

# Information for Conversation Generation: Proposals Utilising Knowledge Graphs

Alex Clay<sup>1,\*</sup>, Ernesto Jiménez-Ruiz<sup>1</sup>

<sup>1</sup>City 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 information, as LLMs use information stored within their parameters from when they were trained [5]. These challenges would still persist when utilizing static knowledge base, as out of date knowledge is known to contribute to hallucinations [2]. Knowledge graphs are able to dynamically add information through the addition of new triples, but require embedding to predict new relationships between existing entities [6]. As such, the adaptation of a knowledge graph embedding (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 could provide 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 user supplied background information which guides the generation of character responses. Grounding an LLM with a knowledge base specifically for providing

---

ISWC 2024 Special Session on Harmonising Generative AI and Semantic Web Technologies, November 13, 2024, Baltimore, Maryland

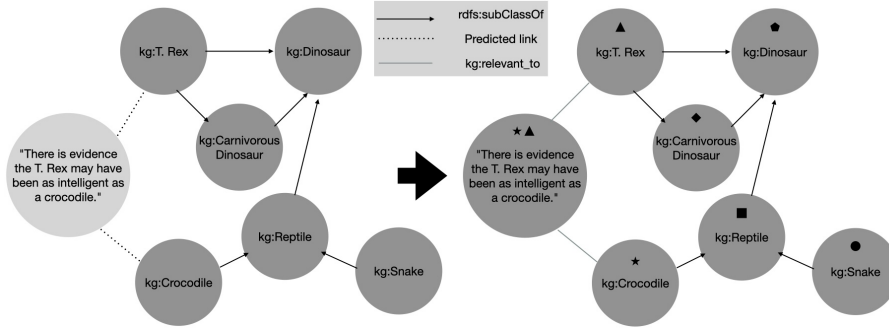
\*Corresponding author.

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

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



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



**Figure 1:** An example showing how an utterance would be dynamically integrated into the KGE, where the symbols represent the embeddings, and the utterance has an average between T Rex and crocodile.

consistent character information could improve on these shallow character representations. We propose a structure to extract character information from novels, delineating between utterances, character facts, and providing 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.

These proposals provide means to improve the user experience with conversational AIs. The active integration of new information would address frustration that users may encounter if the conversational AI displays a lack of knowledge [7] by allowing for knowledge base updates. Human-like AI is more likely to be accepted the user perceives it as having empathy [8] or social emotion, which influences user trust and subsequently acceptance [9]. By increasing the emotional alignment of responses and heightening the perception of emotional capability, the user may be more accepting of the AI. Consistency of character through storing character information could additionally improve user experience by potentially preventing instances where the LLM might hallucinate details or change information within the same interaction.

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 new information can readily be added to knowledge graphs as they do not require training, the same is not true of KGEs, which require training to create embeddings. As language models can neither easily integrate new information or modify their knowledge [10], 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 [11], 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 [7]. Augmenting LLMs with KGs not only supplies external knowledge for inference[12] but also improves their performance and interpretability[13]. Knowledge graphs in particular are able to maintain query performance regardless of the size of the dataset, a challenge encountered by other data structures like NoSQL [13]. Additionally, pretraining LLMs with knowledge graphs has been found to improve evaluation with closed book question answering[14]. As such, the use of KGs with LLMs, either as external knowledge or integrated at training time presents a unique opportunity to improve the information quality of LLM responses.

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 [15, 16, 17], 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 [18]. When a user is previously unknown or otherwise lacks interaction and preference information, it is regarded as a cold start problem [19]. 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 they might make a statement such as "There is evidence the T. Rex may have been as intelligent as a crocodile.". This statement would then be handled as an entity (`kg:utterance1`, `rdf:type`, `kg:Utterance`) (`kg:utterance1`, `kg:text`, "There is evidence the T. Rex may have been as intelligent as a crocodile.") for which a tail would be predicted for the relation, `relevant_to` for a piece of background information. This would ideally return entities like (`kg:T. Rex`, `rdfs:subClassOf`, `kg:CarnivorousDinosaur`) which would be given to the LLM as A T. Rex is a carnivorous dinosaur for incorporation into an information rich response.

