iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** XI **Month of publication:** November 2026

DOI: https://doi.org/10.22214/ijraset.2026.85693

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue XI Nov 2026- Available at www.ijraset.com

Impact of Knowledge Graph Expansion and Pruning on the Performance of Question Answering Graph Neural Networks

Gumelar Purnama Nugraha[1], Novaren Veraldo[2]
Artificial Intelligence, Nanjing University of Information Science and Technology, Nanjing, China

Abstract: ***Question Answering systems, which integrate Large Language Models (LLMs) with Knowledge Graphs (KGs), aim to reduce hallucination and unsubstantiated answers by grounding responses in structured factual data. However, a limitation persists as knowledge graphs often do not contain all the necessary information to address every question This study aim to propose an approach that balances the advantages of expanded reasoning with the potential pitfalls of noise. Integrating a KG-completion model to strategically expand the knowledge graph, and then employ an iterative node pruning mechanism. This allows for a precise observation of the trade-offs between the advantages of expanded reasoning and the potential pitfalls of noise, ultimately aiming to optimize the KG for improved accuracy and robustness in question answering, moving beyond the limitations of relying solely on smaller, pre-defined subgraphs. This expansion notably influenced the QA-GNN model's reasoning, leading to improvements in several instances and suggesting that further parameter tuning could yield even more optimal results.***

Keywords: Question Answering, Knowledge Graph, Machine Learning

I. INTRODUCTION

Large Language Models (LLMs) have recently demonstrated significant progress in the field of natural language processing, leading to remarkable advancements in various applications [1], [2]. Prior research has demonstrated a remarkable profieciency in various natural language understanding and generation tasks without specific fine tuning [3], [4].

However, LLMs often operate as black boxes, making it difficult to understand the reasoning behind their generated answers [5]. Studies [6], [7] also revealed that these LLMs can sometimes lead to the generation of hallucinations or unfaithful knowledge in real world applications due to deficiency in relevant knowledge. This lack of transparency poses a significant challenge, particularly in applications like Question Answering (QA) systems. In such systems, given a natural language question and a Knowledge Graph (KG) as input, the model is tasked with selecting the correct answer. Popular benchmarks used to assess the performance of these systems, such as CommonsenseQA [8] and OpenBookQA [9], typically involve evaluating a model's ability to answer questions that require common sense knowledge and reasoning.

KG are a structured data format representing information as interconnected nodes (entities) and edges (relationships), forming a vast repository of factual knowledge. The structured nature of KG enables them to effectively support complex reasoning processes, The incorporation of these structured facts may offer a solution to the problem of hallucinations observed in LLMs, [10], [11].

Crucially, by utilizing the structured information within a KG, one can gain valuable insights into how the language model infers the answer to a given question. This transparency allows for a deeper understanding of the model's decision-making process, enhancing accountability and enabling us to identify and address potential biases or errors. Furthermore, the use of KGs can improve the interpretability of LLM outputs, making it easier to understand why a particular answer was generated, which is crucial for building trust and ensuring the responsible deployment of these models.

Reasoning is a key aspect of human intelligence. While LLMs excel in language tasks, their reasoning abilities are limited. KGs offer structured knowledge, enabling more robust reasoning. By integrating LLMs with KGs, and leverage their strengths. LLMs can enhance KGs by extracting and integrating new information from text, while KGs can improve LLM reasoning by providing access to structured knowledge and enabling knowledge-based reasoning tasks such as question answering and recommendation systems. While knowledge bases like Freebase [12] and ConceptNet [13] boast an impressive volume of information, they inevitably have limitations. Overcoming this limitation necessitates the development of sophisticated models that can effectively reason with incomplete information, draw inferences, and ultimately arrive at accurate answers.

©IJRASET: All Rights are Reserved

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue XI Nov 2026- Available at www.ijraset.com

II. RELATED WORKS

A. *Prompt engineering*

One of the approach to enhancing KG utilization involves employing prompt engineering techniques. This entails carefully designing prompts that present the LLM with a specific fact retrieved from the KG and then instructing the LLM to perform reasoning tasks based on this provided information.

KG-GPT [14] employs sentence segmentation to decompose complex questions into simpler sub-sentences, enabling it to reason over individual, single-relation sub-sentences more effectively than attempting to process the entire question at once. While this approach offers a valuable strategy, it can lead to shallow reasoning, as the LLM primarily relies on direct information retrieval from the KG and basic linguistic understanding, often failing to capture deeper, more nuanced relationships StructGPT [15] unifies LLM reasoning by integrating diverse structured data like databases, tables, and KGs. This enables more comprehensive reasoning, but can be limited by the quality and completeness of these data sources and the complexity of integrating them seamlessly.

