



HAL
open science

Knowledge Graph for NLG in the context of conversational agents

Hussam Ghanem, Massinissa Atmani, Christophe Cruz

► **To cite this version:**

Hussam Ghanem, Massinissa Atmani, Christophe Cruz. Knowledge Graph for NLG in the context of conversational agents. French Regional Conference on Complex Systems (FRCCS 2023), May 2023, Le Havre, France. hal-04146881

HAL Id: hal-04146881

<https://hal.science/hal-04146881>

Submitted on 3 Jul 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Knowledge Graph for NLG in the context of conversational agents

Hussam Ghanem¹ · Massinissa Atmani¹ ·
Christophe Cruz¹

Received: date / Accepted: date

Abstract The use of knowledge graphs (KGs) enhances the accuracy and comprehensiveness of the responses provided by a conversational agent. While generating answers during conversations consists in generating text from these KGs, it is still regarded as a challenging task that has gained significant attention in recent years. In this document, we provide a review of different architectures used for knowledge graph-to-text generation including: Graph Neural Networks, the Graph Transformer, and linearization with seq2seq models. We discuss the advantages and limitations of each architecture and conclude that the choice of architecture will depend on the specific requirements of the task at hand. We also highlight the importance of considering constraints such as execution time and model validity, particularly in the context of conversational agents. Based on these constraints and the availability of labeled data for the domains of DAVI, we choose to use seq2seq Transformer-based models (PLMs) for the Knowledge Graph-to-Text Generation task. We aim to refine benchmark datasets of kg-to-text generation on PLMs and to explore the emotional and multilingual dimensions in our future work. Overall, this review provides insights into the different approaches for knowledge graph-to-text generation and outlines future directions for research in this area.

Keywords Conversational agents · Knowledge graphs · Natural Language Generation

1 Introduction

Conversational agents, also known as chatbots, are computer programs designed to simulate conversation with human users [55]. These agents can be

ICB, UMR 6306, CNRS, Université de Bourgogne
21000 Dijon, France

integrated with messaging platforms, mobile applications, and websites to provide instant support to customers and handle simple tasks, such as answering questions or helping with bookings. Conversational agents using knowledge graphs (KG) [9] are a type of chatbot that leverages structured data stored in a knowledge graph to generate human-like responses. The knowledge graph is a graph-based representation of entities and their relationships, providing a structured source of information for the chatbot to access and use during conversation. This enables the chatbot to provide more accurate and comprehensive answers to user’s questions. The use of knowledge graphs can greatly enhance the capabilities of conversational agents and make interactions more informative and useful [57].

Generating answers during conversations consists in generating text from data. Data-to-text processes require algorithms that generate linguistically correct sentences for humans and express the semantics and structure of non-linguistic data (sequence, tree, graph, etc.). In addition, the generation of textual answers requires supporting several languages. And for a better interaction with a conversational agent, the emotional context of the conversation (Common Ground) is fundamental. The aim of this work in collaboration with the company DAVI is to integrate a socio-emotional dimension into human-machine interactions which complement the technical and “business” skills linked to professional expertise. The company DAVI is a software publisher in SaaS mode which has expertise in the fields of AI, Affective Computing, and Human Machine Interactions (HMIs). The following picture presents the composite AI of DAVI’s solution including Natural Language Understanding, Emotion detection, skills modeling, Natural Language Generation, and Body Language Generation.




Fig. 1 Composite AI pipeline of virtual agents at DAVI

For now, the natural language generation (NLG) of a conversational engine does not benefit from the latest technological advances in natural language processing (NLP). The Natural Language processing step is based on a manual process to define the template of the answer. This process requires a costly amount of time. Thus, the purpose of this project is to automate and reduce the burden of the generation of template-based responses (as the responses are manually written through a set of rules) in the implementation of conversa-

tional agents. The template-based responses are modelized and stored in the skills' database. Regarding the emotional dimension, emotional responses are injected automatically in the answer depending the emotional analysis of the user. To automate NLG answers from Skills, knowledge graphs were identified.

