



City Research Online

City, University of London Institutional Repository

Citation: Clay, A. & Jiménez-Ruiz, E. (2025). Information for Conversation Generation: Proposals Utilising Knowledge Graphs. Proceedings of the Special Session on Harmonising Generative AI and Semantic Web Technologies, 3953, ISSN 1613-0073

This is the published version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/35214/>

Link to published version:

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

Information for Conversation Generation: Proposals Utilising Knowledge Graphs

Alex Clay¹, Ernesto Jiménez-Ruiz¹

¹City St George's, University of London, Northampton Square, London, EC1V 0HB, United Kingdom

Abstract

LLMs are frequently used tools for conversational generation. Without additional information LLMs can generate lower quality responses due to lacking relevant content and hallucinations, as well as the perception of poor emotional capability, and an inability to maintain a consistent character. Knowledge graphs are commonly used forms of external knowledge and may provide solutions to these challenges. This paper introduces three proposals, utilizing knowledge graphs to enhance LLM generation. Firstly, dynamic knowledge graph embeddings and recommendation could allow for the integration of new information and the selection of relevant knowledge for response generation. Secondly, storing entities with emotional values as additional features may provide knowledge that is better emotionally aligned with the user input. Thirdly, integrating character information through narrative bubbles would maintain character consistency, as well as introducing a structure that would readily incorporate new information.

Keywords

Conversational AI, Retrieval augmented generation, theoretical proposal, information recommendation

1. Introduction

Large Language Models (LLMs) have quickly become the standard for the generation of conversational AI responses, due in part to the popularity of ChatGPT. However, without augmentation, LLMs are prone to a number of pitfalls, namely that of responses lacking valuable content[1] and hallucinations[2], as well as lack of emotional capability[3], and inconsistent character[4]. External knowledge such as knowledge graphs (KG) can be used to improve the first two instances through supplying additional information, and may be key to resolving the latter two through retrieval augmented generation.

Responses lacking content and hallucinations are often due to incomplete or out of date information, as LLMs use information stored within their parameters from when they were trained[5]. These challenges still persist when utilizing static knowledge bases. Knowledge graphs themselves avoid such difficulties yet lack the ability to interpret data beyond the relationships introduced by the data set, a difficulty solved by knowledge graph embeddings (KGE). As such, the adaptation of a KGE to readily incorporate new information without requiring retraining would utilize the benefits of both KGs and KGEs, and better address content poor responses and hallucinations. The selection of the information is as crucial as the information itself. Therefore, a recommender could utilise the information stored in a KG to determine what would be most relevant to a user input.

Supplying the emotion through auxiliary information to the LLM at time of generation has yet to be explored as a means for emotional integration. By selecting information that is more emotionally aligned with the user input through comparing emotional scores stored as features in the KG, it may be possible to generate a response that is more emotionally appropriate and increases the users' perception of the agent's emotional capability.

Additions of character to LLMs often utilize shallow character information in order to guide the generation of the response. Many well known approaches in industry are likely forms of prompt engineering. Grounding an LLM with a knowledge base specifically for providing consistent character information could improve on these shallow character representations. A proposed structure is to

ISWC: Special Session on LLMs

✉ alex.clay@city.ac.uk (A. Clay); ernesto.jimenez-ruiz@city.ac.uk (E. Jiménez-Ruiz)

🆔 0009-0004-2237-7412 (A. Clay); 0000-0002-9083-4599 (E. Jiménez-Ruiz)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

extract character information from novels delineating between utterances, character facts, and provide summaries to act as character memory. This information would then be stored as entities within a KG and organized in a narrative bubble structure to maintain information about what was learned when and improve the integration of personal experience.

This paper will introduce proposals to utilise KGs to support the integration of new information, emotion, and character based personal experience into LLM generated responses.

