

# Knowledge-Enriched Natural Language Generation

Wenhao Yu<sup>1</sup>, Meng Jiang<sup>1</sup>, Zhiting Hu<sup>2</sup>, Qingyun Wang<sup>3</sup>, Heng Ji<sup>3</sup>, Nazneen Rajani<sup>4</sup>

<sup>1</sup>University of Notre Dame, <sup>2</sup>University of California, San Diego,

<sup>3</sup>University of Illinois Urbana-Champaign, <sup>4</sup>Salesforce Research

{wyu1, mjiang2}@nd.edu, zhitinghu@gmail.com,

{qingyun4, hengji}@illinois.edu, nazneen.rajani@salesforce.com

## 1 Introduction

Natural Language Generation (NLG) aims at deliberately constructing a natural language text in order to meet specified communicative goals. NLG has been applied in many real-world applications, including dialogue systems, biography generation, technical paper draft generation, and multimedia news summarization. Neural language models have achieved impressive successes at automatic NLG, especially on creative writing such as story completion and poetry generation. However, in many downstream applications such as news summarization, counter-argument narrative generation, and knowledge base description, the generation process needs to be guided by certain level of knowledge such as sentiment (Hu et al., 2017), topic (Xing et al., 2017), and style (Tikhonov et al., 2019).

The usage of supportive knowledge in NLG can be roughly divided into the following two levels: (1) knowledge description (KD), which transforms structured data into unstructured text, such as topic-to-text (Dong et al., 2021; Yu et al., 2021), knowledge base description (Gardent et al., 2017; Liu et al., 2018a; Qin et al., 2019; Zeng et al., 2021), table-to-text generation (Liu et al., 2018b; Moryossef et al., 2019; Wang et al., 2020) and its variants in low-resource (Ma et al., 2019) and multi-lingual setting (Kaffee et al., 2018), data-to-text (Wiseman et al., 2017; Puduppully et al., 2019), and graph-to-text (Song et al., 2018; Zhu et al., 2019; Yao et al., 2020); (2) knowledge synthesis (KS), which obtain knowledge from external knowledge resources (e.g. knowledge base) and integrate it into text generation, such as image or video caption generation (Whitehead et al., 2018; Lu et al., 2018), knowledge graph-supported dialogue generation (Liu et al., 2019; Zhang et al., 2020), knowledge-guided comment generation (Li et al., 2019), and scientific paper generation (Wang

et al., 2019; Koncel-Kedziorski et al., 2019).

Knowledge-enriched text generation poses unique challenges in modeling and learning, driving active research in several core directions, ranging from integrated modeling of neural representations and symbolic information in the sequential/hierarchical/graphical structures, learning without direct supervisions due to the cost of structured annotation, efficient optimization and inference with massive and global constraints, to language grounding on multiple modalities, and generative reasoning with implicit commonsense knowledge and background knowledge. In this tutorial we will present a roadmap to line up the state-of-the-art methods to tackle these challenges on this **cutting-edge** problem. We will dive deep into various technical components (as shown in Figure 1): how to represent knowledge, how to feed knowledge into a generation model, how to evaluate generation results, and what are the remaining challenges?

## 2 Brief Tutorial Outline

### 2.1 Motivation and Overview [20 mins]

At the beginning of the tutorial we will motivate the task of knowledge-driven NLG by showing a large variety of applications (e.g., KD and KS) in academia and industry which have been mentioned in the Introduction. We will present examples about the shortcomings of pure Seq2Seq or language models as well as the opportunities of using knowledge to enrich the generation. We categorize the input source knowledge and related advanced machine learning technologies in Figure 1. We will present the overview of this tutorial including language models (LMs) and knowledge representation, general learning and generation frameworks, a variety of NLG methods enriched by knowledge sources including semantics and structures, real-world applications, and discussions.

