

Two Birds, One Stone: A Simple, Unified Model for Text Generation from Structured and Unstructured Data

Hamidreza Shahidi,^{1,2} Ming Li,^{1,2} and Jimmy Lin^{1,2}

¹ David R. Cheriton School of Computer Science, University of Waterloo

² RSVP.ai

{h24shahi, mli, jimmylin}@uwaterloo.ca

Abstract

A number of researchers have recently questioned the necessity of increasingly complex neural network (NN) architectures. In particular, several recent papers have shown that simpler, properly tuned models are at least competitive across several NLP tasks. In this work, we show that this is also the case for text generation from structured and unstructured data. We consider neural table-to-text generation and neural question generation (NQG) tasks for text generation from structured and unstructured data, respectively. Table-to-text generation aims to generate a description based on a given table, and NQG is the task of generating a question from a given passage where the generated question can be answered by a certain sub-span of the passage using NN models. Experimental results demonstrate that a basic attention-based seq2seq model trained with the exponential moving average technique achieves the state of the art in *both* tasks. Code is available at <https://github.com/h-shahidi/2birds-gen>.

1 Introduction

Recent NLP literature can be characterized as increasingly complex neural network architectures that eke out progressively smaller gains over previous models. Following a previous line of research (Melis et al., 2018; Mohammed et al., 2018; Adhikari et al., 2019), we investigate the necessity of such complicated neural architectures. In this work, our focus is on text generation from structured and unstructured data by considering description generation from a table and question generation from a passage and a target answer.

More specifically, the goal of the neural table-to-text generation task is to generate biographies based on Wikipedia infoboxes (structured data). An infobox is a factual table with a number of fields

Bernard Keen	
Born	September 5, 1890
Died	August 5, 1981 (aged 90)
Nationality	British
Awards	Fellow of the Royal Society ^[1]
Scientific career	
Fields	Soil scientist
Institutions	University College London

Target Output:

Sir Bernard Augustus Keen FRS (5 September 1890 – 5 August 1981) was a British soil scientist and Fellow of University College London.

Figure 1: An example infobox from the WIKIBIO dataset and the corresponding target output description.

Passage: Hydrogen is commonly used in power stations as a coolant in generators due to a number of favorable properties that are a direct result of its light diatomic molecules.

Answer: as a coolant in generators

Question: How is hydrogen used at power stations?

Table 1: A sample (passage, answer, question) triple from the SQuAD dataset.

(e.g., name, nationality, and occupation) describing a person. For this task, we use the WIKIBIO dataset (Lebret et al., 2016) as the benchmark dataset. Figure 1 shows an example of a biographic infobox as well as the target output textual description.

Automatic question generation aims to generate a syntactically correct, semantically meaningful and relevant question from a natural language text and a target answer within it (unstructured data). This is a crucial yet challenging task in NLP that has received growing attention due to its application in improving question answering systems (Duan et al., 2017; Tang et al., 2017, 2018), providing material for educational purposes (Heilman and

Smith, 2010), and helping conversational systems to start and continue a conversation (Mostafazadeh et al., 2016). We adopt the widely used SQuAD dataset (Rajpurkar et al., 2016) for this task. Table 1 presents a sample (passage, answer, question) triple from this dataset.

Prior work has made remarkable progress on both of these tasks. However, the proposed models utilize complex neural architectures to capture necessary information from the input(s). In this paper, we question the need for such sophisticated NN models for text generation from inputs comprising structured and unstructured data.

Specifically, we adopt a bi-directional, attention-based seq2seq model (Bahdanau et al., 2015) equipped with a copy mechanism (Gu et al., 2016) for both tasks. We demonstrate that this model, together with the exponential moving average (EMA) technique, achieves the state of the art in both neural table-to-text generation and NQG. Interestingly, our model is able to achieve this result even without using any linguistic features.

Our contributions are two-fold: First, we propose a *unified* NN model for text generation from structured and unstructured data and show that training this model with the EMA technique leads to the state of the art in neural table-to-text generation as well as NQG. Second, because our model is, in essence, the primary building block of previous models, our results show that some previous papers propose needless complexity, and that gains from these previous complex neural architectures are quite modest. In other words, the state of the art is achieved by careful tuning of simple and well-engineered models, not necessarily by adding more complexity to the model, echoing the sentiments of Lipton and Steinhardt (2018).

2 Related Work

In this section, we first discuss previous work for neural table-to-text generation and then NQG.