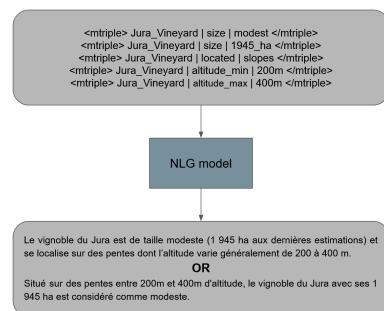


Fig. 2 Automatic text generation from the knowledge graph

In NLG, Two criteria [33] are used to assess the quality of the produced answers. The first criterion is semantic consistency (Semantic Fidelity) which quantifies the fidelity of the data produced against the input data. The most common indicators are 1/ **Hallucination**: It is manifested by the presence of information (facts) in the generated text that is not present in the input data; 2/ **Omission**: It is manifested by the omission of one of the pieces of information (facts) in the generated text; 3/ **Redundancy**: This is manifested by the repetition of information in the generated text; 4/ **Accuracy**: The lack of accuracy is manifested by the modification of information such as the inversion of the subject and the direct object complement in the generated text; 5/ **Ordering**: It occurs when the sequence of information is different from the input data. The second criterion is linguistic coherence (Output Fluency) to evaluate the fluidity of the text and the linguistic constructions of the generated text, the segmentation of the text into different sentences, the use of anaphoric pronouns to reference entities and to have linguistically correct sentences.

Today, neural approaches offer performances exceeding all classical methods for linguistic coherence. However, limits are still present to maintain semantic consistency, and their performance deteriorates even more on long texts [30]. Another limitation due to the complexity of neural approaches is that text generation is non-parameterized with no control over the structure of the generated text. Thus, most of the current neural approaches arrive behind template-based approaches on the criterion of semantic consistency [31], but they are far superior to them on the criterion of linguistic consistency. This

can be explained by the fact that large language models manage to capture certain syntactic and semantic properties of the language.

The present review is organized as follows, Section 2 presents a comprehensive overview of the current state-of-the-art approaches for knowledge graph-to-text generation. In Section 3, we present the latest architectures and techniques that have been proposed in this field. Finally, Section 4 critically examines the strengths and limitations of these techniques in the context of conversational agents


2 Knowledge Graph-to-Text Generation

KG-to-text generation aims at producing easy-to-understand sentences in natural language from knowledge graphs (KGs) while maintaining semantic consistency between the generated sentences and the KG triplets. Compared to the traditional text generation task (Seq2Seq), generating text from a knowledge graph is an additional challenge to guarantee the authenticity of the words in the generated sentences. The existing methods can be classified according to three categories (Figure 3) and will be detailed later:

- **Linearisation with Sequence-to-sequence (Seq2Seq)**: convert the graph to a sequence which is fed to a sequence-to-sequence model;
- **Graph Neural Networks (GNNs)**: encode topological structures of a graph and learn the representation of an entity by the aggregation of the features of the entities and neighbors. They are not used as a standalone model and require a decoder to complete the encoder-decoder architecture;
- **Graph Transformer (GT)**: the enhanced version of the original transformer adapted to handle graphs.

The term "knowledge graph" has been around since 1972, but its current definition can be traced back to Google's 2012 announcement of their Knowledge Graph. This was followed by similar announcements from companies such as Airbnb, Amazon, eBay, Facebook, IBM, LinkedIn, Microsoft, and Uber, among others, leading to an increase in the adoption of knowledge graphs by various industries [4]. As a result, academic research in this field has seen a surge in recent years, with an increasing number of scientific publications on knowledge graphs [4]. These graphs utilize a graph-based data model to effectively manage, integrate, and extract valuable insights from large and diverse datasets [5].

Knowledge graphs, which are composed of nodes that represent different types of entities and edges that denote various types of relationships between those entities, are known as heterogeneous graphs. The integration of information from multiple sources and domains in knowledge graphs leads to an even greater degree of heterogeneity. To address this, recent research has applied heterogeneous graph embedding methods to represent knowledge graphs



Le vignoble du Jura est de taille modeste (1 945 ha aux dernières estimations) et se localise sur des pentes dont l'altitude varie généralement de 200 à 400 m.
Or
 Situé sur des pentes entre 200m et 400m d'altitude, le vignoble du Jura avec ses 1 945 ha est considéré comme modeste.

