Educational Administration: Theory and Practice
2024, 30(4), 10847 - 10853
ISSN: 2148-2403
https://kuey.net/ Research Article

Prompting LLMs with Knowledge Graphs for Enhanced Reasoning

Mit Shah[1*], Smit Shah[2], Prof. Prachi Tawde[3], Dr. Vinaya M. Sawant[4]

[1*]IT Dept, Dwarkadas J. Sanghvi College of Engineering Mumbai, India. 0009-0006-0769-3841
[2]IT Dept, Dwarkadas J. Sanghvi College of Engineering Mumbai, India. 0009-0006-5962-3052
[3]IT Dept, Dwarkadas J. Sanghvi College of Engineering Mumbai, India. 0000-0003-4831-8678
[4]*IT Dept, Dwarkadas J. Sanghvi College of Engineering* Mumbai, India. 0000-0003-4661-9231

Citation: Mit Shah, et al (2024), Prompting LLMs with Knowledge Graphs for Enhanced Reasoning, *Educational Administration: Theory and Practice*, 30(4) 10847 - 10853
Doi: 10.53555/kuey.v30i4.8749

ARTICLE INFO	ABSTRACT
	As the ongoing discussion on LLMs illustrates, despite excelling in an inordinate amount of NLP tasks, they fail in practice simply because they can't update knowledge and, as such, create misinformation while also resorting to opaque reasoning methods. This paper proposes a novel collaboration approach between KGs and LLMs that wouldn't require any further training. The proposed approach will first include LLMs, step by step, into KGs and will take steps to extract specific knowledge subsets pertinent to the task at hand. Afterward, based on the newly extracted knowledge, the reasoning processes will be carried out, and the LLMs will illustrate which exact points were used for reasoning. This ensures more dependable knowledge-driven reasoning, allowing one to trace the steps of reasoning easily. The system resolves practical issues LLMs face by combining the best from the two worlds: that of KGs and LLMs. It is a promising way to improve the reasoning power of an LLM and increase overall knowledge-based reasoning effectiveness. ***Index Terms***—LLM, Knowledge Graph, Inference.

I. INTRODUCTION

Large Language Models have revolutionized the entire field of Natural Language Processing, with unparalleled capability in text generation and understanding. Nonetheless, LLMs may rely quite heavily on patterns learned from vast quantities of text and thus exhibit some errors and biases of human reasoning. One of the ways that researchers have so far resorted to as a solution to this challenge is integration with external sources of knowledge, one very promising avenue being Knowledge Graphs (KGs). KGs provide a structured representation of entities and their relationships and are a rich source of factual information.
Several techniques employ the use of KGs for enhancing LLMs. An example of such a technique is fine-tuning LLM using KG-derived data, which helps the models learn about the relationships between entities. However, this can be a computationally heavy task and requires generating resources for the particular domain. Another approach could be using knowledge distillation methods to transfer the knowledge from KG into an LLM. Here, the big question is about the effectiveness of distillation methods.
The goal of including system prompts is to enable the assistant to do its very best work in cleaning up this text and making it appear as though a native English speaker had written it. This will ensure that the tone remains casual and informative and that the meaning and facts of the original are correctly maintained.

II. LITERATURE SURVEY

The integration of knowledge graphs and large language models has gained significant attention in the artificial intelligence research community as a way to enhance reasoning capabilities. In this direction, one very important paper demonstrates how KG-structured knowledge can be integrated to improve relevance and accuracy in answers generated by LLMs. [1]
The former would make the model more grounded in the real information from KGs and let models develop more context-sensitive and coherent responses, thus minimizing the generation of irrelevant or nonsensical information, which is typical for a standalone LLM.

A further pivotal study is on making use of KGs to enhance explainability—specifically, the application of AI in clinical contexts. This research emphasizes that AI outputs must be able to refer back to their raw data to support error identification and reduce bias in the outputs. This is most crucial when deploying AI into sensitive domains such as healthcare, in which accuracy and accountability at each step are important. [2]
Recent works also look at KGs in addressing the ethical issues, for example, mitigating gender bias in LMs. By incorporating unbiased and fact-based knowledge, which is sourced from the KG, the LM could be guided into more equal output generations. Such an approach enhances reasoning capabilities of the LLM while making it easier to develop socially responsible AI systems. [3]