2.1 Neural Table-to-Text Generation

Recently, there have been a number of end-to-end trainable NN models for table-to-text generation. Lebret et al. (2016) propose an n-gram statistical language model that incorporates field and position embeddings to represent the structure of a table. However, their model is not effective enough to capture long-range contextual dependencies while generating a description for the table.

To address this issue, Liu et al. (2018) suggest a structure-aware seq2seq model with local and global addressing on the table. While local addressing is realized by content encoding of the model’s encoder and word-level attention, global addressing is accomplished by field encoding using a field-gating LSTM and field-level attention. The field-gating mechanism incorporates field information when updating the cell memory of the LSTM units.

Liu et al. (2019b) utilize a two-level hierarchical encoder with coarse-to-fine attention to model the field-value structure of a table. They also propose three joint tasks (sequence labeling, text auto-encoding, and multi-label classification) as auxiliary supervision to capture accurate semantic representations of the tables.

In this paper, similar to Lebret et al. (2016), we use both content and field information to represent a table by concatenating the field and position embeddings with the word embedding. Unlike Liu et al. (2018), we don’t separate local and global addressing by using specific modules for each, but rather adopt the EMA technique and let the bi-directional model accomplish this implicitly, exploiting the natural advantages of the model.

2.2 Neural Question Generation

Previous NQG models can be classified into rule-based and neural-network-based approaches. Du et al. (2017) propose a seq2seq model that is able to achieve better results than previous rule-based systems without taking the target answer into consideration. Zhou et al. (2017) concatenate answer position indicators with the word embeddings to make the model aware of the target answer. They also use lexical features (e.g., POS and NER tags) to enrich their model’s encoder. In addition, Song et al. (2018) suggest using a multi-perspective context matching algorithm to further leverage information from explicit interactions between the passage and the target answer.

More recently, Kim et al. (2019) use answer-separated seq2seq, which replaces the target answer in the passage with a unique token to avoid using the answer words in the generated question. They also make use of a module called keyword-net to extract critical information from the target answer. Similarly, Liu et al. (2019a) propose using a clue word predictor by adopting graph convolution networks to highlight the imperative aspects of the input passage.

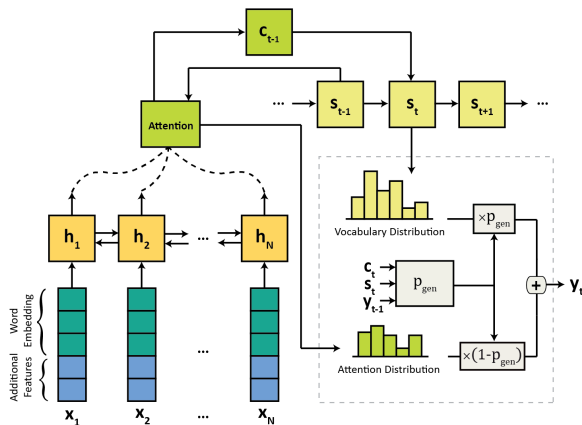


Figure 2: An overview of our model.

Our model is architecturally more similar to Zhou et al. (2017), but with the following distinctions: (1) we do not use additional lexical features, (2) we utilize the EMA technique during training and use the averaged weights for evaluation, (3) we do not make use of the introduced maxout hidden layer, and (4) we adopt LSTM units instead of GRU units. These distinctions, along with some hyperparameter differences, notably the optimizer and learning rate, have a considerable impact on the experimental results (see Section 5).

3 Model: Seq2Seq with Attention and a Copy Mechanism

In this section, we introduce a simple but effective attention-based seq2seq model for both neural table-to-text generation and NQG. Figure 2 provides an overview of our model.

3.1 Encoder

Our encoder is a bi-directional LSTM (BiLSTM) whose input x_t at time step t is the concatenation of the current word embedding e_t with some additional task-specific features.

For neural table-to-text generation, additional features are field name f_t and position information p_t , following Lebre et al. (2016). The position information itself is the concatenation of p_t^+ , which is the position of the current word in its field when counting from the left, and p_t^- , when counting from the right. Considering the word *University*, in Figure 1, as an example, it is the first word from the left and the third word from the right in the *Institutions* field. Hence, the structural information of this word would be $\{Institutions, 1, 3\}$. Thus, the input to the encoder at time step t for this task is

$x_t = [e_t; f_t; p_t^+; p_t^-]$, where $[\cdot; \cdot]$ denotes concatenation along the feature dimension.

For NQG, similar to Zhou et al. (2017), we use a single bit b_t , indicating whether the t^{th} word in the passage belongs to the target answer, as an additional feature. Hence, the input at time step t is $x_t = [e_t; b_t]$. Remarkably, unlike previous work (Song et al., 2018; Kim et al., 2019), we do not use a separate encoder for the target answer to have a *unified* model for both tasks.