Fig. 3 The architecture of KG-to-text generation with the three categories of representation, a) Linearization + Seq2Seq, b) GNNs with decoder (e.g. LSTM), and c) Graph Transformer (GT)

effectively. For example, ERNIE [40] and KnowBERT [41] employ knowledge graph embedding techniques such as TransE [42] and Tucker [58] to encode knowledge graphs.

Generating text and learning alignments between source entities/relationships and target tokens from scratch is a challenging task for standard language models because of the limited amount of parallel graph-text data [17,34]. To overcome this limitation, recent research has focused on developing generalist pre-trained language models for KG-to-text generation. A common approach is to linearize input graphs into text sequences and fine-tune pre-trained seq2seq Transformer models such as GPT [35], BART [10], or T5 [36] based on KG-to-text datasets [1,37]. These pre-trained language models can generate high-quality texts with a simple fine-tuning to the target task, thanks to their self-supervised pre-training on large-scale corpora of unlabeled texts. In fact, pre-trained language models outperform other models with sophisticated structures in KG-to-text generation tasks. This type of approach will be detailed in section 3.

According to [14], text generation tasks using KG-to-text models mainly fall under three aspects:

- **Encoder modification:** To reduce the loss of structural information in sequence encoders with graph inputs that have been linearized [24,6,32], proposals concentrate on constructing more intricate encoder structures to improve the representation of graphs, including GNNs and GTs;
- **Unsupervised training:** These proposals consist in designing unsupervised training methods to jointly learn graph-to-text and text-to-graph con-

- version tasks with non-parallel graph-to-text data [39,17,18]. This makes it possible to compare the final result of the process with the input data;
- **Build pre-trained models:** With the development of pre-trained Natural Language Generation (NLG) models such as GPT, BART, and T5, some recent work directly refines these models on graph-to-text datasets and reports significant performance [1,37,15,2].

Compared to existing work on pre-trained models for KG-text generation, the JoinGT model [14] uses pre-training methods to explicitly learn graph-text alignments instead of directly tuning the pre-trained models seq2seq on KG-to-text datasets.

3 Architectures

In this section, the different architectures used in data-to-text tasks will be presented. As the nature of the data greatly influences the choice of the architecture of the neural approaches, most works either try to adapt the inputs to the architectures of the models or propose new architectures better adapted to the types of input data. Due to the nature of sentences and the tree structure of its representations, several works have proposed to model data structures of this type to enhance performance.

3.1 Graph linearisation and sequence to sequence models (Seq2seq)

Recent years have been marked by significant achievements in the field of PLMs (Pretrained Language Models) [38,44]. Pre-trained on massive corpora, PLMs exhibit good generalization ability to solve related NLG (Natural Language Generation) downstream tasks [43]. However, most existing PLMs were trained on textual data [44,10] without ingesting any structured data input. The seq-to-seq category consists of linearizing the KG [1,25,24,3] and then formulating a Seq2seq generation task using PLMs like GPT [35], BART [10] or T5 [36] with linearized KG nodes as input to generate sentences. The use of pre-trained language models (PLMs) in KG-to-text generation has shown superior performance, but still faces two major challenges: 1) loss of structural information during encoding, as existing models like BERT do not explicitly take into account the relationship between input entities; and 2) lack of explicit graph-text alignments, as complex knowledge graph structures make it difficult to learn graph-text alignments through text reconstruction-based pre-training tasks. Despite attempts to retain as much of the graph topology as possible with seq2seq methods, the Transformer-based seq2seq models' cost is not cheap (especially in the pretraining phase). Also, the computational cost of linearization can be high for large knowledge graphs. Hence, and so to better keep the graph topology, Graph Neural Networks (GNNs) have been proposed, which will be discussed in the next section.