2. Proposals for Information for Conversation Generation

2.1. New Information Integration and Recommendation

Active integration of new information in knowledge bases is crucial for maintaining relevant and up-to-date information, and could reduce hallucinations on the part of the LLM[2]. While such information can readily be added to knowledge graphs, the same is not true of trained KGEs. As language models can neither easily integrate new information or modify their knowledge[6], it is a crucial facet of the external knowledge base. Therefore, a Dynamic KGE (DKGE) would be ideal, as DKGEs are KGEs which are able to add new data into the embeddings without requiring retraining of the entire graph[7], and could provide means to leverage the additional value of KGEs without the knowledge base being static. This reduces instances where the LLM lacks the necessary information to generate a correct response and causes frustration to the user[8].

In order to utilize the information stored within the DKGE, a recommender could take the user input and supply relevant information to the LLM at the time of generation. Knowledge Graph Embeddings (KGEs) have been utilised for recommendation in a number of approaches[9, 10, 11], which highlights the feasibility of integrating recommendation with a DKGE.

In recommendation, typically there are users and products with the intention of using available information to recommend the best product to the target user[12]. When a user is previously unknown or otherwise lacks interaction and preference information, it is regarded as a cold start problem[13]. To utilise this approach for information recommendation, the user input would take the role of an unknown user, for which supporting factual information would be provided instead of a product recommendation. This would be achieved through the use of a link prediction task where utterances and information are stored as different entity types in a heterogeneous knowledge graph, and a relation is predicted between the user input as an utterance type entity to a piece of information.

For example, if the user were to initiate a conversation about dinosaurs and make a statement such as “There is evidence the T. Rex may have been as intelligent as a crocodile.”, in the ideal scenario information pertaining to dinosaurs and ideally reptiles would be returned to aid in a reply. Additionally, the utterance and entities T. Rex and crocodile would be added with a relation along the lines of “as intelligent as” between the two, as well as relations to entities like “dinosaur” and “reptile” with a dynamic embedding generated by using the embedding values of similar existing entities potentially improving the information returned as the conversation continues.

2.2. Emotional Features for Entities

The display of emotion and emotional understanding are crucial for the success of conversational AI due to their necessity in human conversation[14, 15, 16]. Such cues are often expected by the user[3] and can increase the perception of friendliness and intelligence of an AI[17].

A potential means of integrating emotion into generated text is through influencing the supplemental information given to the LLM at the time of generation to be more emotionally in line with that of the user. In taking an otherwise purely factual knowledge base and adding emotional values as features to the entities prior to training, and to the user input at runtime, a recommendation structure based off of this may yield better supplementary data.

Valence, Arousal, and Dominance (VAD) scores [18] are commonly used to denote emotions as a point in space, which provides more specificity than emotional labels.

The introduction of the emotional score then not only increases the likelihood of the returned information being more emotionally similar to the input, but also supplies an additional feature to the input entity which may improve the cold start recommendation. For instance, in a situation without the emotional values, a statement of “The weather is lovely today” could theoretically return sunny and rainy with equal likelihood. However, with the introduction of an emotion value if the input above was assigned a value closer to that of pleasant, and with the expectation that sunny would have a closer value to pleasant than rainy, the recommended information would be more pertinent to the conversation.

2.3. Character Dimension Through Experience From Novels

Companies like Character.AI in industry supply a means by which to add a character atop an existing LLM, which is likely achieved through a form of prompt engineering. However, these characters often fail to maintain consistency at the time of generation, and are able to retain a limited number of tokens. A potential solution is to introduce character information from a structured storage to the LLMs generation of the conversational agent’s responses. Such a method could allow for the incorporation of more character information and responses potentially more similar to that of the character.

Novels lend the opportunity to gain abundant character information as well as the chance to evaluate the approach against existing expectations for the character. In research, some similar investigations have been made, with [19] deriving character information from biographies and [20] which utilised a set of five sentences to provide the basis of the agents characterisation. [21] investigated a similar concept in supplying personal experience through constructing scenes with which to fine-tune an LLM. However, it stands that by structuring the stored character information like memory the conversational integration of this information may seem more natural. Additionally [22] found that in conversation it may be necessary to indicate recall explicitly, therefore delineating between character facts and character memory may be crucial in supplying this information to an LLM to generate dialogue.