3.2 Attention-Based Decoder

Our decoder is an attention-based LSTM model (Bahdanau et al., 2015). Due to the considerable overlap between input and output words, we use a copy mechanism (Gu et al., 2016) that integrates the attention distribution over the input words with the vocabulary distribution.

3.3 Exponential Moving Average

The exponential moving average (EMA) technique, also referred to as temporal averaging, was initially introduced to be used in optimization algorithms for better generalization performance and reducing noise from stochastic approximation in recent parameter estimates by averaging model parameters (Polyak and Juditsky, 1992; Moulines and Bach, 2011; Kingma and Ba, 2015).

In applying the technique, we maintain two sets of parameters: (1) training parameters θ that are trained as usual, and (2) evaluation parameters $\bar{\theta}$ that are an exponentially weighted moving average of the training parameters. The moving average is calculated using the following expression:

$$\bar{\theta} \leftarrow \beta \times \bar{\theta} + (1 - \beta) \times \theta \quad (1)$$

where β is the decay rate. Previous work (Szegedy et al., 2016; Merity et al., 2018; Adhikari et al., 2019; Liu et al., 2019a) has used this technique for different tasks to produce more stable and accurate results. In Section 5, we show that using this simple technique considerably improves the performance of our model in both of the tasks.

4 Experimental Setup

In this section, we introduce the datasets first, then explain additional implementation details, and finally describe the evaluation metrics.

4.1 Datasets

We use the WIKIBIO dataset (Lebre et al., 2016) for neural table-to-text generation. This dataset

contains 728,321 articles from English Wikipedia and uses the first sentence of each article as the ground-truth description of the corresponding infobox. The dataset has been divided into training (80%), validation (10%), and test (10%) sets.

For NQG, we use the SQuAD dataset v1.1 (Rajpurkar et al., 2016) in our experiments, containing 536 Wikipedia articles with over 100K question-answer pairs. The test set of the original dataset is not publicly available. Thus, Du et al. (2017) and Zhou et al. (2017) re-divide available data into training, validation, and test sets, which we call split-1 and split-2, respectively. In this paper, we conduct experiments and evaluate our model on both of the data splits.

4.2 Implementation Details

For the sake of reproducibility, we discuss implementation details for achieving the results shown in Tables 2 and 3. We train the model using cross-entropy loss and retain the model that works best on the validation set during training for both tasks. We replace unknown tokens with a word from the input having the highest attention score. In addition, a decay rate of 0.9999 is used for the exponential moving average in both of the tasks.

For the neural table-to-text generation task, we train the model up to 10 epochs with three different seeds and a batch size of 32. We use a single-layer BiLSTM for the encoder and a single-layer LSTM for the decoder and set the dimension of the LSTM hidden states to 500. Optimization is performed using the Adam optimizer with a learning rate of 0.0005 and gradient clipping when its norm exceeds 5. The word, field, and position embeddings are trainable and have a dimension of 400, 50, and 5, respectively. The maximum position number is set to 30. Any higher position number is therefore counted as 30. The most frequent 20,000 words and 1,480 fields in the training set are selected as word vocabulary and field vocabulary, respectively, for both the encoder and the decoder. Ultimately, we conduct greedy search to decode a description for a given input table.

For the NQG task, we use a two-layer BiLSTM for the encoder and a single-layer LSTM for the decoder. We set the dimension of the LSTM hidden states to 350 and 512 for split-1 and split-2, respectively. Optimization is performed using the AdaGrad optimizer with a learning rate of 0.3 and gradient clipping when its norm exceeds 5. The

word embeddings are initialized with pre-trained 300-dimensional GloVe embeddings (Pennington et al., 2014), which are frozen during training. We train the model up to 20 epochs with five different seeds and a batch size of 50. We further employ dropout with a probability of 0.1 and 0.3 for data split-1 and split-2, respectively. Moreover, we use the vocabulary set released by Song et al. (2018) for both the encoder and the decoder. During decoding, we perform beam search with a beam size of 20 and a length penalty weight of 1.75.

4.3 Evaluation

Following previous work, we use BLEU-4 (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), ROUGE-4, and ROUGE-L (Lin, 2004) to evaluate the performance of our model. BLEU and METEOR were originally designed to evaluate machine translation systems, and ROUGE was designed to evaluate text summarization systems.

