

Stage-wise Fine-tuning for Graph-to-Text Generation

Qingyun Wang^{1*}, Semih Yavuz², Xi Victoria Lin³,
Heng Ji¹, Nazneen Fatema Rajani²

¹ University of Illinois at Urbana-Champaign ² Salesforce Research ³ Facebook AI

syavuz, nazneen.rajani@salesforce.com

victorialin@fb.com

{qingyun4, hengji}@illinois.edu

Abstract

Graph-to-text generation has benefited from pre-trained language models (PLMs) in achieving better performance than structured graph encoders. However, they fail to fully utilize the structure information of the input graph. In this paper, we aim to further improve the performance of the pre-trained language model by proposing a structured graph-to-text model with a two-step fine-tuning mechanism which first fine-tunes the model on Wikipedia before adapting to the graph-to-text generation. In addition to using the traditional token and position embeddings to encode the knowledge graph (KG), we propose a novel tree-level embedding method to capture the inter-dependency structures of the input graph. This new approach has significantly improved the performance of all text generation metrics for the English WebNLG 2017 dataset.¹

1 Introduction

In the graph-to-text generation task (Gardent et al., 2017), the model takes in a complex KG (an example is in Figure 1) and generates a corresponding faithful natural language description (Table 1). Previous efforts for this task can be mainly divided into two categories: sequence-to-sequence models that directly solve the generation task with LSTMs (Gardent et al., 2017) or Transformer (Castro Ferreira et al., 2019); and graph-to-text models (Trisedya et al., 2018; Marcheggiani and Perez-Beltrachini, 2018) which use a graph encoder to capture the structure of the KGs. Recently, Transformer-based PLMs such as GPT-2 (Radford et al., 2019), BART (Lewis et al., 2020),

*This research was conducted during the author’s internship at Salesforce Research.

¹The programs, data and resources are publicly available for research purpose at: <https://github.com/EagleW/Stage-wise-Fine-tuning>

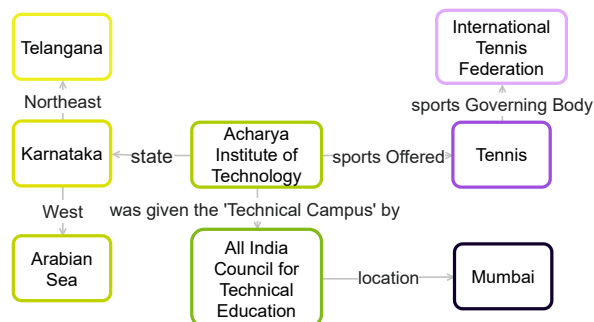


Figure 1: Input RDF Knowledge Graph

and T5 (Raffel et al., 2020) have achieved state-of-the-art results on WebNLG dataset due to factual knowledge acquired in the pre-training phase (Harkous et al., 2020; Ribeiro et al., 2020b; Kale, 2020; Chen et al., 2020a).

Despite such improvement, PLMs fine-tuned only on the clean (or labeled) data might be more prone to hallucinate factual knowledge (e.g., “Visvesvaraya Technological University” in Table 1). Inspired by the success of domain-adaptive pre-training (Gururangan et al., 2020), we propose a novel two-step fine-tuning mechanism graph-to-text generation task. Unlike (Ribeiro et al., 2020b; Herzig et al., 2020; Chen et al., 2020a) which directly fine-tune the PLMs on the training set, we first fine-tune our model over noisy RDF graphs and related article pairs crawled from Wikipedia before final fine-tuning on the clean/labeled training set. The additional fine-tuning step benefits our model by leveraging triples not included in the training set and reducing the chances that the model fabricates facts based on the language model.

Meanwhile, the PLMs might also fail to cover all relations in the KG by creating incorrect or missing facts. For example, in Table 1, although the T5-large with Wikipedia fine-tuning successfully removes the unwanted contents, it still ignores the “sports Governing Body” relation and incorrectly