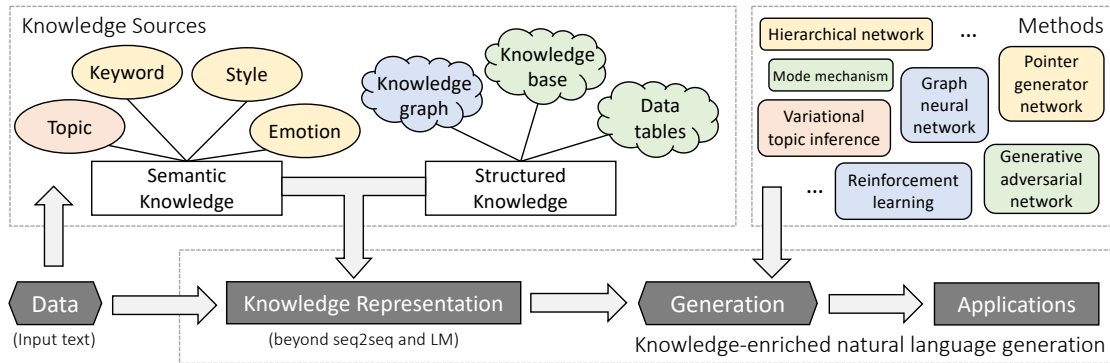


Figure 1: In this tutorial, we will present advanced NLG methods that inject knowledge from a variety of sources.

## 2.2 General Learning and Generation Frameworks [40 mins]

We will present the general methods of knowledge-enriched NLG, which provide the methodological foundations for incorporating different types of knowledge presented in the subsequent parts. Those methods are categorized into three major paradigms which incorporate knowledge through (1) *model architectures* that facilitate the use of knowledge, such as attention methods, copy/pointer mechanisms, graph neural networks (GNNs), knowledge-enriched embedding, etc; (2) *learning frameworks* that inject knowledge information into the generation models through training, such as posterior regularization, constraint-driven learning, semantic loss, knowledge-informed weak supervision, etc; (3) *inference methods* which imposes on the inference process different knowledge constraints to guide decoding, such as lexical constraints, task-specific objectives, global inter-dependency, etc.

### 2.3 NLG Methods Enhanced by Various Knowledge Sources: Part I [30 mins]

In this part, we present *semantic knowledge*-driven natural language generation. The semantic knowledge sources mainly contain keywords, topics, linguistic features, and other semantic constraints (e.g., style, emotion, sentiment). We introduce how the knowledge in each source can be encoded and how the represented knowledge can be decoded into natural language of high quality.

### 2.4 Coffee Break [30 mins]

### 2.5 NLG Methods Enhanced by Various Knowledge Sources: Part II [30 mins]

In this part, we present *structured knowledge*-driven natural language generation. The struc-

tured knowledge sources mainly contain tables, knowledge bases, and knowledge graphs. We introduce how the knowledge in each source can be represented and integrated into generation frameworks. Then, we introduce the methods that (i) find relevant knowledge (e.g., a relational path) from huge knowledge bases and knowledge graphs, and (ii) construct structured knowledge from text, e.g., OpenIE. Lastly, we introduce recent work that integrates multiple types of knowledge ranging from semantic, unstructured, to structured knowledge.

We will give a review of the available structured knowledge representation method, most of which focus on the structured tables. Traditionally, researchers tend to linearize the table for the input with the concatenation of type information. With computational advances in recent years, pre-trained language model based approaches for the linearized input have achieved significant success by combining type information as additional position embedding. However, those methods fail to consider the inter-dependency between different entities. We will discuss two major ways to learn those relations: self-attention mechanism and GNNs.

### 2.6 Applications [30 mins]

In this application session, we first review existing *potential applications* using the knowledge-driven generations. On one hand, the structured knowledge provides additional guidance for the major tasks such as dialogue systems, video captions, and summarizations. On the other hand, researchers have built independent knowledge guided generation tasks, starting from the data-to-text tasks such as Wikibio generation tasks in low-resource and multilingual setting, Webnlg contests, and ROTOWIRE, to more complex graph-to-text tasks such as AMR-to-text generation, scientific paper

generation tasks, and news comment generation. Then, we will cover various post-processing approaches to enhance the quality of generation results for specific applications, such as coverage mechanism, self-attention mechanism, and table-text optimal-transport matching loss. Finally we will briefly present how knowledge-enriched NLG is being used in several conversational AI systems including Amazon Alexa. Other commercial applications for NLG include systems that can retrieve and summarize information from a relational database into natural language text such as Salesforce’s Einstein and Tableau.

## 2.7 Remaining Challenges and Future Directions [30 mins]

At the end of the tutorial we will discuss the remaining challenges and some of the future directions, including the challenge of capturing the interdependency of knowledge elements to make generated output coherent, knowledge reasoning, representing time and number, duplicate removal, augmenting massive pre-trained language models with external commonsense and background knowledge, and developing effective automatic evaluation metrics, and rigorous and efficient human evaluation procedures. We will provide pointers to resources, including data sets, software and on-line demos.