3.2 Graph Neural Networks (GNNs)

Different approaches use different variants of GNNs architectures such as GCNs (Graph Convolutional Networks) [27] or extensions of GCNs such as Syn-GCNs [45] or DCGCNs [11]. Other approaches use the variant GATs (Graph Attention Networks) [8]. Or approaches that use the GGNNs (Gated Graph Neural Networks) variant [16,21,47]. Graph Neural Networks (GNNs) are a type of neural network that are well-suited for processing graph-structured data. In the context of knowledge graph-to-text generation, GNNs can be used to model the relationships between entities in a knowledge graph and generate text based on those relationships. Recent research on using GNNs for knowledge graph to text generation has shown promising results. Some studies have used graph convolutional networks (GCNs) [27] to encode the relationships between entities in a knowledge graph into a low-dimensional representation, or extensions of GCNs such as Syn-GCNs [45] and DCGCNs [11]. Other studies have used graph attention networks (GATs) [8] to dynamically weight the importance of different entities and relationships in the knowledge graph when generating text. Other studies have used a gating mechanism (Gated Graph Neural Networks or GGNNs) that allows for effectively controlling the flow of information between nodes in the graph, which is useful for incorporating contextual information [16,21,47]. Additionally, some researchers have combined GNNs with reinforcement learning to generate text that maximizes a reward function defined over the generated text and the knowledge graph [21].

Overall, the use of GNNs for knowledge graph to text generation is an active area of research, with many recent studies exploring different architectures and training methods. The results so far suggest that GNNs can effectively capture the relationships between entities in a knowledge graph and generate high-quality text based on that information. A limitation related to KG-to-Text generation with GNNs, is that GNNs can be computationally expensive and may struggle to handle large knowledge graphs. Additionally, their performance may degrade for graphs with complex relationships or structures. Despite these limitations, GNNs remain a promising direction for knowledge graph-to-text generation.

3.3 Graphs Transformers (GTs)

In order to benefit from the power of models based on Transformer and to be able to model tree or graph-type data structures as with GNNs, and to overcome the limitations of local neighborhood aggregation while avoiding strict structural inductive biases, recent works have proposed to adapt the Transformer architecture. As Graph Transformers are equipped with self-attention mechanisms, they can capture global context information by applying these mechanisms to the nodes of the graph.

According to [12], GT differs from GNNs in that it allows direct modeling of dependencies between any pair of nodes regardless of their distance in the input graph. An undesirable consequence is that it essentially treats any graph as a fully connected graph, greatly reducing the explicit structure of the graph. To maintain a structure-aware view of the graph, their proposed model introduces an explicit relationship encoding and integrates it into the pairwise attention score computation as a dynamic parameter.

From the GNNs pipeline, if we make several parallel heads of neighborhood aggregation and replace the sum on the neighbors by the attention mechanism, e.g. a weighted sum, we would get the Graph Attention Network (GAT). Adding normalization and MLP feed-forward, we have a Graph Transformer [46]. For the same reasons as Graph Transformer, [13] presents the K-BERT model, they introduce four components to augment the Transformer architecture and to be able to handle a graph as input. The first knowledge layer component takes a sentence and a set of triplets as input and outputs the sentence tree by expanding the sentence entities with their corresponding triplets. They also add the Seeing Layer component to preserve the original sentence structure and model the relationships of the triples by building a Visibility Matrix. Another component is the Mask-Transformer where they modify the self-attention layer to consider the visibility matrix when calculating attention.

The use of Graph Transformers for Knowledge graph to text generation has gained popularity in recent years due to their ability to effectively handle graph structures and capture the relationships between nodes in the graph. Additionally, Graph Transformers can handle large graphs and have the ability to model long-range dependencies between nodes in the graph. Despite the advantages, the training of Graph Transformers can be computationally expensive and the interpretability of the model remains a challenge. Overall, the use of Graph Transformers for Knowledge graph to text generation is a promising area of research and can lead to significant improvements in the generation of text from knowledge graphs.

4 Discussion

The Graph Neural Network (GNN), the Graph Transformer, and linearization with seq2seq models are three different architectures for knowledge graph to text generation, a task that involves generating natural language text from structured knowledge representations like knowledge graphs. GNNs are a type of deep learning model that are well-suited for processing graph-structured data. They provide a flexible and scalable way to model the graph structure and relationships, but they may not be able to efficiently handle large and complex graphs. The Graph Transformer is a specialized version of the Transformer, designed for graph-to-sequence learning tasks. It provides a more direct and efficient way to process the graph structure compared to linearization

with seq2seq models, but it may require more training data and computation resources. Linearization with seq2seq models is a simpler and easier to implement approach that involves converting the knowledge graph into a linear sequence, such as a sentence, and then using a seq2seq model to generate text from the linearized representation. However, this approach can lose some of the structural information and relationships in the knowledge graph during the linearization process.

