

Logic2Text: High-Fidelity Natural Language Generation from Logical Forms

Zhiyu Chen¹, Wenhua Chen¹, Hanwen Zha¹, Xiyu Zhou¹, Yunkai Zhang¹,
Sairam Sundaresan², and William Yang Wang¹

¹University of California, Santa Barbara

²Intel AI

{zhiyuchen, wenhuchen, hwzha, xiyu, yunkai_zhang, william}@cs.ucsb.edu,
sairam.sundaresan@intel.com

Abstract

Previous studies on Natural Language Generation (NLG) from structured data have primarily focused on surface-level descriptions of record sequences. However, for complex structured data, e.g., multi-row tables, it is often desirable for an NLG system to describe interesting facts from logical inferences across records. If only provided with the table, it is hard for existing models to produce controllable and high-fidelity logical generations. In this work, we formulate high-fidelity NLG as generation from logical forms in order to obtain controllable and faithful generations. We present a new large-scale dataset, LOGIC2TEXT, with 10,753 descriptions involving common logic types paired with the underlying logical forms. The logical forms show diversified graph structure of free schema, which pose great challenges on the model’s ability to understand the semantics. We experiment on (1) Fully-supervised training with the full datasets, and (2) Few-shot setting, provided with hundreds of paired examples; We compare several popular generation models and analyze their performances. We hope our dataset can encourage research towards building an advanced NLG system capable of natural, faithful, and human-like generation. The dataset and code is available at <https://github.com/czyssrs/Logic2Text>.

1 Introduction

Natural language generation (NLG) from structured data has been an important research problem in many applications. Recent data-driven methods have achieved good performances on various NLG tasks (Liu et al., 2018; Freitag and Roy, 2018; Chen et al., 2019b). However most studies focus on surface descriptions of simple record sequences, for example, attribute-value pairs of fixed or very limited schema, like E2E (Novikova et al., 2017) and

WikiBio (Lebret et al., 2016). In real-world cases for multi-row tables, it is often more desirable and plausible to provide descriptions involving higher-level logical inference across data records. For example, in Figure 1, instead of plain restatements, human readers would be more favorable to abstract descriptions that can summarize or conclude information over the table records. To produce such logical-level generations of high fidelity, it is not yet appropriate to provide only the table as the input in a real-world NLG system, based on the following reasons:

1) *Low Fidelity*. Given only the table, it is challenging for existing neural models to produce such logically correct generations involving reasoning and symbolic calculations, e.g., max, min, counting, averaging, etc.

2) *Uncontrollable content selection*. Given a table, the space of logically entailed descriptions is exponentially large, due to vast number of combinations of different operations and arguments from the table, e.g., count, comparison, superlative, etc. It is hard and uncontrollable for neural models to decide a valid, favorable choice of logical selections solely based on the table, due to the difficulty of imposing high-level semantic constraints in the compositional generation process.

To combat with the above problems, we argue that it is necessary to leverage intermediate meaning representations to achieve faithful and controllable logical generations. To this end, we formulate the task of logical-level NLG as a **logical form to text** problem. Specifically, besides the table information, the generation module is provided with a logical form representing the semantics of the target text (see Figure 1 for an example). By separating logical reasoning and language realization, the correctness of the intermediate logical form is guaranteed, and the challenge for the realization

module is fully shifted to semantic understanding.

To facilitate research in this direction, we propose a new dataset named LOGIC2TEXT, consisting of 5.6k open-domain tables, 10.8k manually annotated (logical form, description) pairs. Our dataset is of high quality in terms of (1) natural and interesting descriptions; (2) accurate logical forms with 100% execution correctness. In our dataset, the coarse logic types are 7 common ones to describe multi-row tables: *count*, *superlative*, *comparative*, *aggregation*, *majority*, *unique*, and *ordinal*. We employ a Python-like program to serve as our logical forms, which can be easily converted to other types of logical forms. Figure 1 shows two examples of our dataset. Compared with previous surface-level NLG datasets, one major distinction of our dataset is the free schema of the logical forms, which can be represented as diversified graph structures. The new dataset poses great challenges on the model’s ability to understand the structural semantics in graph representation.

We employ an array of popular generation models as the baseline approaches. The experiments are conducted in (1) *Fully-supervised setting*. We train the models using the full dataset to analyze their performances. (2) *Few-shot setting*. We simulate the low-resource scenario in real-world use cases. Experimental results show that the logical forms are critical to acquiring high-fidelity generations. The pre-trained language model outperforms other baselines (pointer-generator, graph2seq, transformer, etc.), but still makes factual and logical errors.

In summary, our contributions are the following:

- We propose a new large-scale dataset, LOGIC2TEXT, with descriptions of common logic types accompanied by the underlying logical forms. The logical forms present diversified graph structures, which raises more challenges on semantic understandings.
- We surveyed several popular generation models as the baselines under fully-supervised and few-shot settings, as well as analyze their pros and cons.

Our dataset can also be used in the reverse way (text to logical form) to facilitate tasks related to semantic parsing. Chen et al. (2019a) propose the task of fact verification against tables, however the performance is greatly limited due to the lack of

country	region	joined opec	population (july 2012)	area (km square)
algeria	africa	1969	37367226	2381740
angola	africa	2007	18056072	1246700
iraq	middle east	1960	31129225	437072
libya	africa	1962	5613380	1759540
nigeria	africa	1971	170123740	923768
...

Surface-level NLG

Description: angola, from the region africa, joined opec in 2007, with an population of 18056072 in 2012.

Description: algeria, from the region africa, joined opec in 1969, with an population of 37367226 in 2012.

Logical-level NLG with logical forms (our dataset)

logical form: eq { count { filter_eq { all_rows ; region ; africa } ; 4 } = True

Description: In 2012 in opec, there were 4 member countries from africa.

logical form: and { eq { hop { argmax { all_rows ; joined opec } ; region } ; africa } ; eq { hop { argmax { all_rows ; joined opec } ; country } ; angola } } = True

Description: In 2012 in opec, angola, from africa, was the latest country to join.

Figure 1: Examples of surface-level NLG compared with NLG with logical forms of our dataset. Here are two examples with logic type *count* and *superlative*. The function nodes are in blue, and the text nodes in grey.

the ground truth logical forms. This can be one direct application of our dataset. In this work, we focus on NLG.

2 Related Work

NLG from structured data or knowledge has been studied for many years. There are various applications, such as the automatic generations of weather reports (Liang et al., 2009), sport reports (Wiseman et al., 2017), clinical and health reports (DiMarco et al., 2007; Lee, 2018), response generation in task-oriented dialogue systems (Wen et al., 2015; Budzianowski et al., 2018; Dušek et al., 2019), etc.

Traditional methods typically employ a pipeline-based approach including content selection, planning and surface realization (Reiter and Dale, 1997; Gatt and Krahmer, 2018). Recent data-driven

methods tend to conflate the pipeline modules into one end-to-end neural networks, such as (Liu et al., 2018; Wiseman et al., 2017, 2018; Gong et al., 2019). Most recently, large-scale pre-trained models (Radford et al., 2019; Song et al., 2019; Raffel et al., 2019) have achieved new state-of-the-arts on various generation tasks. Chen et al. (2019b) demonstrate that a simple pre-training based method can achieve very reasonable performance on the WikiBio dataset (Lebret et al., 2016) under few-shot setting. More recent works begin to focus on fidelity preserving of the generation, such as (Dhingra et al., 2019; Tian et al., 2019). Their work obtains good performances on surface-level NLG. In contrast, our work focus on the fidelity of logical-level generations.