## 3 Diversity Considerations

The topic to be presented is of great interest to diverse group of audience from academics and industry. We will cover a broad diversity of methods and applications in many languages and domains. In particular, enriching modeling and learning with external knowledge, as the core topic in this tutorial, is particularly helpful for low-resource language modeling where no large data are available.

We have a diverse instructor team across multiple institutions (ND, UIUC, UCSD, and Salesforce Inc.) with varying seniority (ranging from junior/senior PhD students to assistant/full professors and senior researchers), two of whom are female researchers. The team has a diverse and broad expertise in natural language processing and generation, machine learning, data mining, and various application domains.

## 4 Prerequisites

This tutorial will present basic and advanced methods in NLG systematically to audience. The audi-

ence may find different useful content when have different levels of prior knowledge: (with the number of ☆ for how much a person may feel comfortable and confident with the subject matter)

- Familiar with Machine Learning from text, e.g., “Understand classification tasks and classical supervised methods on text data” (☆);
- Familiar with basic natural language processing (NLP) frameworks, e.g., “Had experience with LSTM, Seq2Seq, transformer” (☆☆);
- Familiar with some data forms of knowledge, e.g., “Had machine learning experience with topic modeling, knowledge bases, knowledge graphs, data tables, etc.” (☆☆☆).

## 5 Reading List

Full reading list:

<https://github.com/wyu97/KENLG-Reading>

Small reading list:

- Survey: KENLG (Yu et al., 2020)
- General learning and NLG frameworks
  - (1) Seq2Seq (Bahdanau et al., 2015),
  - (2) Transformer (Vaswani et al., 2017),
  - (3) Copy mechanism (Gu et al., 2016);
- Semantic knowledge for enhancing NLG
  - (4) Topic (Xing et al., 2017),
  - (5) Sentiment (Hu et al., 2017),
  - (6) Emotion (Zhou et al., 2018a);
- Structured knowledge for enhancing NLG
  - (7) Wikipedia KB (Liu et al., 2018b),
  - (8) Sports Tables (Wiseman et al., 2017),
  - (9) Commonsense KG (Zhou et al., 2018b),
  - (10) Scientific KG (Koncel et al., 2019).

## 6 Presenters

**Wenhao Yu** is a Ph.D. student in the Department of Computer Science and Engineering at the University of Notre Dame. His research lies in controllable knowledge-driven natural language processing, particularly in natural language generation. His research has been published in top-ranked NLP and data mining conferences such as ACL, EMNLP, AAAI, WWW, and CIKM. Additional information is available at <https://wyu97.github.io/>

**Meng Jiang** is an assistant professor in the Department of Computer Science and Engineering at the University of Notre Dame. He received his B.E. and Ph.D. in Computer Science from Tsinghua University and was a postdoctoral research associate at the University of Illinois at Urbana-Champaign. His research interests focus on knowledge graph

construction and natural language generation for news summarization and forum post generation. The awards he received include Notre Dame Faculty Award in 2019 and Best Paper Awards at ISDSA and KDD-DLG in 2020. Additional information is available at <http://www.meng-jiang.com/>.

**Zhiting Hu** is an assistant professor in Halicioğlu Data Science Institute at UC San Diego. He received his Ph.D. in Machine Learning from Carnegie Mellon University. His research interest lies in the broad area of natural language processing in particular controllable text generation, machine learning to enable training AI agents from all forms of experiences such as structured knowledge, ML systems and applications. His research was recognized with best demo nomination at ACL 2019 and outstanding paper award at ACL 2016. Additional information is available at <http://www.cs.cmu.edu/~zhitingh/>.

**Qingyun Wang** is a Ph.D. student in the Computer Science Department at the University of Illinois at Urbana-Champaign. His research lies in controllable knowledge-driven natural language generation, with a recent focus on the scientific paper generation. He served as a program committee in generation track for multiple conferences including ICML 2020, ACL 2019-2020, ICLR 2021, etc. He previously entered the finalist of the first Alexa Prize competition. Additional information is available at <https://eaglew.github.io/>

**Heng Ji** is a professor at Computer Science Department of University of Illinois at Urbana-Champaign, and Amazon Scholar. She has published on Multimedia Multilingual Information Extraction and Knowledge-enriched NLG including technical paper generation, knowledge base description, and knowledge-aware image and video caption generation. The awards she received include “Young Scientist” by World Economic Forum, “AI’s 10 to Watch” Award by IEEE Intelligent Systems, NSF CAREER award, and ACL 2020 Best Demo Award. She has served as the Program Committee Co-Chair of many conferences including NAACL-HLT2018, and she is NAACL secretary 2020-2021. Additional information is available at <https://blender.cs.illinois.edu/hengji.html>.

