in the LASAGNE framework, we employ an embedding dimension of 300. We utilise the transformer model with six heads for the multi-head attention model with two layers. For the optimisation, we use the Noam optimiser proposed by (Vaswani et al., 2017), where authors use an Adam optimiser (Kingma and Ba, 2015) with several warmup steps for the learning rate. Please refer to the appendix submitted with the paper for more details.

**Models for Comparison** We compare the LASAGNE framework with the last three baselines that have been evaluated on the employed dataset. The first baseline is (Saha et al., 2018) where authors introduce the HRED+KVmem model. The second baseline is D2A (Guo et al., 2018), which

<sup>5</sup><https://amritasaha1812.github.io/CSQA>

Methods		HRED-KVM	D2A	MaSP	LASAGNE (ours)	$\Delta$
Question Type	#Examples	F1 Score				
Overall	206k	9.39%	66.70%	79.26%	<b>82.91%</b>	+3.65%
Clarification	12k	16.35%	35.53%	<b>80.79%</b>	69.46%	-11.33%
Comparative Reasoning (All)	15k	2.96%	48.85%	68.90%	<b>69.77%</b>	+0.87%
Logical Reasoning (All)	22k	8.33%	67.31%	69.04%	<b>89.83%</b>	+20.79%
Quantitative Reasoning (All)	9k	0.96%	56.41%	73.75%	<b>86.67%</b>	+12.92%
Simple Question (Coreferenced)	55k	7.26%	57.69%	76.47%	<b>79.06%</b>	+2.59%
Simple Question (Direct)	82k	13.64%	78.42%	85.18%	<b>87.95%</b>	+2.77%
Simple Question (Ellipsis)	10k	9.95%	81.14%	<b>83.73%</b>	80.09%	-3.64%
Question Type	#Examples	Accuracy				
Overall	66k	14.95%	37.33%	45.56%	<b>64.34%</b>	+18.78%
Verification (Boolean)	27k	21.04%	45.05%	60.63%	<b>78.86%</b>	+18.23%
Quantitative Reasoning (Count)	24k	12.13%	40.94%	43.39%	<b>55.18%</b>	+11.79%
Comparative Reasoning (Count)	15k	8.67%	17.78%	22.26%	<b>53.34%</b>	+31.08%

Table 2: LASAGNE’s performance comparison on the CSQA dataset having 200K dialogues with 1.6M turns and over 12.8M entities. LASAGNE achieves “overall” (weighted average on all question types) new state of the art for both the F1 score and the question type results’ accuracy metric.

uses a semantic parsing approach based on a seq2seq model. Finally, the current state of the art is MaSP (Shen et al., 2019), which is also a semantic parsing approach.

**Evaluation Metrics** We use the same metrics as employed by the authors of the CSQA dataset (Saha et al., 2018) as well as the previous baselines. The “F1-score” is used for questions that have an answer composed of a set of entities. The “Accuracy” metric is used for the question types whose answer is a number or a boolean value (YES/NO). We also provide an overall score for each evaluation metric and their corresponding question categories.

## 4.2 Results

Table 2 summarises the results comparing the LASAGNE framework against the previous baselines. LASAGNE outperforms the previous baselines weighted average on all question types (The row “overall” in the Table 2). Furthermore, LASAGNE is a new SotA in 8 out of 10 question types, and in some cases, the improvement is up to 31 percent.

**What worked in our case?** For question types that require more than two entities for reasoning, such as *Logical Reasoning (All)* and *Verification (Boolean)*, LASAGNE performs considerably better (+20.79% and +18.23% respectively). This is mainly due to the proposed entity recognition module. Furthermore, for question types that require two or more (entity) types and predicates, such as *Quantitative Reasoning (All)*, *Quantitative Reasoning (Count)* and *Comparative Reason-*

*ing (Count)* LASAGNE also outperforms MaSP (+12.92%, +11.79% and +31.08% respectively). Here, the improvement is due to the graph attention-based module, which is responsible for predicting the relevant (entity) types and predicates. Another interesting result is that LASAGNE also performs better in two out of three *Simple Question involving one entity and one predicate* categories. The performance shows the robustness of LASAGNE. **What did not work in our case?** LASAGNE noticeably under-performs on the *Clarification* question type, where MaSP retains the state-of-the-art. The main reason is the spurious logical forms during the annotation process which has further impacted the Simple Questions (Ellipses) performance.