Think-on-Graph [16] facilitates reasoning by representing information as a graph, enabling more sophisticated inference and problem-solving, but can be computationally expensive and may require careful graph construction.

DSLR [17] selective ranking prioritizes relevant information by selectively ranking documents, leading to more focused and informative outputs, but might miss crucial context if the initial selection is inaccurate.

B. *Joint Reasoning*

Recent research also explored the potential of enhancing LLMs with KGs for improved question answering. A key area of focus has been developing joint reasoning models that effectively integrate information from both sources. QA-GNN [18], while demonstrating ease of implementation and commendable performance on negation questions, suffers from a simplistic node representation, limiting its potential.

An enhancement of QA-GNN, JointLK [19], incorporates a bidirectional attention mechanism and iterative pruning to mitigate the impact of noisy subgraphs, leading to improved performance. However, JointLK's efficacy is contingent upon the quality of the underlying KG.

The core idea behind GreaseLM [20] is to allow for multi-layered, bidirectional interaction between the language context and structured knowledge from KGs. Despite achieving competitive results with significantly larger models, it presents a complex architecture.

Similarly, UniKGQA [21], while leveraging abstraction to merge subgraph relations, also exhibits considerable complexity. These findings suggest that while advancements in KG question answering have yielded promising results, there remains a need for models that balance performance with interpretability and computational efficiency.

C. *KG Completion*

Knowledge graph completion, often referred to as link prediction, constitutes a flourishing area of research and plays a pivotal role in the successful operation of many QA systems. By meticulously inferring potential nodes or relationships within the existing KG, the system can broaden the range of questions it can effectively address, increasing its overall utility. KG-BERT [22] leverages the power of BERT to improve KG completion by treating triples as textual sequences. This approach allows it to capture complex relationships and achieve state-of-the-art results. COMET leverages pre-trained language models to generate novel commonsense knowledge, achieving high accuracy in KG completion tasks. This approach excels at capturing complex relationships and generating human-like knowledge. However, it can be computationally expensive and may struggle with factual accuracy and handling contradictory information.

III. METHODOLOGY

Existing research in KG reasoning often prioritizes denoising and subsequent reasoning within a focused subgraph [18], [19]. This approach typically relies on the availability of a high-quality, comprehensive KG. However, it faces limitations when dealing with entities outside the scope of the existing KG, hindering effective reasoning for these entities. This research proposes an approach that aims to overcome these limitations on a question answer system. By dynamically introducing predicted nodes and relations into the working graph using a KG completion model, this research aim to:

1) Expand the reasoning space: The inclusion of predicted nodes can facilitate reasoning beyond the constraints of the initial KG, potentially enabling the model to discover new connections and insights.
2) Bridge reasoning gaps: In cases of incomplete KGs, predicted nodes can act as intermediaries, establishing connections between entities in the question and the answer entities, thereby improving reasoning accuracy.

©IJRASET: All Rights are Reserved

Building upon the QA-GNN approach as shown in Fig. 1, this research refines the initial working graph by incorporating predicted nodes generated from a KG completion model. The model then iteratively updates node representations through message passing and prunes nodes with low confidence at each step. This allows for broader reasoning beyond the limitations of the original KG while mitigating the risk of erroneous predictions through controlled pruning [19], [23].