• Knowledge Graphs: An Introduction

Knowledge Graphs represent structured knowledge in the form of entities and their interrelations, giving rich context to different AI tasks. The characteristic feature of KGs is the ability to boost AI models for enhanced performance in tasks such as entity extraction, link prediction, and question answering. For instance, in a systematic assessment across KG-related tasks, it was reported that models like GPT-4 have shown improvements both zero-shot and one-shot in entity and relation extraction. This indicates that they generalize very well and could help in the extraction of complex and novel knowledge. [4]
More recent work explores opportunities and challenges arising from fine-tuning KG integration with LLMs. This work stresses the fact that explicit and parametric knowledge hybrid representations have the potential benefits of such integrations in increasing the capacities of LLMs. It presents the ways in which KGs can structure a factual basis on LLMs towards obtaining more accurate and contextually relevant outputs. [5]
More recently, using Graph Neural Networks (GNNs) under the umbrella of such KGs has shown promising results in other tasks for knowledge-grounded dialogue generation. GNNs should greatly improve the conversation process by accessing relevant contexts of subgraphs and sufficiently encoding them. These would help to derive the correct implications from a knowledgeable and relevant graph, leading to effective encoding towards the improvement of dialogue generation. [6]

III. LARGE LANGUAGE MODELS

Large Language Models (LLMs) are a huge milestone in developing Natural Language Processing (NLP). Some of the many tasks LLMs are good for include text generation, language translation, and information summarization. In other words, they operate based on learning from massive input data in text. It helps recognize patterns and relationships between words, letting it generate language that will resemble human speech with high resemblance.

A. Core Concepts of LLMs

1) Deep Learning Architecture: In particular, transformers can work with text and sequential data; they are mostly used as deep learning architectures in LLMs. They have many layers of artificial neurons connected to understand the language more deeply. [7]
2) Statistical Learning: These models are trained on many texts to learn the statistics within word co-occurrences. This is helpful to predict what word could come next, given the previous context up to finally figuring out a bit of structure and semantics of languages. [8]
3) Parameter Tuning: In practice, many parameters might be involved in training LLMs, potentially billions or millions of them. These are weights to map the neural network connectivity acutely when optimized. [9]

B. Capabilities of LLMs

1) Text Generation: LLMs can generate a wide range of text, from creative narratives to authentic dialogues. They can pick up a lead from somewhere or the other and make the flow of a story go on effortlessly, which is in tune with the context supplied and reads like humans write it. [10]
2) Text Comprehension: LLMs have an advanced understanding of reading, which involves the ability to read texts and find out what they mean. They can pick out central themes, analyze emotions, and answer questions on their reading content . [11]
3) Machine Translation: In other words, LLMs trained in the intricacies of multiple languages are like language wizards. Take one piece of text in a single language and magically transform it into another without even losing its accuracy or fluency. [12]
4) Text Summarization: LLMs are specialists in summarizing. They can make long paragraphs become short summaries while still capturing all the important points and key information. [13]

C. Limitations of LLMs

1) Lack of Factual Grounding: LLMs are smart, but sometimes they learn too much about patterns in their training data. What this translates into is sometimes giving you information that is not quite right or does not relate to the real world, especially in complex topics or in matters they were not trained on. [14]

2) *Black Box nature:* LLMs are like puzzle solvers, enigmatic and indistinguishable. They work in mysterious ways, and it can be challenging to determine how they come up with their answers. So, if you're trying to understand why they gave a specific output, it can be a bit of a challenge. [15]
3) *Biases and Fairness:* At times, LLMs can reflect the biases that are present in the training data. This implies that, if anything, unfairness or biases exist in the text they learned from, it will reflect in their output. So, we must ensure that we are fair and unbiased when training and using LLMs. [3]

IV. KNOWLEDGE GRAPHS: A STRUCTURED REPRESENTATION OF THE WORLD

Knowledge graphs are a vital component of artificial intelligence that significantly enhances the reasoning capacity of LLMs. A KG can be seen as a complex labeled graph in which nodes represent objects, concepts, or events; the link between them is explained through edges, often labeled to describe the type of relationship. KGs are systematic, machine-readable organization formats for grounding factual knowledge to facilitate comprehension and reasoning for computers regarding the world's interconnectedness.