## 4.3 Ablation Study

**Effect of GAT and Multi-task Learning** Table 3 summarises the effectiveness of the GAT-based module and the multi-task learning. We can observe the advantage of using them together in LASAGNE. To show the effectiveness of GAT-based module, we replace it with two simple classifiers, one for each predicate and type categories. We can observe that the performance drops significantly for the question types that require multiple entity types and predicates (e.g. *Quantitative Reasoning (All)*, *Quantitative Reasoning (Count)* and *Comparative Reasoning (Count)*). When we exclude the multi-task learning and train all the modules independently, there is a negative impact on all question types. In LASAGNE, the filtering and

Methods	Ours	w/o GATs	w/o Multi
<b>Question Type</b>	<b>F1 Score</b>		
Clarification	66.94%	57.33%	59.43%
Comparative	69.77%	57.72%	66.41%
Logical	89.83%	78.52%	86.75%
Quantitative	86.67%	75.26%	82.18%
Simple (Coref)	79.06%	76.46%	77.23%
Simple (Direct)	87.95%	83.59%	85.39%
Simple (Ellipsis)	80.09%	77.19%	78.47%
<b>Question Type</b>	<b>Accuracy</b>		
Verification	78.86%	63.38%	75.24%
Quantitative	55.18%	40.87%	46.27%
Comparative	53.34%	41.73%	45.90%

Table 3: The effectiveness of the GAT and the multi-task learning. The first column contains the results of the LASAGNE framework, where all the modules are trained simultaneously. The second and third columns selectively remove the GAT and the multi-task learning from LASAGNE.

permutation module, along with the GAT-based module, is heavily dependent on the supervision signals received from the previous modules. Therefore it is expected that without the multi-task learning, LASAGNE will underperform on all question types, since each module has to re-learn inherited information.

#### 4.4 Task Analysis

Tasks	Accuracy
Entity Detection	86.75%
Filtering & Permutation	97.49%
Grammar-Guided Decoder for Logical Forms	98.61%
GAT-Based Module for Type/Predicate	92.28%

Table 4: Tasks accuracy of the LASAGNE framework.

Table 4 illustrates the task accuracy of LASAGNE. The Entity Detection task has the lowest accuracy (86.75%). The main reason here is the errors in the entity type prediction. On the other hand, for all other tasks, we have accuracy above 90%.

**Effect of Filtering and Permutation** For justifying the effectiveness and superior performance of LASAGNE’s filtering and permutation module, we compare the overall performance of the entity recognition module to the corresponding module from MaSP. Please note, entity detection modules in both frameworks adopt a similar approach as defined in section 3.3. In Table 5 we can see that the MaSP entity recognition module provides an overall accuracy of 79.8% on test data, while our mod-

ule outperforms it with an accuracy of 92.1%. The main reason for the under-performance of MaSP is that it uses only token embeddings without any entity information. In contrast, our approach avoids re-learning entity information in the question context and relies on the entity detection module’s information.

Model	Entity Recognition Accuracy
MaSP	79.8%
LASAGNE	92.1%

Table 5: Comparing MaSP (Shen et al., 2019) and LASAGNE for entity recognition performance.

#### 4.5 Error Analysis

For the error analysis, we randomly sampled 100 incorrect predictions. We detail the reasons for two types of errors observed in the analysis:

**Entity Ambiguity** Even though our entity detection module assigns (entity) types to each predicted span, entity ambiguity remains the biggest challenge for our framework. For instance, for the question, “Who is associated with Jeff Smith?” LASAGNE entity detection module correctly identifies “Jeff Smith” as an entity surface form and correctly assigns the (entity) type “common name”. However, the Wikidata knowledge graph contains more than ten entities with exactly the same label and type. Our entity linking module has difficulties in such cases. Wikidata entity linking is a newly emerging research domain that has its specific challenges such as entities sharing the same labels, user-created non-standard entity labels and multi-word entity labels (up to 62 words) (Mulang et al., 2020b). Additional entity contexts, such as entity descriptions and other KG contexts, could help resolve the Wikidata entity ambiguity (Mulang et al., 2020a).

**Spurious Logical Form** For specific question categories, we could not identify gold actions for all utterances. Therefore spurious logical form is a standard error that affects LASAGNE. Specifically, we have spurious logical forms for categories such as “Comparative, Quantitative, and Clarification” but still can achieve SotA in the comparative and quantitative categories.