In order to keep the information up to date, the utterance would be added as shown before, as well as the entities T. Rex and crocodile (if not already present) as follows (`kg:T. Rex`, `rdfs:subClassOf`, `kg:Dinosaur`) and (`kg:Crocodile`, `rdfs:subClassOf`, `kg:Reptile`). The best relations found during the initial prediction for recommendation would be used to integrate the new entities. The embedding for the new entity would then be an average of the embeddings of the related entities. Such as the utterance entity (`kg:utterance1`, `kg:text`, "There is evidence the T. Rex may have been as intelligent as a crocodile.") starting with an embedding derived from that of T. Rex and crocodile if they were already present in the knowledge base. A visual representation of this is shown in Figure 1.

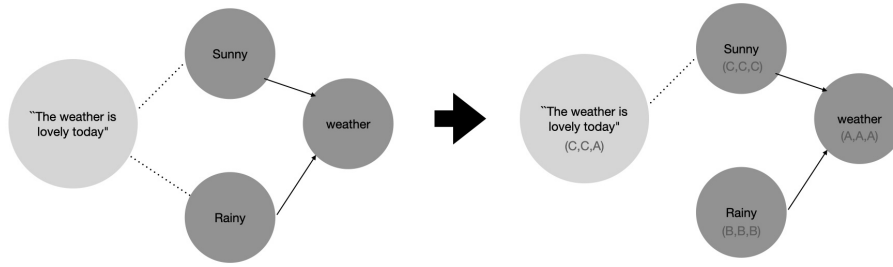
## 2.2. Emotional Features for Entities

The display of emotion and emotional understanding are crucial for the success of conversational AI due to their necessity in human conversation [20, 21, 22]. Such cues are often expected by the user [3] and can increase the perception of friendliness and intelligence of an AI [23].

The incorporation of emotion into generated dialogue has seen significant research in recent years. Some approaches have focused on the integration of both emotion and knowledge[24]often with a focus on commonsense knowledge[25, 26].

Another 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. Supplying more specificity than emotional labels, Valence, Arousal, and Dominance (VAD) scores [27] are commonly used to denote emotions as a point in space and could provide rich feature information. A recommendation structure based off of this may yield better supplementary data.

For instance, compare the statements "Rosalyne is picky" and "Rosalyne is meticulous". A human can recognise that while the meaning is very similar, meticulous has a positive connotation and picky does not, however an LLM would not necessarily have a means of differentiating such words without emotional delineation. As another example, 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. Figure 2 provides a visual



**Figure 2:** The values (C,C,A) represent the VAD scores for the utterance entity, which leads to the selection of sunny over rainy as the VAD scores are more similar.

example of how the addition of VAD score might change the recommendation of information to select a more similar value to that of the input.

The introduction of the emotional score would then not only increase the likelihood of the returned information being more emotionally similar to the input, but also supply an additional feature to the input entity which may improve the cold start recommendation.

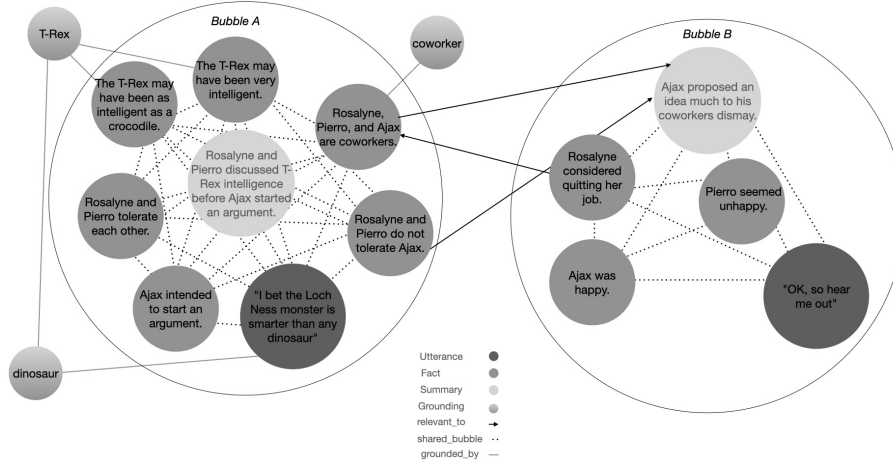
### 2.3. Character Dimension Through Experience From Novels

Companies like Character.AI in industry supply a means by which to add a character atop an LLM [28] through character information supplied by the user. These approaches can lead to generic responses that do not integrate relevant character information [28]. Moreover it is unlikely that the character information prompt is updated during interaction, instead remaining static and relying on further details to be generated in response to user input. This can cause inconsistency and might lead a lengthy conversation to become nonsensical.