Category	Output
Reference	The Acharya Institute of Technology in Karnataka state was given Technical Campus status by All India Council for Technical Education in Mumbai . The school offers tennis which is governed by the International Tennis Federation . Karnataka has the Arabian Sea to its west and in the northeast is Telangana .
T5-large	The state of Karnataka is located southwest of Telangana and east of the Arabian Sea . It is the location of the Acharya Institute of Technology which was granted the Technical Campus status by the All India Council for Technical Education in Mumbai . The Institute is affiliated with the Visvesvaraya Technological University and offers the sport of tennis . [International Tennis Federation]
T5-large + Wiki	The Acharya Institute of Technology is located in the state of Karnataka . It was given the Technical Campus status by the All India Council for Technical Education which is located in Mumbai . The institute offers tennis and has Telangana to its northeast and the Arabian Sea to its west. [International Tennis Federation]
T5-large + Position	The Acharya Institute of Technology is located in the state of Karnataka which has Telangana to its northeast and the Arabian Sea to its west. It was given the Technical Campus status by the All India Council for Technical Education in Mumbai . The Institute offers tennis which is governed by the International Tennis Federation .
T5-large + Wiki + Position	The Acharya Institute of Technology in Karnataka was given the "Technical Campus" status by the All India Council for Technical Education in Mumbai . Karnataka has Telangana to its northeast and the Arabian Sea to its west. One of the sports offered at the Institute is tennis which is governed by the International Tennis Federation .

Table 1: Human and System Generated Description in Figure 1. We use the color box to frame each entity out with the same color as the corresponding entity in Figure 1. We highlight *fabricated facts*, *[missed relations]*, and *incorrect relations* with different color.

links the university to both “Telangana” and “Arabian Sea”. To better capture the structure and interdependence of facts in the KG, instead of using a complex graph encoder, we leverage the power of Transformer-based PLMs with additional position embeddings which have been proved effective in various generation tasks (Herzig et al., 2020; Chen et al., 2020a,b). Here, we extend the embedding layer of Transformer-based PLMs with two additional *triple role* and *tree-level* embeddings to capture graph structure.

We explore the proposed stage-wise fine-tuning and structure-preserving embedding strategies for graph-to-text generation task on WebNLG corpus (Gardent et al., 2017). Our experimental results clearly demonstrate the benefit of each strategy in achieving the state-of-the-art performance on most commonly reported automatic evaluation metrics.

2 Method

Given an RDF graph with multiple relations $G = \{(s_1, r_1, o_1), (s_2, r_2, o_2), \dots, (s_n, r_n, o_n)\}$, our goal is to generate a text faithfully describing the input graph. We represent each relation with a triple $(s_i, r_i, o_i) \in G$ for $i \in \{1, \dots, n\}$, where s_i , r_i , and o_i are natural language phrases that represent the subject, type, and object of the relation,

respectively. We augment our model with addi-




Figure 2: Position Embeddings for the KG in Figure 1

tional position embeddings to capture the structure of the KG. To feed the input for the large-scale Transformer-based PLM, we flatten the graph as a concatenation of linearized triple sequences:

$$|S s_1 |P r_1 |O o_1 \dots |S s_n |P r_n |O o_n$$

following Ribeiro et al. (2020b), where $|S$, $|P$, $|O$ are special tokens prepended to indicate whether the phrases in the relations are subjects, relations, or objects, respectively. Instead of directly fine-tuning the PLM on the WebNLG dataset, we first fine-tune our model on a noisy, but larger corpus crawled from Wikipedia, then we fine-tune the model on the training set.

Positional embeddings Since the input of the WebNLG task is a small KG which describes properties of entities, we introduce additional positional

Model		BLEU(%) \uparrow			METEOR \uparrow			TER \downarrow		
		Seen	Unseen	All	Seen	Unseen	All	Seen	Unseen	All
Without	Gardent et al. (2017)	54.52	33.27	45.13	0.41	0.33	0.37	0.40	0.55	0.47
Pretrained	Moryossef et al. (2019) ²	53.30	34.41	47.24	0.44	0.34	0.39	0.47	0.56	0.51
LM	Zhao et al. (2020)	64.42	38.23	52.78	0.45	0.37	0.41	0.33	0.53	0.42
With	Nan et al. (2021)	52.86	37.85	45.89	0.42	0.37	0.40	0.44	0.59	0.51
Pretrained	Kale (2020)	63.90	52.80	57.10	0.46	0.41	0.44	-	-	-
LM	Ribeiro et al. (2020b)	64.71	53.67	59.70	0.46	0.42	0.44	-	-	-
Our model	T5-large + Wiki + Position	66.07	53.87	60.56	0.46	0.42	0.44	0.32	0.41	0.36