A novel could be structured into narrative ‘bubbles’ within a KG for a particular character, wherein everything encountered during a particular spatial or temporal bound would be stored together, much like human personal experience[23] which is known as episodic memory and handles the ‘what’, ‘when’, and ‘where’ of experiences[24, 25]. For example, the following text has been formatted as it might appear in the KG in Figure 1 for the character, Ajax. The two bubbles are split on the temporal delineation of “the next day”.

“The T-Rex may have actually been very intelligent,” Pierro stated confidently.

“I heard otherwise,” replied Rosalyne. “They may have been as intelligent as a crocodile.”

“I bet the Loch Ness monster is smarter than any dinosaur,” Ajax interjected immediately starting an argument, exactly as he had intended. While Rosalyne and Pierro generally tolerated one another as much as one would a coworker, Ajax’s shenanigans meant he did not receive the same courtesy.

The next day was much the same.

“OK, so hear me out,” Ajax began jovially. Pierro looked despondent. Rosalyne was actively considering quitting her job.

The information within the bounds of the bubble would be separated into entities of the type utterances, facts, and an overall summary. The separation of this information is crucial as it can provide cues to the LLM about how to integrate the information appropriately, such as whether to indicate recall in a response.

Within the bubble each entity would have a relation to every other member entity regardless of type. This relation would be of a type *shared_bubble* and act as the unifying factor of the bubble. In order to recommend information during a conversation, a preliminary response would be generated for which the *shared_bubble* relation would be predicted. The entities within the bubble of the recommended tail entity would then be returned and used to generate a new response with character information.

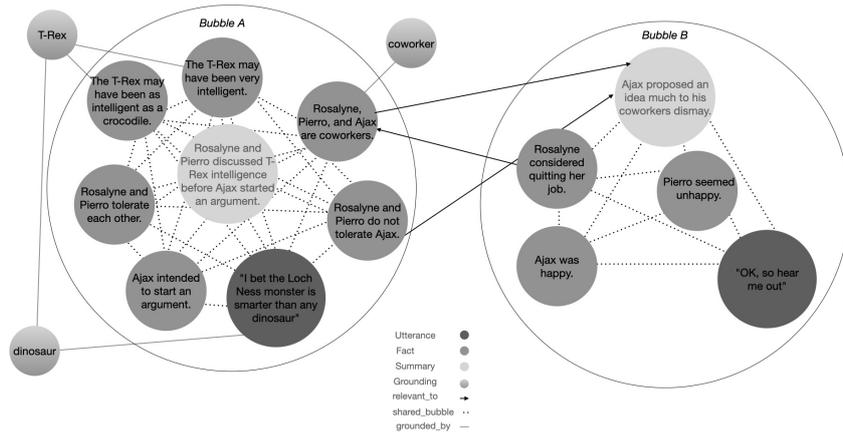


Figure 1: An example of two bubbles for the character Ajax. Each bubble only has one summary entity.

Following the above example, an initial response to a user input about dinosaurs may make a statement like “dinosaurs are cool” which in turn would recommend the A bubble based on the utterance “I bet the Loch Ness monster is smarter than any dinosaur”. This would result in the contents of the bubble being returned as information for the LLM to use to generate a character based response, possibly something like “my coworkers were discussing dinosaurs the other day” replacing the initially generated response.

In order to determine what relations should exist between bubbles, a comparison of the summary entities would be made and given enough relevance, the other entities within those bubbles would then be compared for a potential relation of the type *relevant_to* which would not necessarily be reciprocal. Additional entities would exist outside of the bubbles to ground the bubbles with the relation *grounded_by*, an example of which is the entity *coworker* in Figure 1. Another potential advantage of this approach would be the ability to dynamically add new bubbles. Following an embedding on the initializing information, a new bubble would first have relations formed with existing entities, the embeddings for which would then be aggregated to inform the embedding for the new entity. In a situation where an entity in a new bubble lacks any reference points from existing entities, its value would take an average from the other entities in the bubble. The summary entity of a bubble would always take its embedding from an aggregation of the entities in the bubble as it is the representation of the overall content. In order to reduce the need for retraining, existing entities would only update their embeddings after adding a certain threshold of new relations. If the bubble as a whole has had a number of updates to the member entities, the contents of the bubble would be reembedded with updated values as if it were a new bubble.