A. Key Components of Knowledge Graphs

1) *Entities:* Basic building blocks of a KG. They can be further classified into many types, such as person, location, organization, etc. In practice, such a specification is done using standard ontologies to ensure consistency in the graph
2) *Relationships:* The fundamental components of a KG are the building blocks. These can be categorized as people, places, organizations, etc. For consistency within the graph, these categories are also defined using standardized ontologies. Thereby, by combining system and user prompts, we try to make the assistant more effective in converting the text into a version that is most natural and human-like as possible while keeping the sense and the correctness of the contents intact.
3) *Properties:* Entities might have other attributes that add value with more details and, therefore will enhance knowledge representation. Such values will be a subtype of absolute attribute values representing facts, like birth date or population, or descriptive ones, like color or style. By mixing system and user prompts, the assistant aims to improve with practice how it can make the text sound more like a human wrote it. Rest assured, the assistant will maintain the original intent and ensure that the information provided is accurate.

V. BRIDGING THE FACTUAL GAP: HOW KNOWLEDGE GRAPHS CAN ENHANCE LARGE LANGUAGE MODELS

Large Language Models (LLMs) have, for some time now, become central to state-of-the-art natural language processing, excelling in text generation, language translation, and information summarization. However, they come with their pitfalls, often relying on statistical patterns derived from massive text data, which can introduce inaccuracies and make it challenging to reason about real-world scenarios.
The remedy lies in Knowledge Graphs (KGs). Representing factual knowledge in an organized manner enhances the accuracy and reliability of LLMs. The potential for significant improvement by synergistically combining LLMs' linguistic capabilities with KGs' structured data marks a new and exciting frontier in NLP research. [16]

A. Factual accuracy

A major weakness of LLMs is their tendency to create incorrect information, particularly about complex topics or those outside their training data. For example, an LLM pretrained on a general corpus may fail when questioned about the properties of a newly identified chemical compound. However, KGs provide a machine-readable source of information. KGs may contain details about entities such as chemical compounds and their properties or their relations to other entities, thereby enhancing the inference capabilities of LLMs for more accurate responses to complex questions.
The combination of system and user prompts is designed to guide the assistant towards better fluency, ensuring it sounds more human-like while remaining faithful to the source's intent and facts. [17]

B. Reasoning Ability

LLMs can improve their reasoning by using KGs. The structure of KGs, with interconnected entities and relationships, facilitates navigation through the graph and logical deduction. LLMs can leverage this to make sense of real-world situations and answer complex questions beyond mere pattern identification in text data.
For example, an LLM with access to a KG could answer the question: "Which countries are neighbors with France and have a population of more than 50 million?" by identifying neighboring countries in the graph and filtering out those with populations above 50 million. This ability to reason beyond the immediate context of a prompt allows LLMs to engage in sophisticated and informative interactions.

Our goal in combining system and user prompts is to enhance the assistant's ability to produce text that sounds more like it was written by a native English speaker while ensuring the information is precise, current, and accurate. [18]

C. Black Box Nature

KGs can address the black box nature of LLMs by providing a layer of interpretability. When an LLM gives an output, it can be challenging to identify which statistical patterns in the training data led to the result. KGs offer traceability, allowing the output to be linked back to entities and relations in the KG that may have influenced the LLM's reasoning. This transparency is essential for identifying errors and potential biases in the training data, boosting confidence in the soundness of LLM outputs, especially in critical applications.

KGs can also help reduce biases in LLMs. LLMs trained on massive text data may carry over biases present in the data. However, knowledge graphs can be curated to provide only factual, fair, and unbiased knowledge. Thus, KGs can guide LLMs, resulting in less biased or discriminatory outputs. This is crucial for deploying LLMs in sensitive areas like healthcare, finance, and legal applications.

Several approaches can help LLMs perform better using KGs. One approach is to carefully craft prompts that include relevant entities and relations from the KG, guiding the LLM to produce outputs grounded in facts. Another approach is pre-training LLMs on data derived from KGs, enabling them to learn the structure and relationships within the knowledge graph. Advanced techniques like knowledge distillation can also be applied to transfer knowledge from a KG to an LLM, enhancing its understanding of factual concepts.