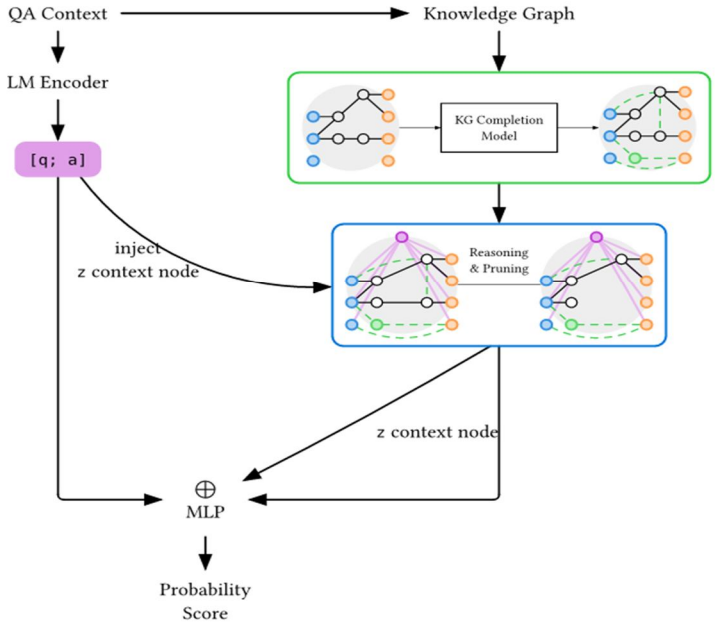

Fig. 1: Overall architecture

A. Working Graph Creation

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be the KG and $\boldsymbol{\mathcal{V}}$ nodes and set of edges $\boldsymbol{\mathcal{E}} \subseteq \boldsymbol{\mathcal{V}} \times \boldsymbol{\mathcal{R}} \times \boldsymbol{\mathcal{V}}$ with $\boldsymbol{\mathcal{R}}$ relation types. Given the question and answer choices, extract the entity $e \in \boldsymbol{\mathcal{E}}$ from the text questions and answer. Form the subgraph $\boldsymbol{\mathcal{G}_{WG}}$ consist of entities appearing in the question and answer, $\{e_q\}$ and $\{e_a\}$ respectively. Indirect relations between both sets of entities are also included, through common neighbor of both set of entities $\{e_o\}$.

This working graph, however, may contain incomplete connections between entities. To address this, a KG completion model is employed to predict missing relations $\{e_x\}$. This prediction process encompasses both one-hop relations predicting $(e_1, ?, e_2)$ and N-hop relations predicting $(e_1, ?, e_?, \dots, e_2)$, effectively identifying potential intermediate entities (Fig. 2). Each predicted relation is associated with a confidence value, reflecting the model's certainty in the prediction.

The $e \in \{e_x\}$ confidence values are obtained through the completion model. The final working graph are consist of $\boldsymbol{\mathcal{G}_{WG}} = \{e_q\} \cup \{e_a\} \cup \{e_o\} \cup \{e_x\}$. Each node is associated with one of the five types $\boldsymbol{\mathcal{T}} = \{\text{Z, Q, A, O, X}\}$.

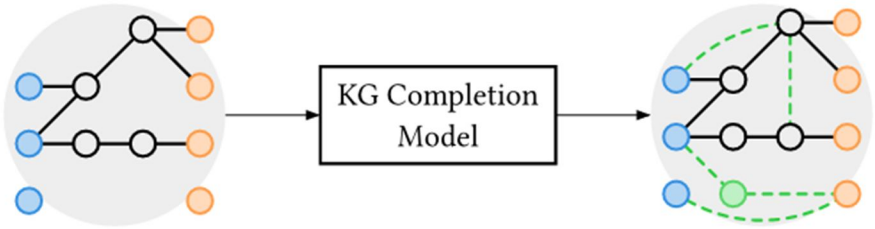

Fig. 2: Working graphs missing relations and nodes are filled with predicted entities *B. Joint Reasoning & Selective Pruning*

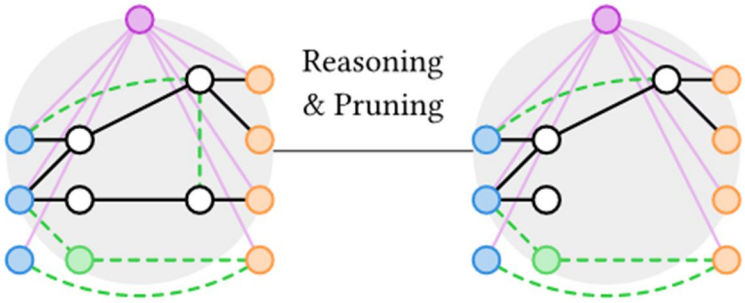

Fig. 3: Working graph Reasoning and Pruning process

Building upon the foundation of QA-GNNs [18], to facilitate for joint reasoning across two distinct knowledge sources, node z is introduced as a context node within the graph structure. This context node encapsulates the concatenated representation of the question and the corresponding answer option $[q; a]$.

The primary function of this context node is to adaptively capture the intricate and dynamic relationships that exist between the question and the candidate answer within the broader context of the KG. To achieve this, the study leverage the Graph Attention Network (GAT) framework [24] to implement an iterative message-passing mechanism. This iterative process enables the propagation and refinement of information across neighboring nodes within the graph, leading to progressively more informative node representations. Furthermore, to enhance the model's ability to assess the reliability of information flow, incorporating confidence scores as node features.

$$h_t^{(l+1)} = f_n\left(\sum_{s \in N_t \cup \{t\}} (\alpha_{st} m_{st})\right) + h_t^{(l)},$$

$$m_{st} = f_m\left(h_s^{(l)}, u_s, r_{st}, c_{st}\right),$$

$$u_t = f_u(u_t), r_{st} = f_r(e_{st}, u_s, u_t,),$$

$$c_{st} = f_c(c_s, cr_{st})$$

where $u_s, u_t \in \{0,1\}^{|\mathcal{T}|}$ one hot vector indicating the node types of s and $t, e_{st} \in \{0,1\}^{|\mathcal{R}|}$ is a one-hot vector indicating the relation type of the edge (s, t) and $c_{st}, cr_{st} \in [0,1]$ is a confidence value of the nodes s and edge of (s, t) respectively.