## 5 Conclusions

In this article, we focus on complex question answering over a large-scale knowledge graph containing conversational context. We provide a



transformer-based framework to handle the task in a multi-task semantic parsing manner. At the same time, we propose a named entity recognition module for entity detection, filtering, and permutation. Furthermore, we also introduce a graph attention-based module, which exploits correlations between (entity) types and predicates for identifying the gold ones for each particular context. We empirically show that our model achieves the best results for numerous question types and also overall. Our ablation study demonstrates the effectiveness of the multi-task learning and of our graph-based module. We also present an error analysis on a random sample of “wrong examples” to discuss our model’s weaknesses. For future work, we believe that reinforcement learning is a viable alternative to explore complex conversational question answering without gold annotations.

## Acknowledgments

The project leading to this publication has received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 812997 (Cleopatra).

## References

- Jason Armitage, Endri Kacupaj, Golsa Tahmasebzadeh, Maria Maleshkova, Ralph Ewerth, Jens Lehmann, et al. 2020. Mlm: A benchmark dataset for multitask learning with multiple languages and modalities. *arXiv preprint arXiv:2008.06376*.
- Philipp Christmann, Rishiraj Saha Roy, Abdalghani Abujabal, Jyotsna Singh, and Gerhard Weikum. 2019. Look before you hop: Conversational question answering over knowledge graphs using judicious context expansion. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 729–738.
- R. Cipolla, Y. Gal, and A. Kendall. 2018. [Multi-task learning using uncertainty to weigh losses for scene geometry and semantics](#). In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7482–7491.
- 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.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2018. Neural approaches to conversational ai. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1371–1374.
- Daya Guo, Duyu Tang, Nan Duan, Ming Zhou, and Jian Yin. 2018. [Dialog-to-action: Conversational question answering over a large-scale knowledge base](#). In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pages 2946–2955.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Comput.*, 9(8):1735–1780.
- Diederik P Kingma and Jimmy Ba. 2015. Adam, a method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, volume 1412.
- Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N. Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef, Sören Auer, and Christian Bizer. 2015. [Dbpedia - a large-scale, multilingual knowledge base extracted from wikipedia](#). *Semantic Web*, 6(2):167–195.
- Chen Liang, Jonathan Berant, Quoc V. Le, Kenneth D. Forbus, and Ni Lao. 2017. [Neural symbolic machines: Learning semantic parsers on freebase with weak supervision](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 23–33. Association for Computational Linguistics.
- Wei Lu, Hwee Tou Ng, Wee Sun Lee, and Luke S. Zettlemoyer. 2008. [A generative model for parsing natural language to meaning representations](#). In *2008 Conference on Empirical Methods in Natural Language Processing, EMNLP 2008, Proceedings of the Conference, 25-27 October 2008, Honolulu, Hawaii, USA, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 783–792. ACL.
- Alexander H. Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. 2016. [Key-value memory networks for directly reading documents](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 1400–1409. The Association for Computational Linguistics.
- Isaiah Onando Mulang, Kuldeep Singh, Chaitali Prabhu, Abhishek Nadgeri, Johannes Hoffart, and Jens Lehmann. 2020a. Evaluating the impact of knowledge graph context on entity disambiguation models. *29th ACM Conference on Information and Knowledge Management (CIKM)*.
- Isaiah Onando Mulang, Kuldeep Singh, Akhilesh Vyas, Saeedeh Shekarpour, Ahmad Sakor, Maria Esther

- Vidal, Soren Auer, and Jens Lehmann. 2020b. Encoding knowledge graph entity aliases in an attentive neural networks for wikidata entity linking. In *Proceedings of 21st Conference on Web Information System and Engineering*.
- 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.
- Amrita Saha, Vardaan Pahuja, Mitesh M. Khapra, Karthik Sankaranarayanan, and Sarath Chandar. 2018. **Complex sequential question answering: Towards learning to converse over linked question answer pairs with a knowledge graph**. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 705–713. AAAI Press.
- Iulian Vlad Serban, Alessandro Sordani, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2016. **Building end-to-end dialogue systems using generative hierarchical neural network models**. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 3776–3784. AAAI Press.
- Tao Shen, Xiubo Geng, Tao Qin, Daya Guo, Duyu Tang, Nan Duan, Guodong Long, and Daxin Jiang. 2019. **Multi-task learning for conversational question answering over a large-scale knowledge base**. 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 2442–2451, Hong Kong, China. Association for Computational Linguistics.
- Kuldeep Singh, Arun Sethupat Radhakrishna, Andreas Both, Saeedeh Shekarpour, Ioanna Lytra, Ricardo Usbeck, Akhilesh Vyas, Akmal Khikmatullaev, Dharmen Punjani, Christoph Lange, et al. 2018. **Why reinvent the wheel: Let’s build question answering systems together**. In *Proceedings of the 2018 World Wide Web Conference*, pages 1247–1256.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. **Dropout: a simple way to prevent neural networks from overfitting**. *The journal of machine learning research*, 15(1):1929–1958.
- Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. **Yago: a core of semantic knowledge**. In *Proceedings of the 16th international conference on World Wide Web*, pages 697–706.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. **Attention is all you need**. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5998–6008.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. **Graph Attention Networks**. *International Conference on Learning Representations*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. **Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation**. *arXiv e-prints*, page arXiv:1609.08144.
- Wen-tau Yih, Matthew Richardson, Christopher Meek, Ming-Wei Chang, and Jina Suh. 2016. **The value of semantic parse labeling for knowledge base question answering**. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 2: Short Papers*. The Association for Computer Linguistics.
- Yuyu Zhang, Hanjun Dai, Zornitsa Kozareva, Alexander J. Smola, and Le Song. 2018. **Variational reasoning for question answering with knowledge graph**. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 6069–6076. AAAI Press.