There are a few popular NLG datasets mostly on surface-level generation. Such as WeatherGov (Liang et al., 2009), E2E (Novikova et al., 2017), WikiBio (Lebret et al., 2016), and ToTTo (Parikh et al., 2020). RotoWire (Wiseman et al., 2017) is a more challenging dataset on generating basketball game reports from multi-row tables. But the reports are still limited to superficial restatements of table records, with very few involving logical inference. Korn et al. (2019) investigate generation of interesting trivia from superlative wikipedia tables. Chen et al. (2020) propose the task of generating arbitrary sentences with logical inference from the table. Their task mainly works for probing purpose, i.e., to test the ability of neural models to produce any logically correct descriptions solely based on the table. However, such a task formulation is not yet appropriate for building a real-world NLG system due to low-fidelity, as we discussed in the introduction. The best-performing model in (Chen et al., 2020) only obtains a factual correctness rate over 20% based on human evaluation, which is clearly far from an acceptable level in real-world systems.

Another line of works related to ours is the text generation from syntactic or semantic sentence structure, such as generation from CCG grammar (White, 2006), UCG grammar (Gardent and Plainfossé, 1990), AMR (Song et al., 2018). There are many early works attempting algorithmic approaches on such kinds of logical formulations (Phillips, 1993; Calder et al., 1989; Shieber et al., 1990; Phillips, 1993), etc. Later proposed datasets include the Groningen Meaning Bank (Bos, 2013), the AMR bank (May, 2016), the

DeepBank (Flickinger et al., 2012), etc. In contrast, our work focus on the logical formulations executed on database style tables, and common symbolic operations on tables, such as count, superlative, comparison. As nowadays much of the production data is stored in table based DB, we believe such a dataset should help building systems with table based data.

3 Dataset Construction

The table source of LOGIC2TEXT is from WikiTables¹ (Bhagavatula et al., 2013), a collection of open-domain tables crawled from Wikipedia. We follow (Chen et al., 2019a) to filter out over-complicated tables and take a subset of tables with less than 20 rows and 10 columns.

In this dataset, we start from 7 types of most commonly used logics (Chen et al., 2019a) to describe multi-row tables: count, superlative, comparative, aggregation, majority, unique, and ordinal. For example, for logic type count, the definition is: counting some rows in the table based on the values in one column, with the scope of all table rows or a subset. Refer to Appendix A for the definitions of all logic types. Each description involves exactly one type of logic. This matches the observation that humans generally do not describe their interested information in tables with over-complicated logics. For logical forms, we use a python-like program, and the function set is an extension of (Chen et al., 2019a). Refer to Appendix B for definitions of all functions.

Our dataset is constructed in 3 stages: §3.1 Description composition and verification, §3.2 Logical form annotation and derivation, §3.3 Logical form execution and verification. We adopt the workflow of composing descriptions first and then deriving the logical forms, because under such an order, the annotators can compose natural descriptions based on the interesting facts in the table, which is hard to be achieved by automatic enumeration of logical forms followed by template re-writing. For all crowd-sourcing tasks we hire Amazon Mechanical Turkers² (AMT) under three requirements: (1) from English native countries (“US”, “CA”, “GB”, “AU”); (2) Approval rate higher than 95% for all HITs; (3) More than 500 approved HITs. We follow the human subject

¹<http://websail-fe.cs.northwestern.edu/wikiTables/about/>

²<https://www.mturk.com/>

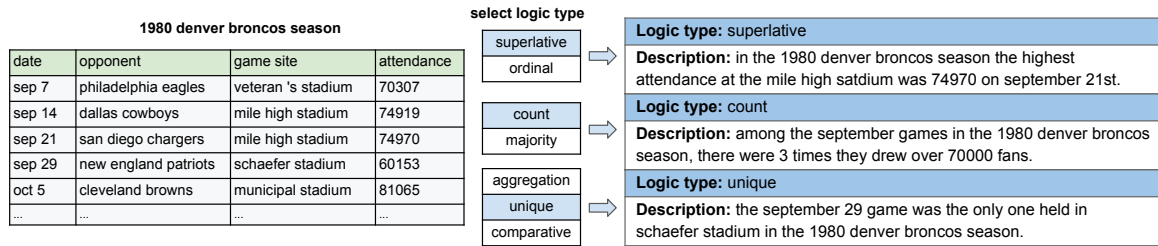


Figure 2: description composition: the workers are asked to select three logic types and compose a statement based on the selected logic type, that describe interesting facts in the table.

research protocols³ to pay the workers. We maintain strict high criterions for approval and review at least 10 random samples for each worker to decide whether to approve or reject all his/her HITs.

3.1 Description Composition & Verification

In this first stage, the human workers are asked to compose statements of a *certain logic type*, that describe *interesting* facts in the table. It’s possible that some logic types cannot be applied to certain tables. Therefore we design the following working procedure: For each table, the 7 logic types are randomly put into three groups (with sizes 2, 2, and 3). The worker is asked to choose one logic type from each group and compose a description based on the chosen logic type. They must follow the requirements (1) try to choose diversified logic types, (2) avoid template-like language and try to compose natural and interesting descriptions, (3) include the information in table captions, so as to compose comprehensive descriptions without unspecified pronouns. An example of the workflow is shown in Figure 2. We provide the workers detailed explanations for each logic type by their corresponding definitions, accompanied by examples. After collecting the descriptions, we add a verification stage to filter out descriptions of low quality. We redistribute the collected descriptions grouped by each logic type, then ask three questions: Is this description (1) of the correct logic type presented? (2) factually correct? (3) grammatically correct and fluent? We filter out the description if any question receives a negative response.

3.2 Logical Form Annotation & Derivation

As the core step of our dataset construction pipeline, we design a workflow to obtain the semantic information via conversations with human workers, then use the information to derive the logical forms. The

questions in the conversation are specifically designed for each logic type. Here we go through the example of logic type *superlative* given in Figure 3 to illustrate our annotation process.

The logical form structure prototype is shown in the right grey part, consisting the description of the superlative value, and other mentioned columns on the row with the superlative value. Then we ask the follow-up questions to derive the complete logical form based on the prototype, shown on the left part of Figure 3: Q1. What is the scope of the superlative operation? If the scope is a subset of all table rows, we perform another round of conversation to annotate the scope. Q2. What is the table column of the superlative operation? Q3. What is the specific type of the superlative operation: maximum or minimum. Q4. What is the table row with the superlative value. Q5. Is the superlative value itself mentioned in the description or not? Q6. What are the other columns mentioned in the description? After collecting the answers of the above questions, we can derive the logical form, as shown in the middle part of Figure 3.

We provide the workers with detailed explanations of the prototype for each logical types, as well as several examples. Note that the prototype covers most, but not all of the logical descriptions due to their diverse nature. Thus we also provide the option to skip the example if it cannot be formulated by the given question set. Check Appendix A for the annotation process of other logic types.