**Nazneen Rajani** is a senior research scientist at Salesforce Research. She got her PhD in Computer Science from UT Austin in 2018. Several of her work has been published in ACL, EMNLP,

NACCL, and IJCAI including work on generating explanations for commonsense and physical reasoning. Nazneen was one of the finalists for the VentureBeat Transform 2020 women in AI Research. Her work has been covered by several media outlets including Quanta Magazine, VentureBeat, SiliconAngle, ZDNet. More information on <https://www.nazneenrajani.com>

## 6.1 Selected Past Tutorials

### Heng Ji:

- ACL’18 and CCL’18: Multi-lingual Entity Discovery and Linking
- SIGMOD’16: Automatic Entity Recognition and Typing in Massive Text Data.
- ACL’15: Successful Data Mining Methods for NLP.
- ACL’14 and NLPCC’14: Wikification and Beyond: The Challenges of Entity and Concept Grounding.
- COLING’12: Temporal Information Extraction and Shallow Temporal Reasoning.

### Meng Jiang:

- KDD’20: Scientific Text Mining and Knowledge Graphs.
- KDD’20: Multi-modal Network Representation Learning: Methods and Applications.
- KDD’17: Mining Entity-Relation-Attribute Structures from Massive Text Data.
- KDD’17: Data-Driven Approaches towards Malicious Behavior Modeling.
- SIGMOD’17: Building Structured Databases of Factual Knowledge from Massive Text.
- WWW’17: Constructing Structured Information Networks from Massive Text Corpora.

### Zhiting Hu:

- KDD’20: Learning from All Types of Experiences: A Unifying Machine Learning Perspective.
- AACL’20: Modularizing Natural Language Processing.

## References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference for Learning Representation (ICLR)*.
- Xiangyu Dong, Wenhao Yu, Chenguang Zhu, and Meng Jiang. 2021. Injecting entity types into entity-guided text generation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017. [Creating training corpora for NLG micro-planners](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 179–188, Vancouver, Canada. Association for Computational Linguistics.
- 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, Berlin, Germany. Association for Computational Linguistics.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P. Xing. 2017. Toward controlled generation of text. In *International Conference on Machine Learning (ICML)*.
- Lucie-Aimée Kaffee, Hady Elsahar, Pavlos Vougiouklis, Christophe Gravier, Frédérique Laforest, Jonathon Hare, and Elena Simperl. 2018. [Learning to generate Wikipedia summaries for underserved languages from Wikidata](#). 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 640–645, New Orleans, Louisiana. Association for Computational Linguistics.
- Rik Koncel, Kedzierski, Dhanush Bekal, Yi Luan, Mirella Lapata, and Hannaneh Hajishirzi. 2019. Text generation from knowledge graphs with graph transformers. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*.
- Rik Koncel-Kedzierski, Dhanush Bekal, Yi Luan, Mirella Lapata, and Hannaneh Hajishirzi. 2019. [Text Generation from Knowledge Graphs with Graph Transformers](#). 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 2284–2293, Minneapolis, Minnesota. Association for Computational Linguistics.
- Wei Li, Jingjing Xu, Yancheng He, ShengLi Yan, Yunfang Wu, and Xu Sun. 2019. [Coherent comments generation for Chinese articles with a graph-to-sequence model](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4843–4852, Florence, Italy. Association for Computational Linguistics.
- Shuman Liu, Hongshen Chen, Zhaochun Ren, Yang Feng, Qun Liu, and Dawei Yin. 2018a. [Knowledge diffusion for neural dialogue generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1489–1498, Melbourne, Australia. Association for Computational Linguistics.
- Tianyu Liu, Kexiang Wang, Lei Sha, Baobao Chang, and Zhifang Sui. 2018b. Table-to-text generation by structure-aware seq2seq learning. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- Zhibin Liu, Zheng-Yu Niu, Hua Wu, and Haifeng Wang. 2019. [Knowledge aware conversation generation with explainable reasoning over augmented graphs](#). 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 1782–1792, Hong Kong, China. Association for Computational Linguistics.
- Di Lu, Spencer Whitehead, Lifu Huang, Heng Ji, and Shih-Fu Chang. 2018. Entity-aware image caption generation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Shuming Ma, Pengcheng Yang, Tianyu Liu, Peng Li, Jie Zhou, and Xu Sun. 2019. [Key fact as pivot: A two-stage model for low resource table-to-text generation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2047–2057, Florence, Italy. Association for Computational Linguistics.
- Amit Moryossef, Yoav Goldberg, and Ido Dagan. 2019. [Step-by-step: Separating planning from realization in neural data-to-text generation](#). 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 2267–2277, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ratish Puduppully, Li Dong, and Mirella Lapata. 2019. [Data-to-text generation with entity modeling](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2023–2035, Florence, Italy. Association for Computational Linguistics.
- Libo Qin, Yijia Liu, Wanxiang Che, Haoyang Wen, Yangming Li, and Ting Liu. 2019. [Entity-consistent end-to-end task-oriented dialogue system with KB retriever](#). 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 133–142, Hong Kong, China. Association for Computational Linguistics.
- 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 (Volume 1: Long Papers)*, pages 1616–1626, Melbourne, Australia. Association for Computational Linguistics.