## A Grammar

We propose a new grammar to annotate the dataset with a gold logical form to perform the semantic parsing task. We consider the work by (Guo et al., 2018) as a starting point for generating them. While we differ in many actions regarding their semantic and therefore their implementation. Our goal was to define more precise and richer actions, which gives us a more flexible grammar in terms of being used to annotate a wider range of question’s complexities. For instance, for a couple of actions, we also define their reverse occurrence (e.g. *find*, *find\_reverse*). We do this in order to match the knowledge graph triple direction

(*subject-predicate-object*). In some questions, we might have the subject or the object entity. Having both normal and reverse actions helps us to identify directly the correct answer based on the action the model predicted. Furthermore, we also define actions that do not exist in (Guo et al., 2018). Some of them are *find\_tuple\_counts*, *atmost*, *atleast*. Table 1 illustrates the complete grammar with all the defined actions. Following (Lu et al., 2008), we define each action with a function that can be executed on the knowledge graph. Finally, in order to execute a sequence of actions, we have to parse it into a tree structure. There our executor starts from the tree leaves and it recursively executes the leftmost non-terminal node until the whole tree is complete.

## B Case Study

Table 9 shows examples from different question types in the CSQA dataset and the logical forms generated from our model. As we can see, our actions can cover reasoning for every question type by following certain patterns depending on them. The sequences can cover all the different complexities of the questions. For example, the logical forms pattern of *Simple Questions* to *Quantitative* or *Comparative* is slightly different due to increased complexity of the latter questions. The reasoning over *Quantitative* or *Comparative* question involves more actions in order to reach the correct answer.

For the question type *Simple Question (Direct)*, we can see the question “Which administrative territory is the birthplace of Antonio Reguero?”. The correct logical form for this example is “filter\_type(find(Antonio Reguero, place of birth), administrative territorial entity)”. Here we can distinguish two different actions; the first one is the *filter\_type* and the other one is the *find* action. The *find* action receives as input an entity subject and a predicate and provides the set of object entities from the Knowledge Graph. Whereas, the *filter\_type* action receives as input a set of entities along with an entity type and results to a set of entities that belong to that particular entity type.

## C Hyperparameters and module configurations

Table 6 summarizes the hyperparameters used across the LASAGNE framework. For the transformer module, we use the configurations from

Hyperparameters	Value
epochs	20
batch size	64
dropout ratio	0.1
learning rate	0.001
warmup steps	4000
optimizer	Adam
$\beta_1$	0.9
$\beta_2$	0.999
$\varepsilon$	1e-09
model dimension	300
model pretrained embeddings	GloVe
non-linear activation	LeakyRelu
GAT input dimension	3072
GAT node dimension	300
GAT pretrained embeddings	BERT

Table 6: Hyper-parameters for LASAGNE framework.

(Vaswani et al., 2017). Our model dimension is  $d_{model} = 300$ , with a total number of  $H = 6$  heads and  $L = 2$  layers. The inner feed-forward linear layers have dimension  $d_{ff} = 600$ . Following the base transformer parameters, we apply residual dropout (Srivastava et al., 2014) to the summation of the embeddings and the positional encodings in both encoder and decoder stacks with a rate of 0.1. The entity detection module has a dimension of 300. Our base LSTM here is followed with a LeakyReLU, dropout, and a linear layer. The output of the linear layer is the module prediction while the LSTM hidden state is propagated to the filtering and permutation layer. The filtering and permutation module receives an input of dimension 600 where here a linear layer is responsible to reduce it to 300 which is the framework dimension. Like in the previous module, a LeakyReLU, dropout, and a linear layer are used for the final predictions. Finally, for the GAT-based module, we use pre-trained BERT embeddings for type and predicate labels. Hence the input dimension on this module is 3072. The GAT layer will produce representations with an embedding size of 300. Next, multiple linear, dropout, and LeakyReLU layers are used to produce the final predictions.