3.3 Logical Form Execution & Verification

After the collection of logical forms, we use the Stanford CoreNLP toolkits⁴ to tokenize all text content (all table information, the descriptions, and the texts in the logical forms). To remove incorrect logical forms, we execute the logical forms and perform another round of semantic verification.

³https://en.wikipedia.org/wiki/Minimum_wage_in_the_United_States

⁴<https://stanfordnlp.github.io/CoreNLP/index.html>

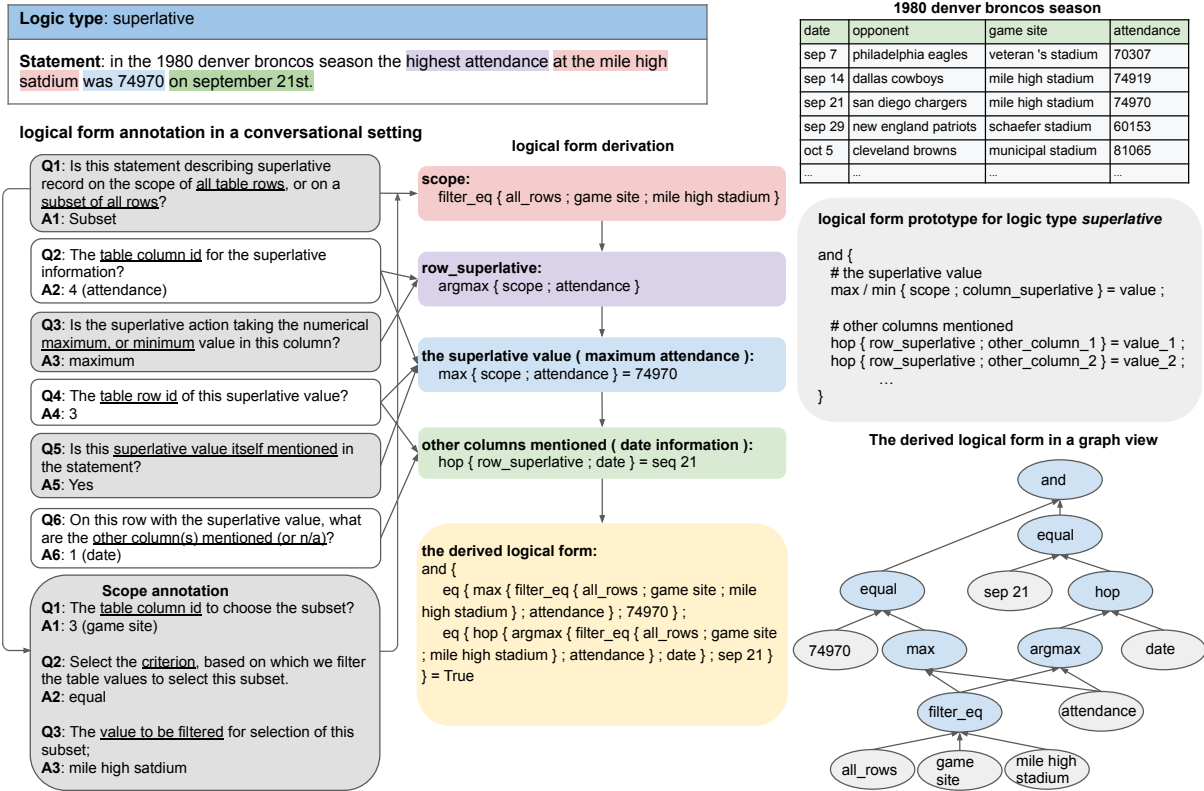


Figure 3: logical form annotation & derivation: Note that in this example the questions are all in concise forms. In the AMT interface shown to the workers, we write instructions in a more casual and detailed manner, accompanied by several examples.

Logical Form Execution The functionality in our logical form is based on the ones used in (Chen et al., 2019a). We extend the function set to deal with semi-structured table cells (dates, mixed numbers and strings, etc.). We execute all logical forms against the corresponding table, and only keeps the ones that evaluate to True. This guarantees that the logical forms in our dataset achieve 100% execution correctness.

Semantic Verification Note that execution correctness does not guarantee semantic correctness. Therefore we perform another round of semantic verification. Since AMT workers do not have experts knowledge to understand the logical forms, we convert the logical form into natural language interpretation based on the operations of each function. We then ask the workers to verify whether the interpretation correctly matches the meaning of the description, with neither insufficient nor redundant information. Then we remove the examples receiving negative responses.

Expert Evaluation To demonstrate the quality of our dataset, we employ two computer science graduate students to conduct evaluations. We randomly sample 200 examples for each logic type to verify the semantic correctness. Each example

is examined by both students, and the decision is made after discussion. The result shows that each logic type reaches a correct rate no less than 90%.

Tables	5,554
Examples	10,753
Vocabulary	14.0k
Avg. description length	16.77
Avg. # nodes in logical form	9.00
Avg. # function nodes in logical form	3.27
Avg. length of the linearized logical form	24.35

Table 1: General statistics of LOGIC2TEXT.

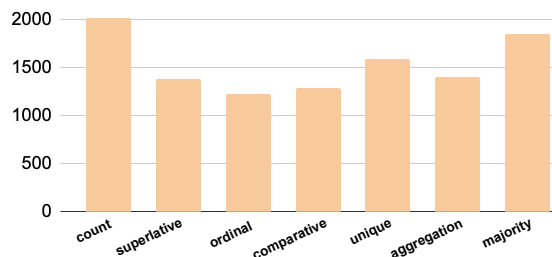


Figure 4: Distribution of logic types.

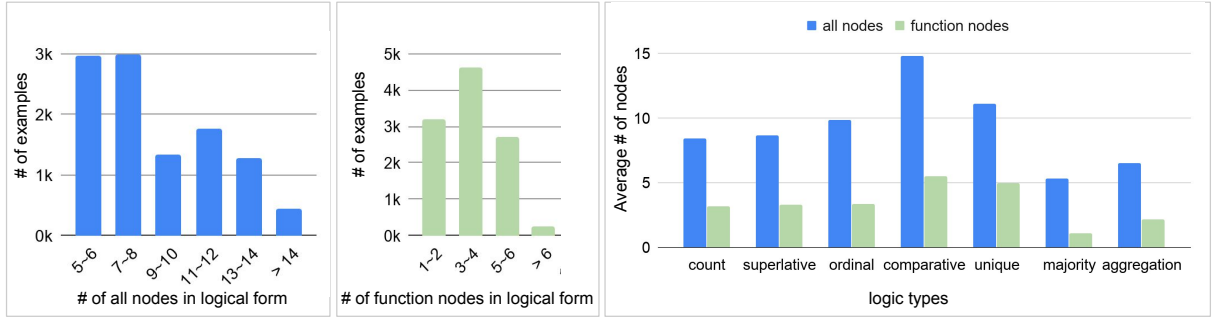


Figure 5: The distribution of our dataset regarding the number of all nodes (*Left*) and function nodes (*Mid*) in the logical form. *Right*: average number of all nodes and function nodes in the logical forms for each logic type.