Table 2: System Results on WebNLG Test Set Evaluated by BLEU, METEOR, and TER with Official Scripts

embeddings to enhance the flattened input of pre-trained Transformer-based sequence-to-sequence models such as BART and TaPas (Herzig et al., 2020). We extend the input layer with two position-aware embeddings in addition to the original position embeddings³ as shown in the Figure 2:

- Position ID, which is the same as the original position ID used in BART, is the index of the token in the flattened sequence $|S s_1 |P r_1 |O o_1 \dots |S s_n |P r_n |O o_n$.
- Triple Role ID takes 3 values for a specific triple (s_i, r_i, o_i) : 1 for the subject s_i , 2 for the relation r_i , and 3 for the object o_i .
- Tree level ID calculates the distance (the number of relations) from the root which is the source vertex of the RDF graph.

Two-step Fine-tuning To get better domain adaptation ability (Gururangan et al., 2020; Herzig et al., 2020), following TaPas and Wikipedia Person and Animal Dataset (Wang et al., 2018), we perform intermediate pre-training by coupling noisy English Wikipedia data with Wikidata triples, both of which are crawled in March 2020. We select 15 related categories (Astronaut, University, Monument, Building, ComicsCharacter, Food, Airport, SportsTeam, WrittenWork, Athlete, Artist, City, MeanOfTransportation, CelestialBody, Politician) that appear in the WebNLG dataset (Gardent et al., 2017) and collect 542,192 data pairs. For each Wikipedia article, we query its corresponding Wikidata triples and remove sentences which contain no values in the Wikidata triples to form graph-text pairs. Unlike (Chen et al., 2020a) which focuses on individual entity-sentence pairs for distant supervision, our pre-training corpus, on the other hand,

²For this baseline, we use the results reported from Zhao et al. (2020) who also use official evaluation scripts.

³For T5 models, we only keep the Triple Role and Tree-level embeddings.

is designed to better adapt to translating deeper graph structure into text. We remove triples and description pairs that have already appeared in the WebNLG dataset. After intermediate pre-training on this noisy corpus, we continue with fine-tuning our model on the WebNLG dataset.

3 Experiments

3.1 Dataset and Implementation details

Model	BLEU \uparrow	P \uparrow	R \uparrow	F1 \uparrow
BART-base	57.8	68.7	68.9	67.0
+ Wikipedia	59.7	69.6	70.7	68.4
+ Position	58.8	68.7	69.9	67.6
+ Wiki + Position	57.3	67.8	69.0	66.6
BART-large	58.3	67.9	69.4	66.8
+ Wikipedia	59.0	68.0	70.4	67.4
+ Position	58.1	67.6	69.4	66.6
+ Wiki + Position	60.0	68.6	69.2	67.1
distill-BART-xsum	59.1	69.9	70.6	68.5
+ Wikipedia	59.8	69.7	71.1	68.8
+ Position	59.2	69.8	70.2	68.3
+ Wiki + Position	59.9	70.1	70.1	68.7
T5-base	61.2	72.3	72.0	70.6
+ Wikipedia	60.9	72.0	71.8	70.2
+ Position	60.8	72.4	72.4	70.8
+ Wiki + Position	60.3	72.2	72.0	70.5
T5-large	60.0	71.6	72.1	70.2
+ Wikipedia	61.3	72.2	72.0	70.5
+ Position	60.6	72.1	72.4	70.6
+ Wiki + Position	61.9	72.8	73.5	71.6

Table 3: Results with both Wikipedia Fine-tuning and Positional Embedding for Various Pre-trained Models over All Categories on Development Set Evaluated by average of PARENT⁴ precision, recall, F1 and BLEU (%)

We use the original version of English WebNLG2017 (Gardent et al., 2017) dataset which contains 18,102/2,268/4,928 graph-description pairs for training, validation, and testing set respectively. For this task, we investigate a variety of the BART and T5 models with our novel tree-

⁴<https://github.com/KaijuML/parent>

level embeddings. The statistics and more details of those models are listed in Appendix A.