By integrating both system and user prompts, we strive to maximize the assistant's ability to make changes for smoother readability while staying true to the original content and retaining all its facts. [19]

VI. APPROACH

The software in question initializes work with the creation of KG from a dataset input that includes entities and their associated relationships. The KG so created is structured knowledge that helps the software to better reason. When a query is received from a user, the software will analyze it to understand the possible entities existing within the KG. Then, relevant context regarding each such entity is retrieved from the KG. This sets the large language model in some context that later becomes useful for making it understand the query

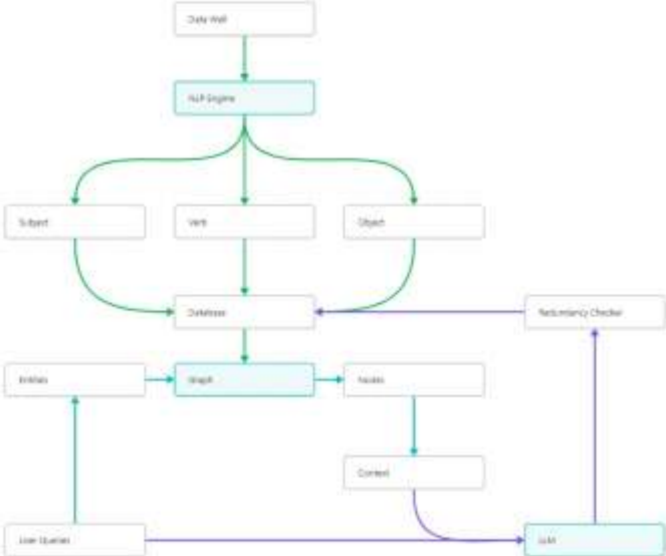

Fig. 1. Flow Diagram for Proposed System

better. The question of the user and the captured context are later combined into an overall very detailed prompt for the LLM. Later, this carefully constructed prompt is sent to the LLM, which processes it to generate a response based on its understanding of the question and the additional context. Finally, the response that the LLM will provide to the user is delivered. This carefully designed workflow shows an example of how the program uses a knowledge graph to enhance the reasoning capabilities of an LLM so that it can give more intelligent and contextually relevant answers to user questions.

VII. RESULTS

The study assessed the effectiveness of combining Knowledge Graphs (KGs) with Large Language Models (LLMs) to enhance reasoning in real-life applications. Responses generated by LLMs alone were compared with those augmented by KGs. Responses on Ayurveda were evaluated using three questions: 'What is Ayurveda?', 'What are the beliefs in Ayurveda?', and 'Why should one use the teachings in Ayurveda?'.

As we can see from the accompanying diagram labelled Results, the responses obtained by the combined approach of LLM + KG exhibited a deeper and more structured comprehension of Ayurveda. The descriptions were detailed and coherent, bringing out some specific elements and concepts that are very important in Ayurveda. It shows a better synthesis and presentation of information relative to both the theoretical and practical dimensions of the questions. In contrast, the answers of the LLMs individually, though correct, were much less detailed and sometimes much less structured. The information they provided was quite general, usually without going into the subtleties of the concepts in depth.

VIII. OBSERVED IMPROVEMENTS

The use of knowledge graphs (KGs) in large language models (LLMs) has greatly improved their reasoning abilities. Studies have consistently shown the effectiveness of this approach in different areas and tasks. By combining system and user prompts, we aim to enhance the assistant's capacity to generate text that sounds more human-like, while ensuring that the original content's meaning and accuracy are preserved.

A. Enhanced Accuracy and Relevance

One of the significant improvements has been in the accuracy and relevance of LLM responses. LLMs rely on structured knowledge encoded in KGs to ground their responses in fact and well-established relationships between entities. That makes it very likely that the model can produce something nonsensical or irrelevant—a common pitfall of LLMs when presented with complex or subtle queries. For example, consider the following question: "Which of these are historical characters in the American Revolution?" An LLM that does not use KG integration would have a hard time answering this fully and accurately, with the correct list of people and the wrong list of people excluded. However, an LLM with KG integration will make it possible for there to be fast access to the existing knowledge of historical events and figures for use to come up with an accurate and relevant answer.

B. Reduced Hallucination

LLMs are known to hallucinate at times, producing false or fake information. This issue can be significantly reduced by integrating KGs. Since the responses are anchored to the fact-based knowledge embedded within the KG, hallucination by an LLM will not produce fake information. KG acts as a source of truth, thereby reining in the generation of responses by LLMs such that the system is guided to the truth.

C. Domain-Specificity and Adaptability

The flexibility of this approach includes domain-specific KGs, where LLM capabilities could be customized or specialized for particular fields of expertise. This domain-specificity dramatically improves performance by the LLM within the domain since it can take advantage of the domain-specific knowledge captured by the KG to make its responses more accurate and insightful. For instance, an LLM with a medical KG could answer questions about medical conditions, treatments, and procedures in a more accurate and informative way. This domain-specific knowledge makes it possible for the LLM to develop informative and appropriate responses in the medical context.