## D Graph Attention Networks

Figure 3 shows the aggregation process of graph attention layer between the (entity) types and predicates from Wikidata. The KB types and predicates are the nodes of the graph, and there exist an edge only between types and predicates with the condition that there exist a triple which involved the predicate and an entity of that type. We use GATs (Veličković et al., 2018) to capture dif-

Methods		HRED-KVM		D2A		MaSP		LASAGNE (ours)	
Question Type	#Examples	Precision	Recall	Precision	Recall	Precision	Recall	Precision	Recall
Overall	206k	6.30%	18.40%	66.57%	66.83%	80.48%	78.07%	87.08%	80.31%
Clarification	12k	12.13%	25.09%	33.97%	37.24%	77.66%	84.18%	81.80%	60.35%
Comparative Reasoning (All)	15k	4.97%	2.11%	54.68%	44.14%	81.20%	59.83%	83.88%	59.73%
Logical Reasoning (All)	22k	5.75%	15.11%	68.86%	65.82%	78.00%	61.92%	98.67%	82.43%
Quantitative Reasoning (All)	9k	1.01%	0.91%	60.63%	52.74%	79.02%	69.14%	79.66%	95.02%
Simple Question (Coreferenced)	55k	5.09%	12.67%	56.94%	58.47%	76.01%	76.94%	72.71%	86.62%
Simple Question (Direct)	82k	8.58%	33.30%	77.37%	79.50%	84.29%	86.09%	94.94%	81.92%
Simple Question (Ellipsis)	10k	6.98%	17.30%	77.90%	84.67%	82.03%	85.50%	93.74%	69.91%
Question Type	#Examples	Accuracy							
Overall	66k	14.95%		37.33%		45.56%		64.34%	
Verification (Boolean)	27k	21.04%		45.05%		60.63%		78.86%	
Quantitative Reasoning (Count)	24k	12.13%		40.94%		43.39%		55.18%	
Comparative Reasoning (Count)	15k	8.67%		17.78%		22.26%		53.34%	

Table 7: Precision and recall comparison with baselines.

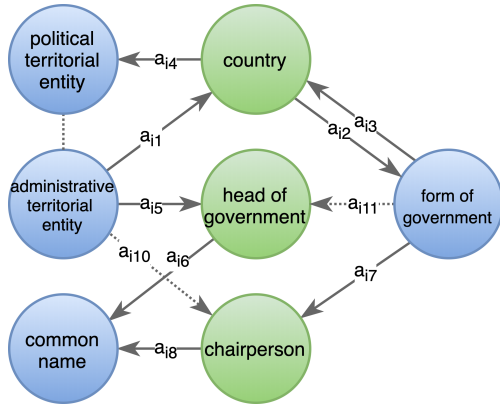


Figure 3: The aggregation process of graph attention layer between the (entity) types and predicates from Wikidata knowledge graph. The dashed lines represent an auxiliary edge, while  $a_{ij}$  represents relative attention values of the edge. We also incorporate the predicates (relations) as nodes of the graph instead of edges.

ferent level of information for a node, based on the neighborhood in the graph. We denote with  $h^{(g)} = \{h_1^{(g)}, \dots, h_n^{(g)}\}$  the initial representations of the nodes, which will also be the input features for the GAT layer. To denote the influence of node  $j$  to the node  $i$ , an attention score  $e_{ij}$  is computed as  $e_{ij} = a(\mathbf{W}h_i^{(g)}, \mathbf{W}h_j^{(g)})$ , where  $\mathbf{W}$  is a parameterized linear transformation, and  $a$  is an attention function. In our case, we follow the GAT paper, and compute  $e_{ij}$  score as follows,

$$e_{ij} = \text{LeakyReLU}(\mathbf{a}^T[\mathbf{W}h_i^{(g)} \parallel \mathbf{W}h_j^{(g)}]), \quad (10)$$