In summary, each of the three architectures has its own advantages and disadvantages, and the choice of architecture will depend on the specific requirements of the actual task.

As our project is part of a context of conversational agents, we must take into account all the consequent constraints such as the execution time (to respect the instant conversation constraint) and the validity of the model where the answer must not be incorrect or ambiguous.

If most of the current models have satisfactory validity performances, the inference time of models based on GNN and GraphTransformer exceeds the limit threshold found in a fluid and natural conversation and requires a huge memory load that violates the standards of current industrialization (MLOps) of neural models.

In light of these elements, and with the constraint of the data labellisation for the domains of DAVI, we choose to go further with seq2seq Transformer based models (PLMs) in our Knowledge Graph-to-Text Generation. We also want to shed light on the fact that DAVI already handles such models in their pipeline and they have the knowledge and infrastructure to optimize the integration and deployment of the seq2seq models. Hence, DAVI should still remain in control of the time to market of the NLG solution.

5 Conclusion

In conclusion, the document discusses different data-to-text architectures and highlights their advantages and limitations in the context of graph-to-text project with DAVI. The Graph Neural Network (GNN), Graph Transformer, and seq2seq models are three approaches that have been applied to the task of generating natural language text from structured knowledge representations like knowledge graphs. Each approach has its own advantages and disadvantages, and the choice of architecture will depend on the specific requirements of the task. Considering the constraints of our project, which includes developing a conversational agent that must generate valid responses in real-time, we have decided to move forward with seq2seq Transformer-based models (PLMs) as they have satisfactory performance on validity and execution time. Additionally, DAVI already handles such models in their pipeline and can optimize their integration and deployment. Our next step will be to explore state-of-the-art

approaches that take into account the emotional and multilingual dimensions to achieve the objectives of the graph-to-text project.

References

1. Ribeiro, L. F., Schmitt, M., Schütze, H., & Gurevych, I. (2020). Investigating pretrained language models for graph-to-text generation. arXiv preprint arXiv:2007.08426.
2. Mager, M., Astudillo, R. F., Naseem, T., Sultan, M. A., Lee, Y. S., Florian, R., & Roukos, S. (2020). Gpt-too: A language-model-first approach for amr-to-text generation. arXiv preprint arXiv:2005.09123.
3. Hoyle, A., Marasović, A., & Smith, N. (2020). Promoting graph awareness in linearized graph-to-text generation. arXiv preprint arXiv:2012.15793.
4. Hogan A, Blomqvist E, Cochez M, d’Amato C, Melo GD, Gutierrez C, Kirrane S, Gayo JE, Navigli R, Neumaier S, Ngomo AC. Knowledge graphs. ACM Computing Surveys (CSUR). 2021 Jul 2;54(4):1-37.
5. N. F. Noy, Y. Gao, A. Jain, A. Narayanan, A. Patterson, and J. Taylor. 2019. Industry-scale knowledge graphs: Lessons and challenges. ACM Queue 17, 2 (2019).
6. Distiawan, B., Qi, J., Zhang, R., & Wang, W. (2018, July). 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) (pp. 1627-1637).
7. Ribeiro, L. F., Zhang, Y., & Gurevych, I. (2021). Structural adapters in pretrained language models for amr-to-text generation. arXiv preprint arXiv:2103.09120.
8. Ribeiro, L. F., Zhang, Y., Gardent, C., & Gurevych, I. (2020). Modeling global and local node contexts for text generation from knowledge graphs. Transactions of the Association for Computational Linguistics, 8, 589-604.
9. Moon, S., Shah, P., Kumar, A., & Subba, R. (2019, July). Opendialkg: Explainable conversational reasoning with attention-based walks over knowledge graphs. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (pp. 845-854).
10. Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., ... & Zettlemoyer, L. (2019). Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. arXiv preprint arXiv:1910.13461. (not graph)
11. Guo, Z., Zhang, Y., Teng, Z., & Lu, W. (2019). Densely connected graph convolutional networks for graph-to-sequence learning. Transactions of the Association for Computational Linguistics, 7, 297-312.
12. Cai, D., & Lam, W. (2020, April). Graph transformer for graph-to-sequence learning. In Proceedings of the AAAI conference on artificial intelligence (Vol. 34, No. 05, pp. 7464-7471).
13. Liu, W., Zhou, P., Zhao, Z., Wang, Z., Ju, Q., Deng, H., & Wang, P. (2020, April). K-bert: Enabling language representation with knowledge graph. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 34, No. 03, pp. 2901-2908).
14. Ke, P., Ji, H., Ran, Y., Cui, X., Wang, L., Song, L., ... & Huang, M. (2021). Jointgt: Graph-text joint representation learning for text generation from knowledge graphs. arXiv preprint arXiv:2106.10502.
15. Chen, W., Su, Y., Yan, X., & Wang, W. Y. (2020). Kgpt: Knowledge-grounded pre-training for data-to-text generation. arXiv preprint arXiv:2010.02307.
16. Chen, Y., Wu, L., & Zaki, M. J. (2020). Toward subgraph guided knowledge graph question generation with graph neural networks. arXiv preprint arXiv:2004.06015.
17. Guo, Q., Jin, Z., Qiu, X., Zhang, W., Wipf, D., & Zhang, Z. (2020). Cyclegt: Unsupervised graph-to-text and text-to-graph generation via cycle training. arXiv preprint arXiv:2006.04702
18. Jin, Z., Guo, Q., Qiu, X., & Zhang, Z. (2020, December). Genwiki: A dataset of 1.3 million content-sharing text and graphs for unsupervised graph-to-text generation. In Proceedings of the 28th International Conference on Computational Linguistics (pp. 2398-2409).
19. Wang, Q., Yavuz, S., Lin, V., Ji, H., & Rajani, N. (2021). Stage-wise Fine-tuning for Graph-to-Text Generation. arXiv preprint arXiv:2105.08021.