Model	P \uparrow	R \uparrow	F1 \uparrow
Gardent et al. (2017)	88.35	90.22	89.23
Moryossef et al. (2019)	85.77	89.34	87.46
Nan et al. (2021)	89.49	92.33	90.83
Ribeiro et al. (2020b)	89.36	91.96	90.59
T5-large + Wiki + Position	96.36	96.13	96.21

Table 4: System Results on WebNLG Test Set Evaluated by BERTScore precision, recall, F1 (%)

3.2 Results and Analysis

We use the standard NLG evaluation metrics to report results: BLEU (Papineni et al., 2002), METEOR (Lavie and Agarwal, 2007), and TER (Snover et al., 2006), as shown in Table 2. Because Castro Ferreira et al. (2020) has found that BERTScore (Zhang* et al., 2020) correlates with human evaluation ratings better, we use BERTScore to evaluate system results⁵ as shown in Table 4. When selecting the best models, we also evaluate each model with PARENT (Dhingra et al., 2019) metric which measures the overlap between predictions and both reference texts and graph contents. Dhingra et al. (2019) show PARENT metric has better human rating correlations. Table 3 shows the pre-trained models with 2-step fine-tuning and position embeddings achieve better results.⁶ We conduct paired t-test between our proposed model and all the other baselines on 10 randomly sampled subsets. The differences are statistically significant with $p \leq 0.008$ for all settings.

Results with Wikipedia fine-tuning. The Wikipedia fine-tuning helps the model handle unseen relations such as “*inOfficeWhileVicePresident*”, and “*activeYearsStartYear*” by stating “*His vice president is Atiku Abubakar.*” and “*started playing in 1995*” respectively. It also combines relations with the same type together with correct order, e.g., given two death places of a person, the model generates: “*died in Sidcup, London*” instead of generating two sentences or placing the city name ahead of the area name.

Results with positional embeddings. For the KG with multiple triples, additional positional embeddings help reduce the errors introduced by pro-

⁵We only use BERTScore to evaluate baselines which have results available online.

⁶For more examples, please check Appendix for reference.

noun ambiguity. For instance, for a KG which has “*leaderName*” relation to both country’s leader and university’s dean, position embeddings can distinguish these two relations by stating “*Denmark’s leader is Lars Løkke Rasmussen*” instead of “*its leader is Lars Løkke Rasmussen*”. The tree-level embeddings also help the model arrange multiple triples into one sentence, such as combining the city, the country, the affiliation, and the affiliation’s headquarter of a university into a single sentence: “*The School of Business and Social Sciences at the Aarhus University in Aarhus, Denmark is affiliated to the European University Association in Brussels*”.

3.3 Remaining Challenges

However, pre-trained language models also generate some errors as shown in Table 5. Because the language model is heavily pre-trained, it is biased against the occurrence of patterns that would enable it to infer the right relation. For example, for the “*activeYearsStartYear*” relation, the model might confuse it with the birth year. For some relations that do not have a clear direction, the language model is not powerful enough to consider the deep connections between the subject and the object. For example, for the relation “*doctoralStudent*”, the model mistakenly describes a professor as a Ph.D. student. Similarly, the model treats an asteroid as a person because it has an epoch date. For KGs with multiple triples, the generator still has a chance to miss relations or mixes the subject and the object of different relations, especially for the unseen category. For instance, for a soccer player with multiple clubs, the system might confuse the subject of one club’s relation with another club.

4 Related Work

The WebNLG task is similar to Wikibio generation (Lebret et al., 2016; Wang et al., 2018), AMR-to-text generation (Song et al., 2018) and RO-TOWIRE (Wiseman et al., 2017; Puduppully et al., 2019). Previous methods usually treat the graph-to-text generation as an end-to-end generation task. Those models (Trisedya et al., 2018; Gong et al., 2019; Shen et al., 2020) usually first linearize the knowledge graph and then use attention mechanism to generate the description sentences. While the linearization of input graph may sacrifice the inter-dependency inside input graph, some papers (Ribeiro et al., 2019, 2020a; Zhao et al., 2020)