While integrating new information alone has been found to better support LLM generation, maintaining a structure focused on highlighting the relations between information learned and used together could support a more narratively focused information suggestion which could potentially more create more natural sounding dialogue.

3. Conclusion

LLMs provide a powerful tool for generating conversational responses, however their responses lack valuable content and hallucinations, as well as lack of emotional capability, and inconsistent character. This paper has introduced proposals for solving these challenges. Firstly, through utilizing dynamic knowledge graph embeddings with recommendation to provide up-to-date and relevant information to the LLM at time of generation to improve the content of responses and reduce likelihood of hallucinations. Additionally, using emotional values as features for entities to improve alignment with user input to increase the perception of emotional capability. Finally, the concept of storing character information in narrative bubbles was introduced, which could be updated without requiring retraining, to provide means for the portrayal of richer characters in LLM based conversation generation.

References

- [1] P. Parthasarathi, J. Pineau, Extending neural generative conversational model using external knowledge sources, in: E. Riloff, D. Chiang, J. Hockenmaier, J. Tsujii (Eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 690–695. URL: <https://aclanthology.org/D18-1073>. doi:10.18653/v1/D18-1073.
- [2] L. Huang, W. Yu, W. Ma, W. Zhong, Z. Feng, H. Wang, Q. Chen, W. Peng, X. Feng, B. Qin, T. Liu, A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions, 2023. arXiv:2311.05232.
- [3] E. Chen, H. Zhao, B. Li, X. Zha, H. Wang, S. Wang, Affective feature knowledge interaction for empathetic conversation generation, *Connection Science* 34 (2022) 2559–2576. URL: <https://doi.org/10.1080/09540091.2022.2134301>. doi:10.1080/09540091.2022.2134301. arXiv:<https://doi.org/10.1080/09540091.2022.2134301>.
- [4] Y. Xiao, Y. Cheng, J. Fu, J. Wang, W. Li, P. Liu, How far are llms from believable ai? a benchmark for evaluating the believability of human behavior simulation, 2024. URL: <https://arxiv.org/abs/2312.17115>. arXiv:2312.17115.
- [5] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel, D. Kiela, Retrieval-augmented generation for knowledge-intensive nlp tasks, in: *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS'20*, Curran Associates Inc., Red Hook, NY, USA, 2020.
- [6] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel, D. Kiela, Retrieval-augmented generation for knowledge-intensive nlp tasks, in: *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*, Curran Associates Inc., Red Hook, NY, USA, 2020.
- [7] Y. Cui, Y. Wang, Z. Sun, W. Liu, Y. Jiang, K. Han, W. Hu, Lifelong embedding learning and transfer for growing knowledge graphs, 2023. URL: <https://arxiv.org/abs/2211.15845>. arXiv:2211.15845.
- [8] D. Buhl, D. Szafarski, L. Welz, C. Lanquillon, Conversation-driven refinement of knowledge graphs: True active learning with humans in the chatbot application loop, in: H. Degen, S. Ntoa (Eds.), *Artificial Intelligence in HCI*, Springer Nature Switzerland, Cham, 2023, pp. 41–54.
- [9] Q. Ai, V. Azizi, X. Chen, Y. Zhang, Learning heterogeneous knowledge base embeddings for explainable recommendation, *Algorithms* 11 (2018) 137. URL: <http://dx.doi.org/10.3390/a11090137>. doi:10.3390/a11090137.
- [10] M. Gao, J.-Y. Li, C.-H. Chen, Y. Li, J. Zhang, Z.-H. Zhan, Enhanced multi-task learning and knowledge graph-based recommender system, *IEEE Transactions on Knowledge and Data Engineering* 35 (2023) 10281–10294. doi:10.1109/TKDE.2023.3251897.
- [11] X. Guo, W. Lin, Y. Li, Z. Liu, L. Yang, S. Zhao, Z. Zhu, Dken: Deep knowledge-enhanced network for recommender systems, *Information Sciences* 540 (2020) 263–277. URL: <https://www.sciencedirect.com/science/article/pii/S0020025520306289>. doi:<https://doi.org/10.1016/j.ins.2020.06.041>.
- [12] J. Frej, M. Knezevic, T. Kaser, Graph reasoning for explainable cold start recommendation, 2024. URL: <https://arxiv.org/abs/2406.07420>. arXiv:2406.07420.
- [13] Y. Bi, L. Song, M. Yao, Z. Wu, J. Wang, J. Xiao, A heterogeneous information network based cross domain insurance recommendation system for cold start users, in: *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '20*, Association for Computing Machinery, New York, NY, USA, 2020, p. 2211–2220. URL: <https://doi.org/10.1145/3397271.3401426>. doi:10.1145/3397271.3401426.
- [14] H. Rashkin, E. M. Smith, M. Li, Y.-L. Boureau, Towards empathetic open-domain conversation models: A new benchmark and dataset, in: A. Korhonen, D. Traum, L. Màrquez (Eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Florence, Italy, 2019, pp. 5370–5381. URL: <https://aclanthology.org/P19-1534>. doi:10.18653/v1/P19-1534.