20. Song, L., Wang, A., Su, J., Zhang, Y., Xu, K., Ge, Y., & Yu, D. (2021). Structural information preserving for graph-to-text generation. arXiv preprint arXiv:2102.06749.
21. Chen, Y., Wu, L., & Zaki, M. J. (2019). Reinforcement learning based graph-to-sequence model for natural question generation. arXiv preprint arXiv:1908.04942.
22. Harkous, H., Groves, I., & Saffari, A. (2020). Have your text and use it too! end-to-end neural data-to-text generation with semantic fidelity. arXiv preprint arXiv:2004.06577.
23. Koncel-Kedziorski, R., Bekal, D., Luan, Y., Lapata, M., & Hajishirzi, H. (2019). Text generation from knowledge graphs with graph transformers. arXiv preprint arXiv:1904.02342.
24. Gardent, C., Shimorina, A., Narayan, S., & Perez-Beltrachini, L. (2017, September). The WebNLG challenge: Generating text from RDF data. In Proceedings of the 10th International Conference on Natural Language Generation (pp. 124-133).
25. Yang, Z., Einolghozati, A., Inan, H., Diedrick, K., Fan, A., Donmez, P., & Gupta, S. (2020). Improving text-to-text pre-trained models for the graph-to-text task. In Proceedings of the 3rd International Workshop on Natural Language Generation from the Semantic Web (WebNLG+) (pp. 107-116).
26. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
27. Kipf, T. N., & Welling, M. (2016). Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907.
28. M. Gori, G. Monfardini, and F. Scarselli, "A new model for learning in graph domains," in Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005., vol. 2, pp. 729-734, IEEE, 2005.
29. F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, "The graph neural network model," *IEEE Transactions on Neural Networks*, vol. 20, no. 1, pp. 61-80, 2008.
30. Ratish Puduppully, Li Dong, and Mirella Lapata. 2018. Data-to-text generation with content selection and planning.
31. Yevgeniy Puzikov and Iryna Gurevych. 2018. E2e nlg challenge: Neural models vs. templates.
32. Moryossef, A., Goldberg, Y., & Dagan, I. (2019). Step-by-step: Separating planning from realization in neural data-to-text generation. arXiv preprint arXiv:1904.03396.
33. Ferreira, T. C., van der Lee, C., Van Miltenburg, E., & Krahrmer, E. (2019). Neural data-to-text generation: A comparison between pipeline and end-to-end architectures. arXiv preprint arXiv:1908.09022.
34. Fu, Z., Shi, B., Lam, W., Bing, L., & Liu, Z. (2020). Partially-aligned data-to-text generation with distant supervision. arXiv preprint arXiv:2010.01268.
35. Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training.
36. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., ... & Liu, P. J. (2020). Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140), 1-67.
37. Kale, M., & Rastogi, A. (2020). Text-to-text pre-training for data-to-text tasks. arXiv preprint arXiv:2005.10433.
38. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.
39. Schmitt, M., Sharifzadeh, S., Tresp, V., & Schütze, H. (2019). An unsupervised joint system for text generation from knowledge graphs and semantic parsing. arXiv preprint arXiv:1904.09447.
40. Zhang, Z., Han, X., Liu, Z., Jiang, X., Sun, M., & Liu, Q. (2019). ERNIE: Enhanced language representation with informative entities. arXiv preprint arXiv:1905.07129.
41. Peters, M. E., Neumann, M., Logan IV, R. L., Schwartz, R., Joshi, V., Singh, S., & Smith, N. A. (2019). Knowledge enhanced contextual word representations. arXiv preprint arXiv:1909.04164.
42. Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., & Yakhnenko, O. (2013). Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems*, 26.