5 Results and Discussion

In this section, we present our experimental results for both neural table-to-text generation and NQG. We report the mean and standard deviation of each metric across multiple seeds to ensure robustness against potentially spurious conclusions (Crane, 2018). In Tables 2 and 3, we compare previous work with our results for NQG and neural table-to-text generation, respectively. All results are copied from the original papers except for Liu et al. (2018) in Table 3, where *Repl.* refers to scores from experiments that we conducted using the source code released by the authors, and *Orig.* refers to scores taken from the original paper.

It is noteworthy that a similar version of our model has served as a baseline in previous papers (Liu et al., 2018; Kim et al., 2019; Liu et al., 2019a). However, the distinctions discussed in Section 2, especially the EMA technique, enable our model to achieve the state of the art in all cases but BLEU-4 on the SQuAD split-2, where our score is very competitive; furthermore, Liu et al. (2019a) only report results from a single trial. Our results indicate that a basic seq2seq model is able to effectively learn the underlying distribution of both datasets.

6 Conclusions and Future Work

In this paper, we question the necessity of complex neural architectures for text generation from structured data (neural table-to-text generation) and

Models	Split-1			Split-2		
	BLEU-4	METEOR	ROUGE-L	BLEU-4	METEOR	ROUGE-L
Heilman (2011)	-	-	-	9.47	18.97	31.68
Du et al. (2017)	12.28	16.62	39.75	-	-	-
Zhou et al. (2017)	-	-	-	13.29	-	-
Zhou et al. (2018)	-	-	-	13.02	-	44.0
Yao et al. (2018)	-	-	-	13.36	17.70	40.42
Song et al. (2018)	13.98	18.77	42.72	13.91	-	-
Zhao et al. (2018)	15.32	19.29	43.91	15.82	19.67	44.24
Sun et al. (2018)	-	-	-	15.64	-	-
Kumar et al. (2018)	16.17	19.85	43.90	-	-	-
Kim et al. (2019)	16.20 ± 0.32	19.92 ± 0.20	43.96 ± 0.25	16.17 ± 0.35	-	-
Liu et al. (2019a)	-	-	-	17.55	21.24	44.53
Our Model	14.81 ± 0.47	19.69 ± 0.24	43.01 ± 0.28	16.14 ± 0.25	20.44 ± 0.20	43.95 ± 0.26
+ EMA	16.29 ± 0.04	20.70 ± 0.08	44.18 ± 0.15	17.47 ± 0.10	21.37 ± 0.06	45.18 ± 0.22

Table 2: Experimental results for NQG on the test sets.

Models	BLEU-4	ROUGE-4
KN*	2.21	0.38
Template KN**	19.80	10.70
Lebret et al. (2016)	34.70 ± 0.36	25.80 ± 0.36
Bao et al. (2018)	40.26	-
Sha et al. (2018)	43.91	37.15
Liu et al. (2018) Orig.	44.89 ± 0.33	41.21 ± 0.25
Liu et al. (2018) Repl.	44.45 ± 0.11	39.65 ± 0.10
Liu et al. (2019b)	45.14 ± 0.34	41.26 ± 0.37
Our Model	46.07 ± 0.17	41.53 ± 0.30
+ EMA	46.76 ± 0.03	43.54 ± 0.07

Table 3: Experimental results for neural table-to-text generation on the test set. *KN is Kneser-Ney language model (Heafield et al., 2013). **Template KN is a KN language model over templates. Both models are proposed by Lebret et al. (2016) as baselines.

unstructured data (NQG). We then propose a simple yet effective seq2seq model trained with the EMA technique. Empirically, our model achieves the state of the art in both of the tasks. Our results highlight the importance of thoroughly exploring simple models before introducing complex neural architectures, so that we can properly attribute the source of performance gains. As a potential direction for future work, it would be interesting to investigate the use of the EMA technique on transformer models as well and conduct similar studies to examine needless architectural complexity in other NLP tasks.

Acknowledgments

This research was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada.

References

- Ashutosh Adhikari, Achyudh Ram, Raphael Tang, and Jimmy Lin. 2019. Rethinking complex neural network architectures for document classification. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4046–4051.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of the 3rd International Conference on Learning Representations*.
- Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72.
- Junwei Bao, Duyu Tang, Nan Duan, Zhao Yan, Yuanhua Lv, Ming Zhou, and Tiejun Zhao. 2018. Table-to-Text: Describing table region with natural language. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, pages 5020–5027.
- Matt Crane. 2018. Questionable answers in question answering research: Reproducibility and variability of published results. *Transactions of the Association of Computational Linguistics*, 6:241–252.
- Xinya Du, Junru Shao, and Claire Cardie. 2017. Learning to ask: Neural question generation for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1342–1352.
- Nan Duan, Duyu Tang, Peng Chen, and Ming Zhou. 2017. Question generation for question answering.