Finally, to optimize computational efficiency and focus the reasoning process on the most salient information, the study employ a pruning strategy. After each iteration of the message-passing process, irrelevant nodes as determined by their contribution to the overall reasoning process, are strategically removed from the graph [19].

For each node v, concatenate the entity text(v) with the QA context text(z) and compute the relevance score,

$$p_v = f_{head} = \left(LM(text(z); text(v);)\right)$$

The relevance score p_v quantifies the importance of each KG node within the given QA context. Nodes with scores below a defined threshold are ranked and subsequently pruned. Given the question q and an answer choice a, calculate the probability of it being the answer

$$p(a|q) \propto exp\left(\text{MLP}(z^{\text{LM}}, z^{\text{GNN}}, g)\right)$$

where $z^{\text{GNN}} = h_z^{(L)}$ and g denotes the pooling of $\left\{h_v^{(L)} | v \in V_{\text{sub}}\right\}$.

IV. EXPERIMENTS

The study evaluate the method using CommonsenseQA [25], a multiple-choice question answering dataset that specifically tests a model's ability to understand and utilize commonsense knowledge. Unlike many QA datasets that focus on factual recall or reading comprehension within a given text, CommonsenseQA requires models to go beyond the provided context and draw upon general human understanding of the world to answer questions.

The CommonsenseQA dataset was constructed using ConceptNet, a large-scale semantic network designed to provide computational systems with a wide range of commonsense knowledge. ConceptNet functions as a KG, where nodes represent concepts and natural language phrases, and edges illustrate the diverse semantic relationships between them. This structure offers a detailed representation of everyday understanding, which is essential for tackling the complexities of commonsense reasoning.
Since the CommonsenseQA test set isn't publicly accessible, our primary experiments were conducted using the in-house (IH) data splits from KagNet [26]
ConceptNet was integrated as a key component to enhance question-answering capabilities. For each QA context, utilizing ConceptNet as the underlying KG $\mathcal{G}$ to extract relevant subgraphs $\mathcal{G}_{\text{sub}}$. This involved retrieving information from the ConceptNet graph, extending outwards with a hop size of $k = 2$ from the entiti s found in the QA context. Following this initial retrieval, a pruning step was applied. This step refined the extracted subgraph by evaluating the relevance score of each node, ensuring that only the most pertinent information was retained for further analysis.
The study leverage the robust capabilities of COMET's [27] automatic knowledge graph construction model to predict novel relationships within the $\mathcal{G}_{\text{sub}}$. COMET, being a powerful model for commonsense KG completion, allows us to infer implicit knowledge from pre-trained language models, thus generating new and potentially valuable connections. A key characteristic of COMET's output is its very high quality, indicated by its low perplexity score, signifying the model's strong predictive performance and coherence.
A crucial validation step involves cross-referencing the predicted knowledge with ConceptNet, and checked for redundancy. Any predicted relations that are found to already exist within ConceptNet are subsequently removed to ensure the novelty and non-redundancy of the generated knowledge. To further enhance the quality and relevance of the predictions, the study implement a penalization mechanism, in which the predicted knowledge scores are adjusted based on their rank. This prioritizes higher-ranked, more novel, and potentially more impactful insights, aligning with the goal of expanding the KG with genuinely new and valuable information.

A. Expanding Relations

The initial phase begins by identifying unconnected node pairs within a given subgraph $\mathcal{G}_{\text{sub}}$. For each pair of entities $(e_1, ?, e_2)$ lacking a direct relation, then COMET's link prediction capabilities are leveraged to hypothesize potential relationships. Utilizing a *beam-k*, *top-k*, and *greedy* decoding strategy, COMET generates a singular, most probable relation for each unconnected node pair. The generated candidate relations then undergo a stringent filtering process. This process prioritizes precision by retaining only the top-ranked predictions based on their confidence scores. Specifically, relations corresponding to the top 2, 5, or 10 highest confidence scores, as output by the COMET model, are selected. Subsequently, these selected relations undergo a rigorous validation against the existing KG. Any generated fact that directly contradicts or negates an established fact within the original KG is identified and removed to preserve data integrity.