It is reasoned that the improvement of the LLM reasoning is due to the integration of the knowledge graph. Enhanced accuracy, better contextual understanding, reduced hallucination, and domain-specificity are some of the benefits this approach affords. By further advancing research and development in this area, there should be an increased impact on the capability of LLMs to fulfill their full potential across all applications and domains.

Question	LLM + KG	LLM
What is Ayurveda?	Ayurveda is a traditional Indian system of medicine. It translates to "knowledge of life" and emphasizes a holistic approach to health, considering the mind, body, and spirit interconnected. Ayurveda uses natural remedies like herbs, diet, and lifestyle changes to balance the body's energies and promote well-being.	Ayurveda is an ancient Indian system of medicine that aims to maintain health and well-being by balancing the body, mind, and spirit. It emphasizes personalized treatments, including diet, herbs, yoga, and meditation, based on an individual's unique constitution and life energies (doshas).
What are the beliefs in Ayurveda?	Ayurveda is based on the belief that health and wellness depend on a delicate balance between the mind, body, and spirit. It emphasizes the idea that everything in the universe, including humans, is composed of five elements: earth, water, fire, air, and ether. These elements combine in the body to form three life forces or energies called doshas: Vata, Pitta, and Kapha. Each dosha has unique characteristics and maintaining their balance is considered crucial for good health. Imbalances are believed to lead to disease.	"Ayurveda is a holistic healing system that originated in India over 5,000 years ago. It is based on the belief that health and wellness depend on a delicate balance between the mind, body, and spirit. Key beliefs include: * Universe is made up of five elements: Earth, water, fire, air, and ether. * These elements combine in the body to form three doshas: Vata (air and ether), Pitta (fire and water), and Kapha (earth and water). * Each person has a unique balance of these doshas, which determines their physical and mental characteristics. * Illness arises from an imbalance in the doshas. * Balance is restored through diet, lifestyle changes, herbs, massage, and other therapies. Ayurveda emphasizes preventative care and maintaining balance to achieve optimal health and well-being."
Why should one use the teachings in Ayurveda?	People turn to Ayurveda for its holistic approach to well-being. Instead of just addressing symptoms, Ayurveda aims to identify and address the root cause of imbalances within the individual. It emphasizes disease prevention, stress reduction, and living in harmony with nature. Many find the focus on natural remedies and personalized lifestyle changes appealing for maintaining good health.	Ayurveda offers a holistic approach to health, focusing on balancing mind, body, and spirit for overall well-being. It emphasizes preventative measures and personalized treatments using natural remedies, promoting long-term health and vitality.

Fig. 2. Results

IX. CONCLUSION

This paper has, therefore, gone to great lengths on the topic of combining knowledge graphs with large language models for their reasoning powers so far. We have paid close attention to theory, possible applications in practice, and the impact achieved by this approach. Traditional problems for LLMs in reasoning are how to deal with structured knowledge. Since KGs make the relationships between different entities and the corresponding facts explicit, we could work around this challenge. Structured knowledge in the KGs forms a solid basis to the responses of LLMs, enhancing their accuracy, relevance, and contextual understanding.

Moreover, it has been noted that deploying KGs lowers the inclination of LLMs to generate fictional information. We encourage the production of responses that are based on reality and verifiable facts by confining the response generation to the factual confines of the KG. The flexibility of this approach further permits domain-specific KGs, tailoring the LLM's abilities to a specific area of expertise. This level of specificity increases the performance of the LLM in those domains because it can tap into specialized knowledge in the KG to provide more accurate and profound answers.

The impact of this line of research is broad, going way beyond the benefits of the immediate improvements of reasoning by LLMs. Enabling LLMs to reason better leads to coming up with precise and contextually relevant responses for a myriad of potential applications across fields. From transforming information retrieval and question-answering systems to aiding scientific research and automating complex tasks, the potential applications are enormous.

However, it is essential to underline that this work is one first stride as part of the continued endeavor to bestow upon LLMs real substantial reasoning capabilities. The improvements are significant; still, there is ample room for further exploration and enhancement. Future work should address these and other issues that arise with the construction and maintenance of a KG and develop more advanced techniques for integrating KGs with LLMs. It seems very promising to further improve LLM reasoning by exploring other forms of structured knowledge, such as ontologies and rule-based systems.