Category	Output
T5-large	Andrew White (<i>born in 2003</i>) is a musician who is associated with the band Kaiser Chiefs and Marry Banilow. He is also associated with the label Polydor Records and is signed to B-Unique Records. S Aleksandra Kovač P activeYearsStartYear O 1990
T5-large	Walter Baade was born in the German Empire and graduated from the University of Gottingen. He was the doctoral student of Halton Arp and Allan Sandage and was the discoverer of 1036 Ganymed. S Walter Baade P doctoralStudent O Halton Arp; S Walter Baade P doctoralStudent O Allan Sandage
T5-large +Wiki	11264 Claudiomaccone was <i>born on the 26th of November, 2005</i> . He has an orbital period of 1513.722 days, a periapsis of 296521000.0 kilometres and an apoapsis of 475426000.0 kilometres. S 11264 Claudiomaccone P epoch O 2005-11-26; S Aleksandr Prudnikov P club O FC Amkar Perm
T5-large +Position	The chairman of FC Spartak Moscow is Sergey Rodionov. Aleksandr Prudnikov plays for FC Spartak Moscow and <i>manages FC Amkar Perm</i> . [S FC Amkar Perm P manager O Gadzhi Gadzhiyev; S Aleksandr Prudnikov P club O FC Amkar Perm]

Table 5: System Error Examples. We highlight *fabricated facts*, [missed relations], **incorrect relations**, and **ground truth relations** with different color.

use graph encoder such as GCN (Duvinaud et al., 2015) and graph transformer (Wang et al., 2020a; Koncel-Kedziorski et al., 2019) to encode the input graphs. Others (Shen et al., 2020; Wang et al., 2020b) try to carefully design loss functions to control the generation quality. With the development of computation resources, large scale PLMs such as GPT-2 (Radford et al., 2019), BART (Lewis et al., 2020) and T5 (Raffel et al., 2020) achieve state-of-the-art results even with simple linearized graph input (Harkous et al., 2020; Chen et al., 2020a; Kale, 2020; Ribeiro et al., 2020b). Instead of directly fine-tuning the PLMs, we propose a two-step fine-tuning mechanism to get better domain adaptation ability. In addition, using positional embeddings as an extension for PLMs has shown its effectiveness in table-based question answering (Herzig et al., 2020), fact verification (Chen et al., 2020b), and graph-to-text generation (Chen et al., 2020a). We capture the graph structure by enhancing the input layer with the triple role and tree-level embeddings.

5 Conclusions and Future Work

We propose a new two-step structured generation task for the graph-to-text generation task based on a two-step fine-tuning mechanism and novel tree-level position embeddings. In the future, we aim to address the remaining challenges and extend the framework for broader applications.

Acknowledgement

This work is partially supported by Agriculture and Food Research Initiative (AFRI) grant no. 2020-67021-32799/project accession no.1024178 from the USDA National Institute of Food and Agriculture, and by the Office of the Director of National

Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via contract # FA8650-17-C-9116. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- Thiago Castro Ferreira, Claire Gardent, Nikolai Ilinykh, Chris van der Lee, Simon Mille, Diego Moussallem, and Anastasia Shimorina. 2020. *The 2020 bilingual, bi-directional WebNLG+ shared task: Overview and evaluation results (WebNLG+ 2020)*. In *Proceedings of the 3rd International Workshop on Natural Language Generation from the Semantic Web (WebNLG+)*, pages 55–76, Dublin, Ireland (Virtual). Association for Computational Linguistics.
- Thiago Castro Ferreira, Chris van der Lee, Emiel van Miltenburg, and Emiel Krahmer. 2019. *Neural data-to-text generation: A comparison between pipeline and end-to-end architectures*. 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 552–562, Hong Kong, China. Association for Computational Linguistics.
- Wenhu Chen, Yu Su, Xifeng Yan, and William Yang Wang. 2020a. *KGPT: Knowledge-grounded pre-training for data-to-text generation*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8635–8648, Online. Association for Computational Linguistics.

- Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyong Zhou, and William Yang Wang. 2020b. [Tabfact: A large-scale dataset for table-based fact verification](#). In *Proceedings of the 8th International Conference on Learning Representations*.
- Bhuwan Dhingra, Manaal Faruqui, Ankur Parikh, Ming-Wei Chang, Dipanjan Das, and William Cohen. 2019. [Handling divergent reference texts when evaluating table-to-text generation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4884–4895, Florence, Italy. Association for Computational Linguistics.
- David K Duvenaud, Dougal Maclaurin, Jorge Iparaguirre, Rafael Bombarell, Timothy Hirzel, Alan Aspuru-Guzik, and Ryan P Adams. 2015. [Convolutional networks on graphs for learning molecular fingerprints](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2224–2232. Curran Associates, Inc.
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017. [The WebNLG challenge: Generating text from RDF data](#). In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 124–133, Santiago de Compostela, Spain. Association for Computational Linguistics.
- 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)*, pages 3143–3152, Hong Kong, China. Association for Computational Linguistics.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Hamza Harkous, Isabel Groves, and Amir Saffari. 2020. [Have your text and use it too! end-to-end neural data-to-text generation with semantic fidelity](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2410–2424, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Jonathan Herzig, Pawel Krzysztof Nowak, Thomas Müller, Francesco Piccinno, and Julian Eisenschlos. 2020. [TaPas: Weakly supervised table parsing via pre-training](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4320–4333, Online. Association for Computational Linguistics.
- Mihir Kale. 2020. [Text-to-text pre-training for data-to-text tasks](#). *Computation and Language Repository*, arXiv:2005.10433. Version 2.
- Rik Koncel-Kedziorski, 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.
- Alon Lavie and Abhaya Agarwal. 2007. [METEOR: An automatic metric for MT evaluation with high levels of correlation with human judgments](#). In *Proceedings of the Second Workshop on Statistical Machine Translation*, pages 228–231, Prague, Czech Republic. Association for Computational Linguistics.
- 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*, pages 1203–1213, Austin, Texas. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- 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*, pages 1–9, Tilburg University, The Netherlands. 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.
- Linyong Nan, Dragomir Radev, Rui Zhang, Amrit Rau, Abhinand Sivaprasad, Chiachun Hsieh, Xiangru Tang, Aadit Vyas, Neha Verma, Pranav Krishna, Yangxiaokang Liu, Nadia Irwanto, Jessica

- Pan, Faiaz Rahman, Ahmad Zaidi, Mutethia Mutuma, Yasin Tarabar, Ankit Gupta, Tao Yu, Yi Chern Tan, Xi Victoria Lin, Caiming Xiong, Richard Socher, and Nazneen Fatema Rajani. 2021. [DART: Open-domain structured data record to text generation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 432–447, Online. Association for Computational Linguistics.
- 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, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Ratish Puduppully, Li Dong, and Mirella Lapata. 2019. [Data-to-text generation with content selection and planning](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6908–6915.
- 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. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Leonardo F. R. Ribeiro, Claire Gardent, and Iryna Gurevych. 2019. [Enhancing AMR-to-text generation with dual graph representations](#). 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 3183–3194, Hong Kong, China. Association for Computational Linguistics.
- Leonardo F. R. Ribeiro, Yue Zhang, Claire Gardent, and Iryna Gurevych. 2020a. [Modeling global and local node contexts for text generation from knowledge graphs](#). *Transactions of the Association for Computational Linguistics*, 8:589–604.
- Leonardo FR Ribeiro, Martin Schmitt, Hinrich Schütze, and Iryna Gurevych. 2020b. [Investigating pre-trained language models for graph-to-text generation](#). *arXiv preprint arXiv:2007.08426*.
- Xiaoyu Shen, Ernie Chang, Hui Su, Cheng Niu, and Dietrich Klakow. 2020. [Neural data-to-text generation via jointly learning the segmentation and correspondence](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7155–7165, Online. Association for Computational Linguistics.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and Ralph Weischedel. 2006. [A study of translation error rate with targeted human annotation](#). In *In Proceedings of the Association for Machine Translation in the Americas (AMTA 2006)*.
- 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.
- Bayu Distiawan Trisedya, Jianzhong Qi, Rui Zhang, and Wei Wang. 2018. [GTR-LSTM: A triple encoder for sentence generation from RDF data](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1627–1637, Melbourne, Australia. Association for Computational Linguistics.
- Qingyun Wang, Xiaoman Pan, Lifu Huang, Boliang Zhang, Zhiying Jiang, Heng Ji, and Kevin Knight. 2018. [Describing a knowledge base](#). In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 10–21, Tilburg University, The Netherlands. Association for Computational Linguistics.
- Tianming Wang, Xiaojun Wan, and Hanqi Jin. 2020a. [AMR-to-text generation with graph transformer](#). *Transactions of the Association for Computational Linguistics*, 8:19–33.
- Zhenyi Wang, Xiaoyang Wang, Bang An, Dong Yu, and Changyou Chen. 2020b. [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.
- 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.
- Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Chao Zhao, Marilyn Walker, and Snigdha Chaturvedi. 2020. [Bridging the structural gap between encoding and decoding for data-to-text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2481–2491, Online. Association for Computational Linguistics.