TABLE I: Accuracy test on CommonsenseQA (in house test) for expanding subgraph relations

Method	***k=2***	***k=5***	***k=10***
QA-GNN (baseline)		73.41	
+ COMET *top-k*	69.21	**73.20**	71.72
+ COMET *beam-k*	**73.67**	73.58	72.49
+ COMET greedy		**74.08**	

The experiments results exhibited variability, yet consistently demonstrated that each method possesses an optimal *k* value for performance as shown in Table I. Notably, the greedy method yielded the most favorable outcomes overall. This suggests that the inclusion of a greater number of relations is generally advantageous for enhancing reasoning capabilities within the knowledge graph. However, this parameter requires careful tuning to achieve optimal performance, as an indiscriminately high *k* may not always translate to improved results.

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue XI Nov 2026- Available at www.ijraset.com

B. Expanding Nodes

The process begins with a systematic examination of each individual node present within the initial subgraph $\mathcal{G}_{\text{sub}}$. This comprehensive iteration serves to identify all potential points from which the graph's structure can be meaningfully expanded. For each such node, COMET are utilized in conjunction with a predefined set of 17 merged relation types following QA-GNN [18], to predict the most suitable new triples. This predictive step is designed to infer novel connections and entities, thereby generating a set of candidate relational statements for graph augmentation.

These newly generated triples subsequently undergo a rigorous two-stage validation and selection procedure. Initially, each candidate triple is meticulously checked against the original KG; any prediction that directly contradicts or negates an existing, established fact is promptly identified and removed. This crucial filtering mechanism is essential for maintaining the factual integrity and logical consistency of the knowledge base. Following this contradiction assessment, the remaining generated triples proceed to have their relevance scores (2) computed. Finally, to ensure the integration of only the most impactful and well-supported information, select the triples exhibiting the top 2, 5, or 10 highest relevance scores to be included into the subgraph. This selective integration strategy aims to enhance the graph's completeness while preserving its overall quality.

Thus creating new nodes where each create triples of (e, r', e_x) where e_x are the new nodes, $e \in \{e_q\} \cup \{e_a\} \cup \{e_o\}$ and $r' \in R$ relations type that exist in the KG but this particular edges are not exist in the KG. Triples in format of (e_x, r, e_x) are guaranteed not exist.

TABLE II: Accuracy test on CommonsenseQA (in house test) for expanding subgraph nodes

Method	*k=2*	*k=5*	*k=10*
QA-GNN (baseline)		73.41	
+ COMET *top-k*	71.82	72.19	72.20
+ COMET *beam-k*	72.09	71.57	69.88
+ COMET greedy		73.58	

Table II showed a decline in performance when new nodes were introduced. However, the greedy approach consistently yielded the most notable outcomes, mirroring the findings from COMET evaluations. This suggests that expanding the knowledge graph by an additional hop may not aid in the reasoning process and could, in fact, introduce further ambiguity.

C. Expanding Nodes and Relations

Here knowledge graph expansion strategy integrates both methods explained above. A key principle of this combined approach is to ensure that the generation processes for new relations and the potential introduction of new nodes do not unduly influence one another. Specifically, the prediction of new relations is exclusively performed on node pairs that are already part of the initial subgraph from which the expansion begins.

This design choice ensures that the relational inference mechanism operates solely on the foundational set of nodes available at the outset, thereby preventing any newly predicted nodes from immediately affecting the scope or context for further relational predictions within the current expansion cycle. This deliberate separation aims to facilitate a systematic and controlled augmentation of the knowledge graph, which helps to mitigate the propagation of potential biases or unverified information throughout the overall expansion process.

TABLE III: Accuracy test on CommonsenseQA (in house test) for expanding both relations and nodes

Method	*k=2*	*k=5*	*k=10*
QA-GNN (baseline)		73.41	
+ COMET *top-k*	69.41	70.45	73.12
+ COMET *beam-k*	71.86	72.94	70.22
+ COMET greedy		73.61	

The outcomes as shown in Table III indicate that while the approach of expanding nodes alone yields superior results compared to not expanding nodes at all, it does not surpass the performance achieved by exclusively expanding relations. This suggests that integrating both node and relation expansion features contributes to a more generalized outcome. Furthermore, this combination appears to successfully mitigate the weaknesses associated with solely generating new nodes, even though the overall score does not exceed that of the relations-only strategy.

©IJRASET: All Rights are Reserved

V. CONCLUSIONS

This study introduced a novel methodology for expanding the reasoning capabilities of a subgraph, aiming to extend beyond the confines of the initial knowledge graph. By establishing new connections between entities, the study approach successfully bridged existing reasoning gaps. The experimental results clearly indicate that this expansion significantly impacted the reasoning performance of the QA-GNN model. While some cases demonstrated notable improvements, others suggest the need for further parameter adjustments to achieve optimal outcomes.