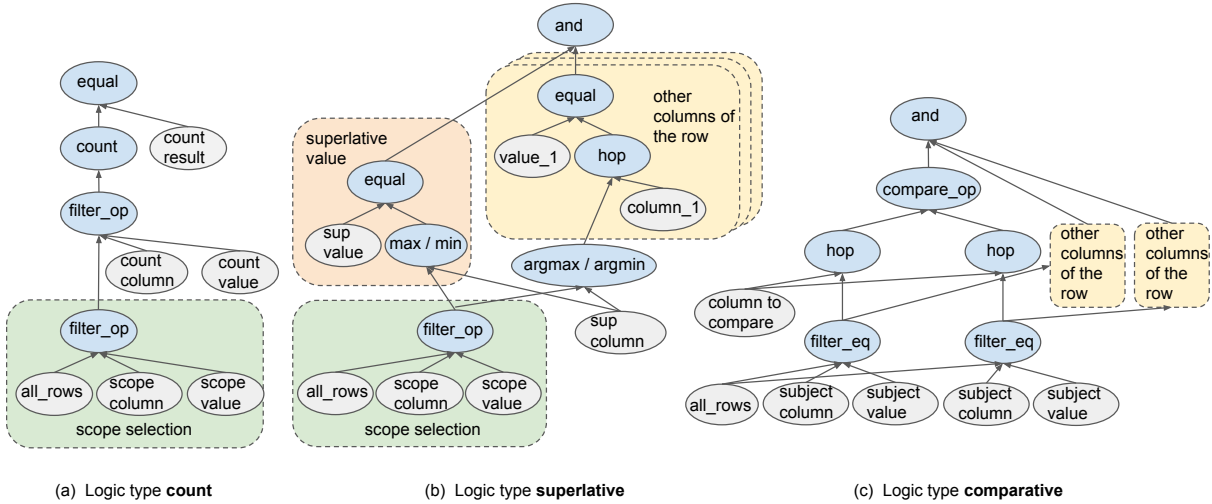


Figure 6: Overview of logical form structures for logic type *count*, *superlative*, and *comparative*. (a) *count*: the structure in the green shadow is optional, representing the scope of counting. It can be all table rows (a single text node) or a subset of rows from a filter operation. (b) *superlative*: the structure in the orange shadow is optional, depending on the presence of the max/minimum value in the description. The structure in the yellow shadow appears 0 or more times.

4 Dataset Statistics and Analysis

We follow a rough ratio of 8:1:1 to split our dataset into 8,566 for training, 1,095 for development, and 1,092 for testing. The train, dev, and test sets have no overlap tables. We show the statistics of the dataset in Table 1 and the distributions of 7 logic types in Figure 4. Each table has 1-3 descriptions with different logic types. Since the logical forms present graph structure nature, we analyze the complexity of the logical forms based on the number of nodes in the graph, regarding the number of function nodes (*count*, *max*, etc.) and the number of all nodes (both function nodes and text nodes), respectively. As shown in Figure 5, the logical forms in LOGIC2TEXT have a minimum of 5 nodes and maximum over 14 nodes. For different logic types, *comparative* has the most number of nodes, because it involves the selection and operation for two table rows. *superlative*, *ordinal*, and *unique* primarily focus on one table row, some-

times with the scope being a subset of all table rows, which makes the logical forms more complex. *count*, *majority*, and *aggregation* are summarization based logic types on multiple table rows. They are the three relatively simpler ones in terms of logical form structures. Figure 6 gives the logical form structures for 3 example logic types.

5 Experiments

In this section we first describe the baseline models of our dataset in §5.1; Then we conduct experiments in fully-supervised setting §5.2; We demonstrate the importance of the logical form in §5.3 and perform ablation studies in §5.4; At last we carry out experiments under few-shot setting §5.5.

5.1 Baseline Models

Apart from the logical forms serving as the primary input to the generation model, the table informa-

tion is also crucial to provide context information. Following human’s order to comprehend the table and produce descriptions, the input C is formulated as the sequence of table captions, table headers, table content, and the logical form. The goal is to generate a sequence w that maximize $P(w | C)$:

$$w = \operatorname{argmax} \prod P(w_t | w_{0:t-1}, C) \quad (1)$$

We employ the following models as our baselines for LOGIC2TEXT:

Template We manually craft generation templates for each logic type based on the logical form.

Seq2seq+att We employ the seq2seq with an attention model from (Bahdanau et al., 2015). The input sequence is formulated as the concatenation of the table caption, table headers, the linearized table content, and the linearized logical form.

Pointer generator (See et al., 2017) adds the copy mechanism upon the seq2seq with an attention model, allowing the decoder to copy tokens from the input directly. Such a mechanism is known to be critical for fidelity-preserving generation with abundant entities, numbers, etc.

Graph2seq+copy There is a line of research for graph neural network based encoders, such as (Marcheggiani and Perez-Beltrachini, 2018; Xu et al., 2018), etc. We employ one representative model, Graph2seq (Xu et al., 2018), to encode the logical forms. The table caption and headers are first fed into a seq2seq, followed by the graph encoder for the logical form. We also add the copy mechanism to allow copying from the input.

Transformer+copy The popular Transformer model (Vaswani et al., 2017) has shown remarkable progress in many tasks including NLG. In addition to the original Transformer structure, we add the copy mechanism where the last hidden layer is used to calculate the attention score and the copy switch. We also add segment embeddings for different input components, similar as (Devlin et al., 2019).

GPT-2 Generally, with Transformer based structures, recent large-scale pre-trained models have achieved new SOTA results in a wide range of NLP tasks. A typical workflow is to use the pre-trained model as initialization, then fine-tune the model on task-specific data. In this work, we employ the generative pre-training model, GPT-2 (Radford et al., 2019), as one of our baselines.

For all neural models we use Byte-Pair Encoding (BPE) (Sennrich et al., 2016) and the subword vocabulary used in (Radford et al., 2019). Refer to Appendix C for more implementation details.

5.2 Fully-Supervised Setting

For automatic evaluations, we employ BLEU-4⁵ (B-4), ROUGE-1, 2, 4, and L (F measure)⁶, noted as R-1, R-2, R-4, and R-L. The results for all baselines are presented in Table 2.

For models without pre-training, the copy mechanism brings a significant improvement, comparing pointer-generator and seq2seq. This is because the descriptions in our dataset involve much factual information from the table and the logical form, e.g., entity names, and numbers. However, the pre-trained language model GPT-2 can mostly accurately produce these factual terms even without a copy mechanism, demonstrating the powerful prior knowledge obtained from large-scale pre-training.

Models	B-4	R-1	R-2	R-4	R-L
Template	17.57	50.56	24.20	6.61	37.81
Seq2seq+att	12.46	36.22	15.91	4.49	31.03
Pointer generator	24.03	56.23	30.51	10.78	46.85
Graph2seq+copy	25.38	58.15	32.79	12.25	49.47
Transformer+copy	26.42	58.77	33.05	12.83	49.01
GPT-2	31.44	64.16	39.48	17.46	53.99

Table 2: Automatic evaluation results for all baseline models under fully-supervised setting.

Compared to the pointer generator, which takes linearized logical form as input, Graph2seq+copy directly models the graph structure and gets a slight improvement. The Transformer+copy model obtains better performance than the Graph2seq+copy model, as the Transformer architecture is indeed a graph neural network with self-attention as aggregation function over the neighbors and regards the input as a fully-connected graph. Recent works (Lin et al., 2019; Rogers et al., 2020; Mager et al., 2020) have shown that Transformer-based structure can capture hierarchical syntactic structures and graph representations. The GPT-2 model obtains the best performance among all with a significantly larger improvement. As a pre-trained language model with the Transformer structure, it combines the strength of both structural modeling and language modeling prior. Some example generations are provided in Appendix E.