- Alexey Tikhonov, Viacheslav Shibaev, Aleksander Nagaev, Aigul Nugmanova, and Ivan P. Yamshchikov. 2019. [Style transfer for texts: Retrain, report errors, compare with rewrites](#). 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 3936–3945, Hong Kong, China. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems (Neruips)*. Curran Associates, Inc.
- Qingyun Wang, Lifu Huang, Zhiying Jiang, Kevin Knight, Heng Ji, Mohit Bansal, and Yi Luan. 2019. [PaperRobot: Incremental draft generation of scientific ideas](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1980–1991, Florence, Italy. Association for Computational Linguistics.
- Zhenyi Wang, Xiaoyang Wang, Bang An, Dong Yu, and Changyou Chen. 2020. [Towards faithful neural table-to-text generation with content-matching constraints](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1072–1086, Online. Association for Computational Linguistics.
- Spencer Whitehead, Heng Ji, Mohit Bansal, Shih-Fu Chang, and Clare Voss. 2018. [Incorporating background knowledge into video description generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3992–4001, Brussels, Belgium. Association for Computational Linguistics.
- Sam Wiseman, Stuart Shieber, and Alexander Rush. 2017. [Challenges in data-to-document generation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2253–2263, Copenhagen, Denmark. Association for Computational Linguistics.
- Chen Xing, Wei Wu, Yu Wu, Jie Liu, Yalou Huang, Ming Zhou, and Wei-Ying Ma. 2017. Topic aware neural response generation. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- Shaowei Yao, Tianming Wang, and Xiaojun Wan. 2020. Heterogeneous graph transformer for graph-to-sequence learning. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Wenhao Yu, Chenguang Zhu, Zaitang Li, Zhiting Hu, Qingyun Wang, Heng Ji, and Meng Jiang. 2020. A survey of knowledge-enhanced text generation. *arXiv preprint arXiv:2010.04389*.
- Wenhao Yu, Chenguang Zhu, Tong Zhao, Zhichun Guo, and Meng Jiang. 2021. Sentence-permuted paragraph generation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Qingkai Zeng, Jinfeng Lin, Wenhao Yu, Jane Cleland-Huang, and Meng Jiang. 2021. Enhancing taxonomy completion with concept generation via fusing relational representations. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*.
- Houyu Zhang, Zhenghao Liu, Chenyan Xiong, and Zhiyuan Liu. 2020. [Grounded conversation generation as guided traverses in commonsense knowledge graphs](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2031–2043, Online. Association for Computational Linguistics.
- Hao Zhou, Minlie Huang, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. 2018a. Emotional chatting machine: Emotional conversation generation with internal and external memory. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018b. Commonsense knowledge aware conversation generation with graph attention. In *International Joint Conference on Artificial Intelligence (IJCAI)*.
- Jie Zhu, Junhui Li, Muhua Zhu, Longhua Qian, Min Zhang, and Guodong Zhou. 2019. [Modeling graph structure in transformer for better AMR-to-text generation](#). 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 5459–5468, Hong Kong, China. Association for Computational Linguistics.