- In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 866–874.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O.K. Li. 2016. Incorporating copying mechanism in sequence-to-sequence learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1631–1640.
- Kenneth Heafield, Ivan Pouzyrevsky, Jonathan H. Clark, and Philipp Koehn. 2013. Scalable modified Kneser-Ney language model estimation. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 690–696.
- Michael Heilman. 2011. *Automatic Factual Question Generation from Text*. Ph.D. thesis, Carnegie Mellon University.
- Michael Heilman and Noah A. Smith. 2010. Good question! Statistical ranking for question generation. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 609–617.
- Yanghoon Kim, Hwanhee Lee, Joongbo Shin, and Kyomin Jung. 2019. Improving neural question generation using answer separation. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, pages 6602–6609.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *arXiv:1412.6980*.
- Vishwajeet Kumar, Ganesh Ramakrishnan, and Yuanfang Li. 2018. A framework for automatic question generation from text using deep reinforcement learning. *arXiv:1808.04961*.
- Rémi Lebreton, 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*, pages 1203–1213.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain.
- Zachary C. Lipton and Jacob Steinhardt. 2018. Troubling trends in machine learning scholarship. *arXiv:1807.03341v2*.
- Bang Liu, Mingjun Zhao, Di Niu, Kunfeng Lai, Yancheng He, Haojie Wei, and Yu Xu. 2019a. Learning to generate questions by learning what not to generate. *arXiv:1902.10418*.
- Tianyu Liu, Fuli Luo, Qiaolin Xia, Shuming Ma, Baobao Chang, and Zhifang Sui. 2019b. Hierarchical encoder with auxiliary supervision for neural table-to-text generation: Learning better representation for tables. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, pages 6786–6793.
- 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*, pages 4881–4888.
- Gábor Melis, Chris Dyer, and Phil Blunsom. 2018. On the state of the art of evaluation in neural language models. In *Proceedings of the 6th International Conference on Learning Representations*.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2018. Regularizing and optimizing LSTM language models. In *Proceedings of the 6th International Conference on Learning Representations*.
- Salman Mohammed, Peng Shi, and Jimmy Lin. 2018. Strong baselines for simple question answering over knowledge graphs with and without neural networks. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 291–296.
- Nasrin Mostafazadeh, Ishan Misra, Jacob Devlin, Margaret Mitchell, Xiaodong He, and Lucy Vanderwende. 2016. Generating natural questions about an image. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1802–1813.
- Eric Moulines and Francis R. Bach. 2011. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *Advances in Neural Information Processing Systems*, pages 451–459.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318.
- 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.
- Boris T. Polyak and Anatoli B. Juditsky. 1992. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392.

- Lei Sha, Lili Mou, Tianyu Liu, Pascal Poupart, Sujian Li, Baobao Chang, and Zhifang Sui. 2018. Order-planning neural text generation from structured data. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, pages 5414–5421.
- Linfeng Song, Zhiguo Wang, Wael Hamza, Yue Zhang, and Daniel Gildea. 2018. Leveraging context information for natural question generation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 569–574.
- Xingwu Sun, Jing Liu, Yajuan Lyu, Wei He, Yanjun Ma, and Shi Wang. 2018. Answer-focused and position-aware neural question generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3930–3939.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2818–2826.
- Duyu Tang, Nan Duan, Tao Qin, Zhao Yan, and Ming Zhou. 2017. Question answering and question generation as dual tasks. *arXiv:1706.02027*.
- Duyu Tang, Nan Duan, Zhao Yan, Zhirui Zhang, Yibo Sun, Shujie Liu, Yuanhua Lv, and Ming Zhou. 2018. Learning to collaborate for question answering and asking. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1564–1574.
- Kaichun Yao, Libo Zhang, Tiejian Luo, Lili Tao, and Yanjun Wu. 2018. Teaching machines to ask questions. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 4546–4552.
- Yao Zhao, Xiaochuan Ni, Yuanyuan Ding, and Qifa Ke. 2018. Paragraph-level neural question generation with maxout pointer and gated self-attention networks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3901–3910.
- Qingyu Zhou, Nan Yang, Furu Wei, Chuanqi Tan, Hangbo Bao, and Ming Zhou. 2017. Neural question generation from text: A preliminary study. In *National CCF Conference on Natural Language Processing and Chinese Computing*, pages 662–671. Springer.
- Qingyu Zhou, Nan Yang, Furu Wei, and Ming Zhou. 2018. Sequential copying networks. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, pages 4987–4995.