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. The use of KGs in such a situation would reduce the number of needed tokens for an initial prompt, as the character would not require an initial explicit definition, but rather incorporate relevant character information as needed throughout the interaction. KGs would also provide consistent information, rather than relying on generation to supply details adhoc which may not be retained between responses. Additionally, providing supporting information from an auxiliary knowledge base allows for a larger amount of character information to be stored and integrated into conversation. A particular attribute of KGs that could further improve the quality of responses is the latent information in relations that KGs are able to utilise which could allow for richer integration of character information with the rest of the knowledge base.

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 [29] deriving character information from biographies and [30] which utilised a set of five sentences to provide the basis of the agents characterisation. [31] investigated a similar concept in supplying personal experience through constructing scenes with which to fine-tune an LLM. When compared with approaches like fine tuning, retrieval augmented generation is more adaptable, and better utilises external knowledge[32] which is crucial for handling large amounts of character information that might see changes over many interactions. As such, it stands that by structuring the stored character information like memory the conversational integration of this information may seem more natural. Additionally [33] 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 [34] which is known as episodic memory and handles the 'what',



**Figure 3:** An example of two bubbles for the character Ajax. Each bubble only has one summary entity.

‘when’, and ‘where’ of experiences [35, 36]. For example, the following text has been formatted as it might appear in the KG in Figure 3 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 Rosalynne. “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 Rosalynne 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. Rosalynne 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. Consider the first sentence of the third line of the example text. From the view of a KG constructed for the character Ajax, the first sentence would become two entities (`kg:utteranceA`, `rdf:type`, `kg:Utterance`)(`kg:utteranceA`, `kg:text`, "I bet the Loch Ness monster is smarter than any dinosaur") and (`kg:factA`, `rdf:type`, `kg:Fact`)(`kg:factA`, `kg:text`, "Ajax intended to start an argument") which would form the triple (`kg:utteranceA`, `kg:shared_bubble`, `kg:factA`) which indicates the entities shared presence within a bubble.

Entity `utterance1` would also be a member of the triple (`kg:utterance1`, `kg:grounded_by`, `kg:Dinosaur`) relating to the grounding entity `dinosaur`, as shown in Figure 3, which grounds it to information beyond the bubble relation and aids in informing its embedding through the relation `grounded_by`.

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. For instance, the entity of type `fact`, `Rosalynne, Pierro, and Ajax are coworkers` forms a triple



with the summary entity Ajax proposed an idea much to his coworkers dismay with the relation of `relevant_to` as the fact that they are coworkers provides useful supporting information.

For recommendation during a conversation, first a preliminary response would be generated for the user input. For example, if the user were to state “what do you think about dinosaurs?” an LLM might generate “dinosaurs are cool” as a response, which may not reflect character accurate information. This initial response would then be treated as an unseen utterance entity “dinosaurs are cool” for which link prediction for the relation `shared_bubble` would be conducted, in this instance leading to a predicted link with the utterance entity “I bet the Loch Ness monster is smarter than any dinosaur” due to the shared presence of `dinosaur`. The contents of the A bubble in Figure 3 would then be returned as information for the LLM to use in generating a new character based response, possibly something like “the Loch Ness monster is cooler than dinosaurs” replacing the initially generated response.

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.

### Declaration on Generative AI

The authors have not employed any Generative AI tools.