Human Evaluation

Automatic scores are not sufficient for precise evaluation of factual and logical correctness. Therefore we conduct human evaluations through (1) crowd-

⁵Standard script NIST mteval-v13a.pl

⁶rouge-1.5.5.

sourcing on Amazon Mechanical Turkers (AMT), and (2) human expert evaluations.

For human evaluations on AMT, we randomly sample 500 examples from each of the top best-performing methods (GPT-2 and Transformer+copy), and the gold references. The evaluations are conducted on two axes: *factual correctness* and *language fluency*. For factual correctness, we ask the workers to verify whether the description is factually supported by the table; For language fluency, we conduct pairwise comparisons between different methods. For both evaluations, we distribute each task to 3 workers to eliminate human variance. The evaluation results of language fluency and factual correctness are shown in Table 4 and the first row of Table 3, respectively. For more details of the evaluation, check Appendix D. To conduct a precise evaluation of semantic

	Gold	GPT-2	Transformer+copy
% factually correct	98.1	82.4	65.1
% semantically correct	92.0	73.0	43.0

Table 3: Human evaluation results of factual correctness (first row) and semantic correctness (second row).

	% win	% loss	% tie
GPT-2 vs Gold	35.6	43.3	21.1
GPT-2 vs Transformer+copy	54.0	25.3	20.7
Gold vs Transformer+copy	61.2	23.6	15.2

Table 4: Human evaluation results of language fluency.

correctness, i.e., whether the generation correctly matches the meaning of the logical form, we invite human experts (two computer science graduate students) to perform the evaluation. We sample 200 examples from each method and ask them to verify whether the description correctly presents the meaning of the logic form. Each example is examined by both students, and the decision is made after discussion. The second row of Table 3 shows the evaluation results.

As we can observe from all evaluation results, the GPT-2 model gives big improvements on both fidelity preserving and language fluency, but there’s still a gap, especially on semantic correctness. We believe our dataset can serve as a valuable resource posing such a challenge on high-fidelity generation with complex semantics.

5.3 Importance of the Logical Form

We conduct experiments without using the logical form, i.e., to generate arbitrary logically correct descriptions solely based on the table, which is the task setting of (Chen et al., 2020). The generation is evaluated with all descriptions of the same table as multi-references, as in their setting. The best performing model of (Chen et al., 2020) obtains a BLEU-4 score of 20.17 and a factual correctness rate of 20.2% based on human evaluation of 500 samples. In contrast, the generations of our best-performing baseline can obtain a factual correctness rate of 82.4% shown in Table 3, which demonstrates the great importance of the logical form on high-fidelity generation. Note that the automatic scores are not directly comparable, since, in our task setting, each generation maps to a unique logical form and is evaluated with a single reference.

5.4 Component-Wise Ablation

Models	B-4	R-1	R-2	R-4	R-L
GPT-2	31.44	64.16	39.48	17.46	53.99
-w/o caption	21.67	54.26	29.16	9.99	45.70
-w/o header	29.86	62.98	38.46	16.64	52.57
-w/o content	30.42	64.17	38.89	16.79	53.63

Table 5: Ablation study on other input components.

We perform ablation studies on other input components: the table caption, header, and content, using the best-performing GPT-2 model. As shown in Table 5, both the table caption and header provide strong context information for generation, and the table content also brings a slight improvement.

5.5 Few-Shot Setting

Considering that acquiring a large amount of (logical form, description) pairs in real-world cases is expensive, we also include a few-shot learning task for our dataset, where the model is only provided with hundreds of paired examples. Previous works have shown that the pre-trained language models obtain strong NLG performance even with a handful of fine-tuning instances (Chen et al., 2019b). Therefore we still use the best-performing GPT-2 model for this study. In our dataset, the amount of unseen logical form structures increases with the reduction of training instances. As shown in Table 6, while there’s still a gap with the fully-supervised result, the result with 1,000 training instances using GPT-2 is comparable to some other baselines

with the full training data. This demonstrates the potential of incorporating generative pre-training for the few-shot learning task.

# of examples	B-4	R-1	R-2	R-4	R-L
Full	31.44	64.16	39.48	17.46	53.99
100	17.09	48.26	23.52	7.47	38.74
200	19.98	51.99	27.02	9.42	41.86
500	23.04	56.64	30.99	11.35	46.86
1000	24.57	57.81	32.64	12.21	47.67

Table 6: Results for few-shot learning setting with 100, 200, 500, and 1000 training examples, using GPT-2.

6 Conclusion

In this work, we formulate the problem of logical-level NLG as generation from logical forms in order to obtain controllable and high-fidelity generations. To this end, we propose a new dataset named LOGIC2TEXT. There are some other potential future directions. 1) Human evaluations are precise but expensive. Our dataset can be used in the reverse direction to train a semantic parser, to assist parsing-based evaluations. 2) In this work, we primarily focus on the step to generate descriptions based on the logical form. Another potential future direction could be the content selections, i.e., how to select and organize the logical forms to construct a discourse plan based on user interests.

Acknowledgment

We thank the anonymous reviewers for their thoughtful comments. This research was sponsored in part by Intel AI Faculty Research Grant and NSF IIS 1528175. The authors are solely responsible for the contents of the paper and the opinions expressed in this publication do not reflect those of the funding agencies.

References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

Chandra Sekhar Bhagavatula, Thanapon Noraset, and Doug Downey. 2013. [Methods for exploring and mining tables on wikipedia](#). In *Proceedings of the ACM SIGKDD Workshop on Interactive Data Exploration and Analytics, IDEA@KDD 2013, Chicago, Illinois, USA, August 11, 2013*, pages 18–26. ACM.

Johan Bos. 2013. [The groningen meaning bank](#). In *Proceedings of the Joint Symposium on Semantic Processing, Textual Inference and Structures in Corpora, JSSP 2013, Trento, Italy, November 20-22, 2013*, page 2. Association for Computational Linguistics.

Pawel Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018. [Multiwoz - A large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 5016–5026. Association for Computational Linguistics.

Jonathan Calder, Mike Reape, and Henk Zeevat. 1989. [An algorithm for generation in unification categorical grammar](#). In *EACL 1989, 4th Conference of the European Chapter of the Association for Computational Linguistics, April 10-12, 1989, University of Manchester, Institute of Science and Technology, Manchester, England*, pages 233–240. The Association for Computer Linguistics.

Wenhu Chen, Jianshu Chen, Yu Su, Zhiyu Chen, and William Yang Wang. 2020. [Logical natural language generation from open-domain tables](#). *arXiv preprint arXiv:2004.10404*.

Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyou Zhou, and William Yang Wang. 2019a. [Tabfact: A large-scale dataset for table-based fact verification](#). *CoRR*, abs/1909.02164.

Zhiyu Chen, Harini Eavani, Yinyin Liu, and William Yang Wang. 2019b. [Few-shot NLG with pre-trained language model](#). *CoRR*, abs/1904.09521.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.

Bhuwan Dhingra, Manaal Faruqui, Ankur P. Parikh, Ming-Wei Chang, Dipanjan Das, and William W. Cohen. 2019. [Handling divergent reference texts when evaluating table-to-text generation](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4884–4895. Association for Computational Linguistics.