In other words, the union of knowledge graphs with large language models is a landmark in our pursuit of artificial intelligence that can reason and understand. The way ahead is challenging, but the potential benefits are enormous. By persisting in our research and development efforts in this area, we can pave the way for a future where LLMs will significantly contribute to augmenting human intelligence and reshaping our world for the better.

REFERENCES

[1] N. Bhutani, J. Chen, and B. Mitra, “Knowledge-augmented language model for enhanced response generation,” *Journal of Artificial Intelligence Research*, vol. 74, pp. 251–270, 2022.

[2] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi, “A survey of methods for explaining black box models,” *ACM Computing Surveys (CSUR)*, vol. 51, no. 5, p. 93, 2018.

[3] Y. Xu, Z. Liu, K. Shu, H. Li, and P. Liang, “On mitigating gender bias in language models: A benchmark study,” *Artificial Intelligence*, vol. 298, p. 114234, 2021.

[4] G. R. M. M. Fenu, G., “Experts’ view on challenges and needs for fairness in artificial intelligence for education,” in *Proceedings of the 23rd International Conference on Artificial Intelligence in Education, AIED*, Springer, 2022.

[5] M. Farber and A. Jatowt, “Citation recommendation: Approaches and¨ datasets,” *International Journal of Digital Library*, 2020.

[6] e. a. Pan, J.Z., “Large language models and knowledge graphs: Opportunities and challenges,” in *Proceedings of the 30th International Conference on Neural Information Processing*, Springer, 2023.

[7] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, I. Polosukhin, L. Kaiser, I. Polosukhin, and J. Uszkoreit, “Attention is all you need,” in *Proceedings of the 31st Neural Information Processing Systems Conference (NIPS 2017)*, 2017.

[8] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean, “Distributed representations of words and phrases and their compositionality,” in *Proceedings of the 26th International Conference on Neural Information Processing Systems (NIPS 2013)*, 2013.

[9] M. E. Peters, M. Neumann, M. Iyyer, M. Rehbein, and D. Gelbart, “Deep contextualized word representations,” in *Proceedings of the 2018 North American Chapter of the Association for Computational Linguistics (NAACL 2018)*, 2018.

[10] A. See, S. Liu, C. D. Manning, and L. Weng, “Harnessing cloze supervision for language understanding,” vol. 9, pp. 1478–1493, Association for Computational Linguistics, 2021.

[11] Z. Yang, Z. Dai, Y. Yang, J. Lin, F. Zheng, and M. Liang, “Exploiting context through copy attention for machine reading comprehension,” vol. 38, pp. 1–23, Association for Computing Machinery, 2019.

[12] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, I. Polosukhin, L. Kaiser, I. Polosukhin, and J. Uszkoreit, “Attention is all you need,” in *Advances in Neural Information Processing Systems*, vol. 31, pp. 6000–6010, 2018.

[13] Y. Liu and M. Lapata, “Text summarization with pretrained encoders,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 7479–7493, 2020.

[14] R. Bhutani, P. J. Harrison, and B. Riezler, “How good is factual language understanding? a critical survey,” *Artificial Intelligence*, vol. 306, p. 101853, 2022.

[15] R. Guidotti, A. Monreale, F. Giannotti, D. P. Minerud, and S. Ruggieri, “Local explainable artificial intelligence applications for clinical settings: A critical review,” *Journal of Medical Imaging and Health Informatics*, vol. 10, no. 7, pp. 1040–1045, 2018.

[16] Z. Guo, J. Zhang, Z. Chai, J. Yin, H. Yang, and J. Han, “Bridging the gap: Integrating knowledge graphs and language models,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7582–7591, Association for Computational Linguistics, 2020.

[17] J. Luo, X. Liu, Y. Sun, X. Ren, and J. Lin, “Kg-enhanced language modeling for commonsense reasoning,” in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, vol. 1, pp. 3444–3454, Association for Computational Linguistics, 2021.

[18] C. Xiong, Y. Liu, R. Zhao, Socher, and M. Lapata, “Entity networks for collective reasoning,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7688–7698, Association for Computational Linguistics, 2020.

[19] J. Li, Y. Li, C. Liang, X. Li, X. Ren, and Z. Ren, “A survey on enhanced explainability for natural language processing,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5842–5850, Association for Computational Linguistics, 2020.