### 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] A. Bordes, N. Usunier, A. Garcia-Durán, J. Weston, O. Yakhnenko, Translating embeddings for modeling multi-relational data, in: *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS'13*, Curran Associates Inc., Red Hook, NY, USA, 2013, p. 2787–2795.
- [7] 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.
- [8] C. Pelau, D.-C. Dabija, I. Ene, What makes an AI device human-like? The role of interaction quality, empathy and perceived psychological anthropomorphic characteristics in the acceptance of artificial intelligence in the service industry, *Computers in Human Behavior* 122 (2021) 106855. doi:10.1016/j.chb.2021.106855.
- [9] S. Zhang, Z. Meng, B. Chen, X. Yang, X. Zhao, Motivation, Social Emotion, and the Acceptance of Artificial Intelligence Virtual Assistants—Trust-Based Mediating Effects, *Frontiers in Psychology* 12 (2021). URL: <https://www.frontiersin.org/articles/10.3389/fpsyg.2021.728495/pdf>. doi:10.3389/fpsyg.2021.728495.
- [10] 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.
- [11] 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](#).
- [12] S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, X. Wu, Unifying large language models and knowledge graphs: A roadmap, *IEEE Transactions on Knowledge and Data Engineering* 36 (2024) 3580–3599. URL: <http://dx.doi.org/10.1109/TKDE.2024.3352100>. doi:10.1109/tkde.2024.3352100.
- [13] N. Ibrahim, S. Aboulela, A. Ibrahim, R. Kashef, A survey on augmenting knowledge graphs (KGs) with large language models (LLMs): models, evaluation metrics, benchmarks, and challenges, *Discover Artificial Intelligence* 4 (2024). doi:10.1007/s44163-024-00175-8.
- [14] F. Moiseev, Z. Dong, E. Alfonseca, M. Jaggi, SKILL: Structured knowledge infusion for large language models, in: M. Carpuat, M.-C. de Marneffe, I. V. Meza Ruiz (Eds.), *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Association for Computational Linguistics, Seattle, United States, 2022, pp. 1581–1588. URL: <https://aclanthology.org/2022.naacl-main.113>. doi:10.18653/v1/2022.naacl-main.113.
- [15] 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.
- [16] 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.
- [17] 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>.
  - [18] J. Frej, M. Knezevic, T. Kaser, Graph reasoning for explainable cold start recommendation, 2024. URL: <https://arxiv.org/abs/2406.07420>. arXiv:2406.07420.
  - [19] 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.
  - [20] 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.
  - [21] 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.
  - [22] Y. Liu, J. Gao, J. Du, L. Zhou, R. Xu, Empathetic response generation with state management, 2022. arXiv:2205.03676.
  - [23] 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.
  - [24] D. Varshney, A. Ekbal, E. Cambria, Emotion-and-knowledge grounded response generation in an open-domain dialogue setting, *Know.-Based Syst.* 284 (2024). URL: <https://doi.org/10.1016/j.knosys.2023.111173>. doi:10.1016/j.knosys.2023.111173.
  - [25] P. Zhong, D. Wang, P. Li, C. Zhang, H. Wang, C. Miao, Care: Commonsense-aware emotional response generation with latent concepts, *Proceedings of the AAAI Conference on Artificial Intelligence* 35 (2021) 14577–14585. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/17713>. doi:10.1609/aaai.v35i16.17713.
  - [26] X. Zheng, Y. Du, X. Qin, Comasa: context multi-aware self-attention for emotional response generation, *Neurocomputing* 611 (2025) 128692. URL: <https://www.sciencedirect.com/science/article/pii/S0925231224014632>. doi:<https://doi.org/10.1016/j.neucom.2024.128692>.
  - [27] 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).
  - [28] X. Wang, H. Dai, S. Gao, P. Li, Characteristic AI agents via large language models, in: N. Calzolari, M.-Y. Kan, V. Hoste, A. Lenci, S. Sakti, N. Xue (Eds.), *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, ELRA and ICCL, Torino, Italia, 2024, pp. 3016–3027. URL: <https://aclanthology.org/2024.lrec-main.269>.
  - [29] 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>.
  - [30] 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.

- [31] 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>.
- [32] Y. Gao, Y. Xiong, X. Gao, K. Jia, J. Pan, Y. Bi, Y. Dai, J. Sun, M. Wang, H. Wang, Retrieval-augmented generation for large language models: A survey, 2024. URL: <https://arxiv.org/abs/2312.10997>. arXiv:2312.10997.
- [33] 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.
- [34] Y. Ezzyat, L. Davachi, What Constitutes an Episode in Episodic Memory?, *Psychological Science* 22 (2011) 243–252. doi:10.1177/0956797610393742.
- [35] E. Tulving, et al., Episodic and semantic memory, *Organization of memory* 1 (1972) 1.
- [36] E. Tulving, Episodic Memory: From Mind to Brain, *Annual Review of Psychology* 53 (2002) 1–25. doi:10.1146/annurev.psych.53.100901.135114.