REFERENCES

[1] T. B. Brown et al., "Language Models are Few-Shot Learners." [Online]. Available: https://arxiv.org/abs/2005.14165
[2] W. X. Zhao et al., "A Survey of Large Language Models." [Online]. Available: https://arxiv.org/abs/2303.18223
[3] L. Ouyang et al., "Training language models to follow instructions with human feedback." [Online]. Available: https://arxiv.org/abs/2203.02155
[4] S. Zhang et al., "OPT: Open Pre-trained Transformer Language Models." [Online]. Available: https://arxiv.org/abs/2205.01068
[5] C. Rudin, "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead." Accessed: Dec. 23, 2024. [Online]. Available: https://arxiv.org/abs/2203.02155
[6] J. Li, X. Cheng, W. X. Zhao, J.-Y. Nie, and J.-R. Wen, "HaluEval: A Large-Scale Hallucination Evaluation Benchmark for Large Language Models." [Online]. Available: https://arxiv.org/abs/2305.11747
[7] A. Gunjal, J. Yin, and E. Bas, "Detecting and Preventing Hallucinations in Large Vision Language Models," Proceedings of the AAAI Conference on Artificial Intelligence, vol. 38, no. 16, pp. 18135–18143, Mar. 2024, doi: 10.1609/aaai.v38i16.29771.
[8] A. Talmor, J. Herzig, N. Lourie, and J. Berant, "CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge," 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), J. Burstein, C. Doran, and T. Solorio, Eds., Minneapolis, Minnesota: Association for Computational Linguistics, June 2019, pp. 4149–4158. doi: 10.18653/v1/N19-1421.
[9] T. Mihaylov, P. Clark, T. Khot, and A. Sabharwal, "Can a Suit of Armor Conduct Electricity? A New Dataset for Open Book Question Answering." Accessed: Dec. 23, 2024. [Online]. Available: http://arxiv.org/abs/1809.02789
[10] C. E. Quintero-Narvaez and R. Monroy, "Integrating Knowledge Graph Data with Large Language Models for Explainable Inference," in Proceedings of the 17th ACM International Conference on Web Search and Data Mining, in WSDM '24. New York, NY, USA: Association for Computing Machinery, Mar. 2024, pp. 1198–1199. doi: 10.1145/3616855.3636507.
[11] X. Guan et al., "Mitigating Large Language Model Hallucinations via Autonomous Knowledge Graph-based Retrofitting." Accessed: Dec. 23, 2024. [Online]. Available: http://arxiv.org/abs/2311.1331
[12] K. Bollacker, C. Evans, P. Paritosh, T. Sturge, and J. Taylor, "Freebase: a collaboratively created graph database for structuring human knowledge," in Proceedings of the 2008 ACM SIGMOD international conference on Management of data, in SIGMOD '08. New York, NY, USA: Association for Computing Machinery, June 2008, pp. 1247–1250. doi: 10.1145/1376616.1376746.
[13] R. Speer, J. Chin, and C. Havasi, "ConceptNet 5.5: An Open Multilingual Graph of General Knowledge." Accessed: Dec. 23, 2024. [Online]. Available: http://arxiv.org/abs/1612.03975
[14] J. Kim, Y. Kwon, Y. Jo, and E. Choi, "KG-GPT: A General Framework for Reasoning on Knowledge Graphs Using Large Language Models." Accessed: Dec. 23, 2024. [Online]. Available: http://arxiv.org/abs/2310.11220
[15] J. Jiang, K. Zhou, Z. Dong, K. Ye, W. X. Zhao, and J.-R. Wen, "StructGPT: A General Framework for Large Language Model to Reason over Structured Data." Accessed: Dec. 23, 2024. [Online]. Available: http://arxiv.org/abs/2305.09645
[16] J. Sun et al., "Think-on-Graph: Deep and Responsible Reasoning of Large Language Model on Knowledge Graph." Accessed: July 30, 2024. [Online]. Available: http://arxiv.org/abs/2307.07697
[17] T. Hwang, S. Jeong, S. Cho, S. Han, and J. C. Park, "DSLR: Document Refinement with Sentence-Level Re-ranking and Reconstruction to Enhance Retrieval-Augmented Generation." Accessed: July 16, 2024. [Online]. Available: http://arxiv.org/abs/2407.03627
[18] M. Yasunaga, H. Ren, A. Bosselut, P. Liang, and J. Leskovec, "QA-GNN: Reasoning with Language Models and Knowledge Graphs for Question Answering," in Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Online: Association for Computational Linguistics, 2021, pp. 535–546. doi: 10.18653/v1/2021.naacl-main.45.
[19] Y. Sun, Q. Shi, L. Qi, and Y. Zhang, "JointLK: Joint Reasoning with Language Models and Knowledge Graphs for Commonsense Question Answering," in Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, M. Carpuat, M.-C. de Marneffe, and I. V. Meza Ruiz, Eds., Seattle, United States: Association for Computational Linguistics, July 2022, pp. 5049–5060. doi: .
[20] X. Zhang et al., "GreaseLM: Graph REASoning Enhanced Language Models for Question Answering," no. arXiv:2201.08860. arXiv, Jan. 2022. . doi: 10.48550/arXiv.2201.08860.
[21] J. Jiang, K. Zhou, X. Zhao, and J.-R. Wen, "UniKGQA: Unified Retrieval and Reasoning for Solving Multi-hop Question Answering Over Knowledge Graph," presented at the The Eleventh International Conference on Learning Representations, Sept. 2022. Accessed: July 18, 2025. [Online]. Available: https://openreview.net/forum?id=Z63RvyAZ2Vh
[22] L. Yao, C. Mao, and Y. Luo, "KG-BERT: BERT for Knowledge Graph Completion." Accessed: Dec. 23, 2024. [Online]. Available: http://arxiv.org/abs/1909.03193
[23] D. Taunk, L. Khanna, P. Kandru, V. Varma, C. Sharma, and M. Tapaswi, "GrapeQA: GRaph Augmentation and Pruning to Enhance Question-Answering." Accessed: Dec. 23, 2024. [Online]. Available: http://arxiv.org/abs/2303.12320
[24] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph Attention Networks." Accessed: Dec. 23, 2024. [Online]. Available: http://arxiv.org/abs/1710.10903