43. Li, J., Tang, T., Zhao, W. X., & Wen, J. R. (2021). Pretrained language models for text generation: A survey. arXiv preprint arXiv:2105.10311.
44. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are unsupervised multitask learners. OpenAI blog, 1(8), 9.
45. Marcheggiani, D., Frolov, A., & Titov, I. (2017). A simple and accurate syntax-agnostic neural model for dependency-based semantic role labeling. arXiv preprint arXiv:1701.02593.
46. Chaitanya K. Joshi. Transformers are Graph Neural Networks. The Gradient.
47. Li, Y., Tarlow, D., Brockschmidt, M., & Zemel, R. (2015). Gated graph sequence neural networks. arXiv preprint arXiv:1511.05493.
48. Dwivedi, V. P., & Bresson, X. (2020). A generalization of transformer networks to graphs. arXiv preprint arXiv:2012.09699.
49. Dušek, O., & Kasner, Z. (2020). Evaluating semantic accuracy of data-to-text generation with natural language inference. arXiv preprint arXiv:2011.10819.
50. Tian, R., Narayan, S., Sellam, T., & Parikh, A. P. (2019). Sticking to the facts: Confident decoding for faithful data-to-text generation. arXiv preprint arXiv:1910.08684.
51. Xiao, Y., & Wang, W. Y. (2021). On hallucination and predictive uncertainty in conditional language generation. arXiv preprint arXiv:2103.15025.
52. 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. <https://doi.org/10.18653/v1/D17-1239>
53. Durmus, E., He, H., & Diab, M. (2020). FEQA: A question answering evaluation framework for faithfulness assessment in abstractive summarization. arXiv preprint arXiv:2005.03754.
54. Katja Filippova. 2020. Controlled Hallucinations: Learning to Generate Faithfully from Noisy Data. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings. 8646870.
55. Wahde, M., & Virgolin, M. (2022). Conversational agents: Theory and applications. In HANDBOOK ON COMPUTER LEARNING AND INTELLIGENCE: Volume 2: Deep Learning, Intelligent Control and Evolutionary Computation (pp. 497-544).
56. Hussain, S., Ameri Sianaki, O., & Ababneh, N. (2019). A survey on conversational agents/chatbots classification and design techniques. In Web, Artificial Intelligence and Network Applications: Proceedings of the Workshops of the 33rd International Conference on Advanced Information Networking and Applications (WAINA-2019) 33 (pp. 946-956). Springer International Publishing.
57. Ait-Mlouk, A., & Jiang, L. (2020). KBot: A Knowledge graph based chatBot for natural language understanding over linked data. IEEE Access, 8, 149220-149230.
58. Balažević, I., Allen, C., & Hospedales, T. M. (2019). Tucker: Tensor factorization for knowledge graph completion. arXiv preprint arXiv:1901.09590.