- [15] H. Zhou, M. Huang, T. Zhang, X. Zhu, B. Liu, Emotional chatting machine: emotional conversation generation with internal and external memory, in: Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence, AAAI'18/IAAI'18/EAAI'18, AAAI Press, 2018.
- [16] Y. Liu, J. Gao, J. Du, L. Zhou, R. Xu, Empathetic response generation with state management, 2022. [arXiv:2205.03676](https://arxiv.org/abs/2205.03676).
- [17] K. Wang, X. Wan, Sentigan: Generating sentimental texts via mixture adversarial networks, in: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18, International Joint Conferences on Artificial Intelligence Organization, 2018, pp. 4446–4452. URL: <https://doi.org/10.24963/ijcai.2018/618>. doi:10.24963/ijcai.2018/618.
- [18] J. A. Russell, A. Mehrabian, Evidence for a three-factor theory of emotions, *Journal of Research in Personality* 11 (1977) 273–294. URL: <https://www.sciencedirect.com/science/article/pii/009265667790037X>. doi:[https://doi.org/10.1016/0092-6566\(77\)90037-X](https://doi.org/10.1016/0092-6566(77)90037-X).
- [19] A. Bogatu, D. Rotarescu, T. Rebedea, S. Ruseti, Conversational agent that models a historical personality, in: Romanian Conference on Human-Computer Interaction, 2015. URL: <https://api.semanticscholar.org/CorpusID:18293004>.
- [20] S. Zhang, E. Dinan, J. Urbanek, A. Szlam, D. Kiela, J. Weston, Personalizing dialogue agents: I have a dog, do you have pets too?, in: I. Gurevych, Y. Miyao (Eds.), Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Melbourne, Australia, 2018, pp. 2204–2213. URL: <https://aclanthology.org/P18-1205>. doi:10.18653/v1/P18-1205.
- [21] Y. Shao, L. Li, J. Dai, X. Qiu, Character-LLM: A trainable agent for role-playing, in: H. Bouamor, J. Pino, K. Bali (Eds.), Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Singapore, 2023, pp. 13153–13187. URL: <https://aclanthology.org/2023.emnlp-main.814>.
- [22] J. Campos, J. Kennedy, J. F. Lehman, Challenges in exploiting conversational memory in human-agent interaction, in: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '18, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2018, p. 1649–1657.
- [23] Y. Ezzyat, L. Davachi, What Constitutes an Episode in Episodic Memory?, *Psychological Science* 22 (2011) 243–252. doi:10.1177/0956797610393742.
- [24] E. Tulving, et al., Episodic and semantic memory, *Organization of memory* 1 (1972) 1.
- [25] E. Tulving, Episodic Memory: From Mind to Brain, *Annual Review of Psychology* 53 (2002) 1–25. doi:10.1146/annurev.psych.53.100901.135114.