Chrysanne DiMarco, HDominic Covvey, D Cowan, V DiCiccio, E Hovy, J Lipa, D Mulholland, et al.

2007. The development of a natural language generation system for personalized e-health information. In *Medinfo 2007: Proceedings of the 12th World Congress on Health (Medical) Informatics; Building Sustainable Health Systems*, page 2339. IOS Press.
- Ondřej Dušek, Jekaterina Novikova, and Verena Rieser. 2019. Evaluating the state-of-the-art of end-to-end natural language generation: The E2E NLG Challenge. *arXiv preprint arXiv:1901.11528*.
- Dan Flickinger, Yi Zhang, and Valia Kordoni. 2012. Deepbank. a dynamically annotated treebank of the wall street journal. In *Proceedings of the 11th International Workshop on Treebanks and Linguistic Theories*, pages 85–96.
- Markus Freitag and Scott Roy. 2018. Unsupervised natural language generation with denoising autoencoders. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 3922–3929. Association for Computational Linguistics.
- Claire Gardent and Agnes Plainfossé. 1990. Generating from a deep structure. In *COLING 1990 Volume 2: Papers presented to the 13th International Conference on Computational Linguistics*.
- Albert Gatt and Emiel Kraemer. 2018. Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *J. Artif. Intell. Res.*, 61:65–170.
- Heng Gong, Xiaocheng Feng, Bing Qin, and Ting Liu. 2019. Table-to-text generation with effective hierarchical encoder on three dimensions (row, column and time). 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 2019, Hong Kong, China, November 3-7, 2019*, pages 3141–3150. Association for Computational Linguistics.
- Flip Korn, Xuezhi Wang, You Wu, and Cong Yu. 2019. Automatically generating interesting facts from wikipedia tables. In *Proceedings of the 2019 International Conference on Management of Data, SIGMOD Conference 2019, Amsterdam, The Netherlands, June 30 - July 5, 2019*, pages 349–361. ACM.
- Rémi Lebret, David Grangier, and Michael Auli. 2016. Neural text generation from structured data with application to the biography domain. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 1203–1213.
- Scott H Lee. 2018. Natural language generation for electronic health records. *NPJ digital medicine*, 1(1):63.
- Percy Liang, Michael I. Jordan, and Dan Klein. 2009. Learning semantic correspondences with less supervision. In *ACL 2009, Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics and the 4th International Joint Conference on Natural Language Processing of the AFNLP, 2-7 August 2009, Singapore*, pages 91–99.
- Yongjie Lin, Yi Chern Tan, and Robert Frank. 2019. Open sesame: Getting inside bert’s linguistic knowledge. *CoRR*, abs/1906.01698.
- Tianyu Liu, Kexiang Wang, Lei Sha, Baobao Chang, and Zhifang Sui. 2018. Table-to-text generation by structure-aware seq2seq learning. 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 4881–4888.
- Manuel Mager, Ramón Fernández Astudillo, Tahira Naseem, Md. Arafat Sultan, Young-Suk Lee, Radu Florian, and Salim Roukos. 2020. Gpt-too: A language-model-first approach for amr-to-text generation. *CoRR*, abs/2005.09123.
- Diego Marcheggiani and Laura Perez-Beltrachini. 2018. Deep graph convolutional encoders for structured data to text generation. In *Proceedings of the 11th International Conference on Natural Language Generation, Tilburg University, The Netherlands, November 5-8, 2018*, pages 1–9. Association for Computational Linguistics.
- Jonathan May. 2016. Semeval-2016 task 8: Meaning representation parsing. In *Proceedings of the 10th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2016, San Diego, CA, USA, June 16-17, 2016*, pages 1063–1073. The Association for Computer Linguistics.
- Jekaterina Novikova, Ondrej Dusek, and Verena Rieser. 2017. The E2E dataset: New challenges for end-to-end generation. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue, Saarbrücken, Germany, August 15-17, 2017*, pages 201–206.
- Ankur P. Parikh, Xuezhi Wang, Sebastian Gehrmann, Manaal Faruqui, Bhuwan Dhingra, Diyi Yang, and Dipanjan Das. 2020. Totto: A controlled table-to-text generation dataset. *CoRR*, abs/2004.14373.
- John D. Phillips. 1993. Generation of text from logical formulae. *Mach. Transl.*, 8(4):209–235.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.

- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *CoRR*, abs/1910.10683.
- Ehud Reiter and Robert Dale. 1997. [Building applied natural language generation systems](#). *Natural Language Engineering*, 3(1):57–87.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. [A primer in bertology: What we know about how BERT works](#). *CoRR*, abs/2002.12327.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). 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 1073–1083.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.
- Stuart M. Shieber, Gertjan van Noord, Fernando C. N. Pereira, and Robert C. Moore. 1990. [Semantic-head-driven generation](#). *Comput. Linguistics*, 16(1):30–42.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. [MASS: masked sequence to sequence pre-training for language generation](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 5926–5936. PMLR.
- Linfeng 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, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 1616–1626. Association for Computational Linguistics.
- Ran Tian, Shashi Narayan, Thibault Sellam, and Ankur P. Parikh. 2019. [Sticking to the facts: Confident decoding for faithful data-to-text generation](#). *CoRR*, abs/1910.08684.
- 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.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrksic, Peihao Su, David Vandyke, and Steve J. Young. 2015. [Semantically conditioned lstm-based natural language generation for spoken dialogue systems](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 1711–1721. The Association for Computational Linguistics.
- Michael White. 2006. [Efficient realization of coordinate structures in combinatory categorial grammar](#). *Research on Language and Computation*, 4(1):39–75.
- Sam Wiseman, Stuart M. Shieber, and Alexander M. Rush. 2017. [Challenges in data-to-document generation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2253–2263.
- Sam Wiseman, Stuart M. Shieber, and Alexander M. Rush. 2018. [Learning neural templates for text generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 3174–3187.
- Kun Xu, Lingfei Wu, Zhiguo Wang, Yansong Feng, and Vadim Sheinin. 2018. [Sql-to-text generation with graph-to-sequence model](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 931–936. Association for Computational Linguistics.

Appendix

A. Logic Type Definitions & Logical Form Annotation

Logic Type Definitions

We define all 7 logic types in our dataset and provide examples based on the following table in Figure 7.

table caption: opec

country	region	joined opec	population (july 2012)	area (km square)
algeria	africa	1969	37367226	2381740
angola	africa	2007	18056072	1246700
iraq	middle east	1960	31129225	437072
kuwait	middle east	1960	2646314	17820
libya	africa	1962	5613380	1759540
nigeria	africa	1971	170123740	923768
qatar	middle east	1961	1951591	11437
saudi arabia	middle east	1960	26534504	2149690
united arab emirates	middle east	1967	5314317	83600
venezuela	south america	1960	28047938	912050

Figure 7: Example table

Count: counting some rows in the table based on the values in one column, with the scope of all table rows or a subset of rows.

Example descriptions: “in opec 2012, there were 4 countries from africa.”, “in opec 2012, among the countries from africa, 2 of them joined after 1970.”, etc.

Superlative: Describing the maximum or minimum value in a column, with the scope of all table rows or a subset of rows. You may also talk about other columns on this row with the superlative value.