where  $\mathbf{a} \in \mathbb{R}^{2d}$  is a single-layer feedforward network, and  $\parallel$  denotes concatenation. This attention scores are normalized using a softmax function and producing the  $\alpha_{ij}$  scores for all the edges in a neighborhood. These normalized attention scores

are used to compute the output features  $\bar{h}_i^{(g)}$  of a node in a graph, by applying a linear combination of all the nodes in the neighborhood as below,

$$\bar{h}_i^{(g)} = \sigma\left(\sum_{j \in \mathcal{N}_i} \alpha_{ij} \mathbf{W}h_j^{(g)}\right) \quad (11)$$

where  $\sigma$  is a non-linear function. Following (Veličković et al., 2018) and (Vaswani et al., 2017) we also apply a multi-head attention mechanism and compute the final output features as,

$$\bar{h}_i^{(g)} = \sigma\left(\frac{1}{K} \sum_{k=1}^K \sum_{j \in \mathcal{N}_i} \alpha_{ij}^k \mathbf{W}^k h_j^{(g)}\right) \quad (12)$$

where  $K$  is equal to the number of heads, and  $\alpha_{ij}^k$ ,  $\mathbf{W}^k$  are the corresponding attention scores and linear transformation by the  $k$ -th attention mechanism. During our experiments, we found out the  $K = 2$  was sufficient for our model.

## E Experiments

Table 7 summarizes precision and recall results comparing LASAGNE framework against the previous baselines. Furthermore, a detailed task analysis for each task on each question type is illustrated on Table 8.



Tasks	Entity Detection	Filt. & Permut.	Logical Form	Type/Predicate
Question Type	Accuracy			
Clarification	92.19%	99.97%	98.36%	86.70%
Comparative Reasoning (All)	92.03%	99.88%	99.00%	97.18%
Logical Reasoning (All)	72.20%	99.44%	98.18%	95.95%
Quantitative Reasoning (All)	87.38%	100.0%	99.56%	95.87%
Simple Question (Coreferenced)	93.50%	96.92%	98.50%	90.12%
Simple Question (Direct)	90.58%	99.34%	98.58%	89.71%
Simple Question (Ellipsis)	77.90%	99.98%	98.81%	90.02%
Verification (Boolean)	79.50%	84.66%	99.79%	98.10%
Quantitative Reasoning (Count)	77.77%	99.80%	97.34%	92.09%
Comparative Reasoning (Count)	92.04%	99.98%	98.66%	96.92%

Table 8: Task accuracy from LASAGNE. We can obtain that entity detection is the task with lowest accuracy while filtering and permutation together with logical form generation are the tasks with highest accuracy.

Question Type	Question	Logical Forms
Simple (Direct)	Q1: Which administrative territory is the birthplace of Antonio Reguero ?	filter_type( find(Antonio Reguero, place of birth), administrative territorial entity)
Simple (Ellipsis)	Q1: Which administrative territories are twin towns of Madrid ? A1: Prague, Moscow, Budapest Q2: And what about Urban Community of Brest?	filter_type( find(Urban Community of Brest, twinned administrative body), administrative territorial entity)
Simple (Coref)	Q1: What was the sport that Marie Pyko was a part of ? A1: Association football Q2: Which political territory does that person belong to ?	filter_type( find(Marie Pyko, country of citizenship), political territorial entity)
Quantitative Reasoning (Count)	Q1: How many beauty contests and business enterprises are located at that city ? A1: Did you mean Caracas? Q2: Yes	count(union( filter_type(find_reverse(Caracas, located in), beauty contest), filter_type(find_reverse(Caracas, located in), business enterprises) ))
Quantitative Reasoning (All)	Q1: Which political territories are known to have diplomatic connections with max number of political territories ?	argmax( find_tuple_counts(diplomatic relation, political territorial entity, political territorial entity))
Comparative Reasoning (Count)	Q1: How many alphabets are used as the scripts for more number of languages than Jawi alphabet ?	count(greater(count( filter_type(find(Jawi alphabet, writing system), language)), find_tuple_counts(writing system, alphabet, language)))
Comparative Reasoning (All)	Q1: Which occupations were more number of publications and works mainly about than composer ?	greater(union( find_reverse_tuple_counts(main subject, occupation, publication), find_reverse_tuple_counts(main subject, occupation, work)), count(filter_multi_types(find_reverse(composer, main subject), publication, work)))
Verification	Q1: Was Geir Rasmussen born at that administrative territory ?	is_in(find(Geir Rasmussen, place of birth), Chicago)

Table 9: Examples from the CSQA dataset (Saha et al., 2018) annotated with gold logical form.