[25] A. Talmor, J. Herzig, N. Lourie, and J. Berant, "CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge," 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), J. Burstein, C. Doran, and T. Solorio, Eds., Minneapolis, Minnesota: Association for Computational Linguistics, June 2019, pp. 4149–4158. . doi: 10.18653/v1/N19-1421.

[26] B. Y. Lin, X. Chen, J. Chen, and X. Ren, "KagNet: Knowledge-Aware Graph Networks for Commonsense Reasoning," 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), K. Inui, J. Jiang, V. Ng, and X. Wan, Eds., Hong Kong, China: Association for Computational Linguistics, Nov. 2019, pp. 2829–2839. doi: 10.18653/v1/D19-1282.

[27] A. Bosselut, H. Rashkin, M. Sap, C. Malaviya, A. Çelikyilmaz, and Y. Choi, "COMET: Commonsense Transformers for Automatic Knowledge Graph Construction," in Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL), 2019.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

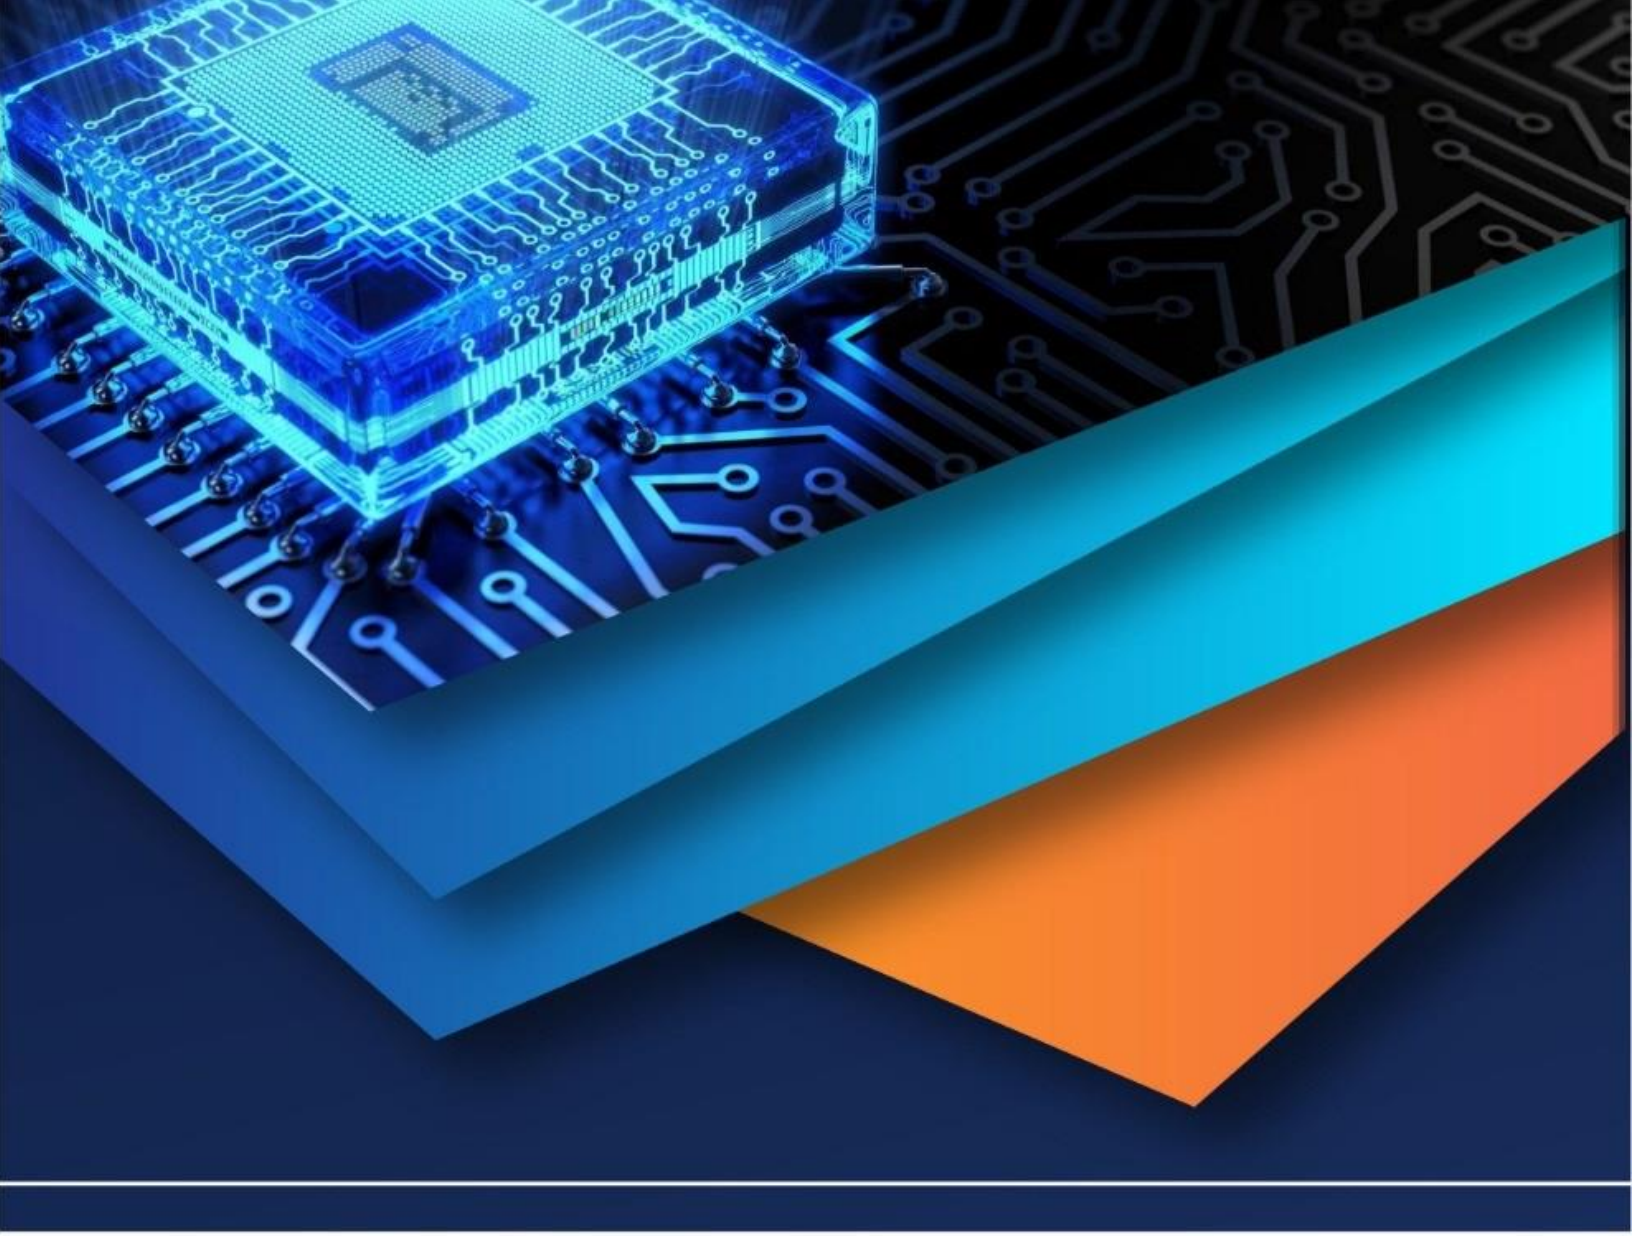

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)