Example descriptions: “in opec in 2012, angola, from africa, was the latest country to join.”, “among the member countries in opec in 2012 from the middle east, qatar was the smallest in area.”, etc.

Ordinal: Describing the n-th maximum or minimum value in a column, with the scope of all table rows or a subset of rows. You may also talk about other columns on this row with the n-th maximum or minimum value.

Example descriptions: “in opec in 2012, qatar was the 5th country to join.”, “Among the africa member countries, algeria was the 2nd earliest to join.”, etc.

Comparative: Comparing two rows in the table, regarding their values in one column. You may also talk about other columns on these two rows.

Example descriptions: “in opec in 2012, libiya joined 2 years later than kuwait.”, “in opec in 2012, algeria, from africa, had a larger population than iraq from the middle east.”

Aggregation: Describing the sum or average value over a column, with the scope of all table rows or a subset of rows.

Example descriptions: “in opec 2012, the countries from africa had an average population of around 57,800,000.”, etc.

Unique: Describing one unique row, regarding one column, with the scope of all table rows or a subset of rows. You may also talk about other columns on this unique row.

Example descriptions: “in opec 2012, angola was the only country to join after 2000.”, “in 2012, among the member countries from africa, the only one to join opec after 2000 is angola.”, etc.

Majority: Describing the majority values (most or all) over one column, with the scope of all table rows or a subset of rows.

Example descriptions: “in opec 2012, most countries joined before 2000.”, “in opec 2012, all of the africa member countries had an area larger than 900,000.”, etc.

Logical Form Annotation

Here we provide the question sets for annotating each logical type.

Count: (1). Choose whether the counting is performed on the scope of all table rows, or on a subset of all rows. (2). Select the table column that the counting is performed on. (3). Select the criterion, based on which we filter the table records to be counted. Here we consider the following criterion: “equal”, “not equal”, “less than”, “less than or equal to”, “greater than”, “greater than or equal to”, “fuzzily match”, “all” (or “other” if none of the above is correct). (4). Based on the selected criterion, write the value to be filtered for counting. (5). Write down the result of the counting.

Superlative: (1). Is the superlative action performed on the scope of all table rows, or on a subset of all rows? (2). What is the table column that the superlative action is performed on? (3). Is the superlative action taking the numerical maximum, or minimum value among the records? (4). What is the table row containing this superlative value? (5). On this row with the superlative value, what are the other column(s) mentioned? If not any other column is mentioned, write ‘n/a’. (6). Is this superlative value itself mentioned in the statement?

Aggregation: (1). Choose whether the aggregation is performed on the scope of all table rows, or on a subset of all rows. (2). Select the table column that the aggregation is performed on. (3). What is the type of this aggregation, sum or average? (4). What is the result of this aggregation?

Comparative: (1). Which column is the statement comparing? (2). What is the first row to be compared? (3). What is the second row to be compared? (4). What is the relationship comparing the records numerically in the first row with the second? (choose from “greater”, “less”, “equal”, “not equal”, “difference value”, or “other” if not any of the above. Here we consider the relationship between actual numerical values between two records, NOT the relationship expressed in the statement)

(5). Is the compared records itself mentioned in the statement? (6). What are the other column(s) of these two rows mentioned in the statement?

Majority: (1). What is the scope of this majority? (2). Which column the statement is describing? (3). Is the statement describing all the records or most frequent records within the scope? (4). Select the criterion, based on which we filter records to describe the majority. Here we consider the following criterion: "equal", "not equal", "less than", "less than or equal to", "greater than", "greater than or equal to", "fuzzily match" (or "other" if none of the above is correct). (5). Based on the selected criterion, write the value to be filtered for describing the majority.

Ordinal: (1). What is the scope that the ordinal description is performed on? (all rows or a subset of rows) (2). What is the table column that the ordinal description is based on? (3). Is the ordinal description based on a numerically max to min or min to max ranking of the column records? (4). What is the order described in the statement, based on this ranking? (5). What is the table row containing this n-th record ? (6). On this row, what are the other column(s) mentioned? If not any other column is mentioned, write 'n/a'. (7). Is this n-th record itself mentioned in the statement?

Unique: (1). What is the scope of this statement describing unique row? (2). What is this unique row? (3). Write the table column that shows the uniqueness of this row (4). Select the criterion, based on which we filter records in this column to find the unique row. Here we consider the following criterion: "equal", "not equal", "less than", "greater than", "fuzzily match" (or "other" if none of the above is correct). (5). Based on the selected criterion, write the value to be filtered for the unique row. (6). On this unique row, what are the other column(s) mentioned (except the column describing the scope)? If not any other column is mentioned, write 'n/a'.

B. Function Definitions

Here we list the function definitions and descriptions for our logical form in table 7. Note that since the tables in WikiTables are not standard database table, but semi-structured tables, the cell values are often not well-formatted with a lot of mixed strings and numbers, dates in different formats, etc. Therefore for some functions involving arithmetic operations on table cell values, we only specify a

coarse "object" type for the arguments, and then parse the numerical or date type values in the function implementations. Refer to our released code for detailed implementations.

C. Model Implementation Details

Here we provide some implementation details of the baseline models.

Template Some example templates are listed below. Texts in braces is optional depending on the logical form.

count:

in [table_caption], (among the ones whose [scope_column] are [equal to/greater than/...] [scope_value]), there are [result] ones whose [column_name] are [equal to/greater than/...] [value]

superlative:

in [table_caption], (among the ones whose [scope_column] are [equal to/greater than/...] [scope_value]), the [max/minimum] [column_name] is [value].

in [table_caption], (among the ones whose [scope_column] are [equal to/greater than/...] [scope_value]), [subject], with ([other_col1] [other_val];...), has the [max/minimum] [column_name], ([value]).

ordinal:

similar as *superlative*, replace max/minimum as n-th max/minimum.

comparative:

in [table_caption], [subject1] has [greater/less/...] [column_name] than [subject2].

in [table_caption], [subject1] has [diff_value] [column_name] [greater/less/...] than [subject2].

in [table_caption], [subject1], with ([other_col1] [other_val];...), has [greater/less/...] [column_name] than [subject2], with ([other_col1] [other_val];...).

unique:

in [table_caption], (among the ones whose [scope_column] are [equal to/greater than/...] [scope_value]), there is only one of them whose [column_name] is [greater/less /...] than [value].

in [table_caption], (among the ones whose [scope_column] are [equal to/greater than/...] [scope_value]), the only one whose [column_name] is [greater/less/...] than [value] is for [subject], with ([other_col1] [other_val];...).

aggregation:

Name	Arguments	Output	Description
count	view	number	returns the number of rows in the view
only	view	bool	returns whether there is exactly one row in the view
hop	row, header string	object	returns the value under the header column of the row
and	bool, bool	bool	returns the boolean operation result of two arguments
max/min/avg/sum	view, header string	number	returns the max/min/average/sum of the values under the header column
nth_max/nth_min	view, header string	number	returns the n-th max/n-th min of the values under the header column
argmax/argmin	view, header string	row	returns the row with the max/min value in header column
nth_argmax/nth_argmin	view, header string	row	returns the row with the n-th max/min value in header column
eq/not_eq	object, object	bool	returns if the two arguments are equal
round_eq	object, object	bool	returns if the two arguments are roughly equal under certain tolerance
greater/less	object, object	bool	returns if argument 1 is greater/less than argument 2
diff	object, object	object	returns the difference between two arguments
filter_eq/not_eq	view, header string, object	view	returns the subview whose values under the header column is equal/not equal to argument 3
filter_greater/less	view, header string, object	view	returns the subview whose values under the header column is greater/less than argument 3
filter_greater_eq/less_eq	view, header string, object	view	returns the subview whose values under the header column is greater/less or equal than argument 3
filter_all	view, header string	view	returns the view itself for the case of describing the whole table
all_eq/not_eq	view, header string, object	bool	returns whether all the values under the header column are equal/not equal to argument 3
all_greater/less	view, header string, object	bool	returns whether all the values under the header column are greater/less than argument 3
all_greater_eq/less_eq	view, header string, object	bool	returns whether all the values under the header column are greater/less or equal to argument 3
most_eq/not_eq	view, header string, object	bool	returns whether most of the values under the header column are equal/not equal to argument 3
most_greater/less	view, header string, object	bool	returns whether most of the values under the header column are greater/less than argument 3
most_greater_eq/less_eq	view, header string, object	bool	returns whether most of the values under the header column are greater/less or equal to argument 3

Table 7: Function definitions

in [table_caption], (among the ones whose [scope_column] are [equal to/greater than/...] [scope_value]), the [average/sum] of [column_name] is [result].

majority:

in [table_caption], (among the ones whose [scope_column] are [equal to/greater than/...] [scope_value]), [most/all] of them has [column_name] [equal to/greater than/ ...] [majority_value].

For all neural models we use Byte-Pair Encoding (BPE) (Sennrich et al., 2016) and the subword vocabulary used in (Radford et al., 2019). We use the pre-trained word embeddings from (Radford et al., 2019) and project to certain smaller dimensions (300) as the word embeddings. The batch size of all models are set to 32. The beam size is set to 3. As the table content only serves as context information for generation, to save GPU memory we set the maximum length of the table content as 200. The hyperparameters are chosen based on manual tuning regarding the BLEU score on the validation set.

Seq2seq+att & pointer-generator The learning rate is set to 0.001. For seq2seq, the training takes around 16000 gradient steps. For pointer generator, training takes around 5000 steps.

Graph2seq+copy we reuse the code skeleton from the released code from (Xu et al., 2018). The table

caption and header are first fed into a seq2seq, then the final hidden state is used to initialize the nodes of the graph encoder. When applying attention and copy, for graph nodes, we concatenate the token embedding and the embedding of its node as the embedding for the token. The learning rate is set to 0.0005. Training takes around 11000 steps.

Transformer+copy we mostly follow the structure setting in the original Transformer model (Vaswani et al., 2017). We use 4 attention heads and 6 layers. The final hidden layer is used for calculating the attention score and the copy switch. We also add the segment embeddings for different input components similar as (Devlin et al., 2019). The learning rate is set to 0.0005. training takes around 32000 steps.

GPT-2 We use the GPT-2 small 117M model from the released code and pre-trained model from (Radford et al., 2019). Word embeddings are fixed during training. The learning rate is set to 0.0003. The training takes around 500 steps to converge.

All the experiments are run on GeForce GTX 1080Ti GPU. Table 8 shows the validation performance of different baselines.

D. Human Evaluation Details

Human Evaluations on AMT We randomly sample 500 examples from the top two best performing methods (GPT-2 and Transformer+copy), and the gold references. The evaluations are conducted on two axes: *factual correctness* and *language fluency*.

Models	B-4	R-1	R-2	R-4	R-L
Template	17.81	51.16	24.89	6.68	38.12
Seq2seq+att	12.26	35.44	15.68	4.81	30.36
Pointer generator	25.43	57.35	31.97	12.33	48.11
Graph2seq+copy	25.65	57.65	31.98	12.29	48.28
Transformer+copy	27.20	59.70	34.06	14.03	48.71
GPT-2	32.98	64.86	40.02	18.38	54.59

Table 8: Automatic evaluation results for validation set.

For factual correctness, we provide the workers with both the table and the description, and ask them to verify whether the description is factually correct based on the table. If the description contains too many grammar errors to be readable, the worker is instructed to select "incorrect". Minor grammar errors can be accepted, as long as the worker can understand the meanings. For language fluency, we conduct pairwise comparison between the three methods. For this evaluation we only present the pair of descriptions to the worker, and ask them to select a better one only based on language fluency (a better description should be fluent, coherent, and free of grammar errors), or select "Tied" if the two descriptions are of similar quality. For both evaluations we distribute each task to 3 workers to eliminate human variance.

Human Expert Evaluation To conduct precise evaluation of semantic correctness, i.e., whether the generation correctly matches the meaning of the logical form, we invite human experts (two computer science graduate students) to perform the evaluation. We sample 200 examples from each method and ask them to verify whether the description correctly presents the meaning of the logic form, with neither insufficient nor redundant information. The description should also be fluent and free of grammar errors. Therefore this evaluation can be seen as a comprehensive evaluation of the generation quality. Each example is examined by both students and the decision is made after discussion.

E. Generation Examples

We provide 2 examples of generations in Figure 8 and Figure 9.

east coast conference

institution	nickname	location	founded	type	enrollment	joined
university of bridgeport	purple knights	bridgeport , connecticut	1927	private	4018	2000
daemen college	wildcats	amherst , new york	1947	private (nonsectarian)	2100	2013
university of the district of columbia	firebirds	washington , dc	1851	public	5471	2011
dowling college	golden lions	oakdale , new york	1963	private	7000	1989
mercy college	mavericks	dobbs ferry , new york	1950	private	10000	1989
molloy college	lions	rockville centre , new york	1955	private	3533	1989
new york institute of technology	bears	old westbury , new york	1955	private	12755	1989
queens college	knights	flushing , new york	1937	public	17639	1989
roberts wesleyan college	redhawks	chili , new york	1866	private (free methodist)	2000	2012

Logical form: greater { hop { filter_eq { all_rows ; institution ; mercy college } ; enrollment } ; hop { filter_eq { all_rows ; institution ; dowling college } ; enrollment } } = true

Gold: in the east coast conference , more people attended school at mercy college than at dowling college .

GPT-2: in the east coast conference , mercy college has a greater enrollment than dowling college .

Transformer+copy: more people attend the enrollment in the north coast conference than dowling college .

Figure 8: Example generations.

sheffield shield

rank	s wicket	player	matches	average
1	513	clarrie grimmitt (vic / sa)	79	25.29
2	441	michael kasprowicz (qld)	101	24.56
3	430	andy bichel (qld)	89	23.24
4	419	jo angel (wa)	105	24.86
5	384	terry alderman (wa)	97	24.21

Logical form: and { eq { max { all_rows ; average } ; 25.29 } ; eq { hop { argmax { all_rows ; average } ; player } ; clarrie grimmitt (vic / sa) } } = true

Gold: clarrie grimmitt had the highest average in the sheffield shield , 25.29 .

GPT-2: clarkrie grimmitt was the player with the highest average in the sheffield shield .

Transformer+copy: in the player that had 25.29 , the highest number of average average average attendance for the player who had 25.29 .

Figure 9: Example generations.