Knowledge-enhanced Response Generation in Dialogue Systems: Current Advancements and Emerging Horizons

Priyanshu Priya[1], **Deeksha Varshney**[1], **Mauajama Firdaus**[2], **Asif Ekbal**[1]
[1]Indian Institute of Technology Patna, India
[2]University of Alberta, Edmonton, Canada
(priyanshu_2021cs26@iitp.ac.in, 1821cs13@iitp.ac.in, mauzama.03@gmail.com, asif@iitp.ac.in)

Abstract

This tutorial provides an in-depth exploration of Knowledge-enhanced Dialogue Systems (KEDS), diving into their foundational aspects, methodologies, advantages, and practical applications. Topics include the distinction between internal and external knowledge integration, diverse methodologies employed in grounding dialogues, and innovative approaches to leveraging knowledge graphs for enhanced conversation quality. Furthermore, the tutorial touches upon the rise of biomedical text mining, the advent of domain-specific language models, and the challenges and strategies specific to medical dialogue generation. The primary objective is to give attendees a comprehensive understanding of KEDS. By delineating the nuances of these systems, the tutorial aims to elucidate their significance, highlight advancements made using deep learning, and pinpoint the current challenges. Special emphasis is placed on showcasing how KEDS can be fine-tuned for domain-specific requirements, with a spotlight on the healthcare sector. The tutorial is crafted for both beginners and intermediate researchers in the dialogue systems domain, with a focus on those keen on advancing research in KEDS. It will also be valuable for practitioners in sectors like healthcare, seeking to integrate advanced dialogue systems.

1. Introduction

In the realm of artificial intelligence, dialogue systems have evolved as crucial interfaces facilitating human-machine interaction through natural language conversations. These systems are broadly categorized into task-oriented and open-domain dialogue systems. While task-oriented systems are designed to assist users in specific tasks like restaurant booking (Firdaus et al., 2020d, 2021c; Varshney and Singh, 2021), open-domain systems engage in a broader spectrum of conversational topics without a defined objective (Firdaus et al., 2020a; Varshney et al., 2020). The integration of deep learning, particularly neural language models, has significantly elevated the performance of these systems, yet challenges like understanding user opinions, integrating visual data, and ambiguity in open-domain interactions persist (Chen et al., 2017). In addressing the limitation of generating bland or generic responses common in traditional dialogue systems, Knowledge Enhanced Dialogue Systems (KEDS) have emerged as a prominent solution. The crux of KEDS lies in grounding the dialogues in external or internal knowledge, thereby enriching the conversation with insightful and contextually relevant responses. This tutorial provides an in-depth examination of KEDS, shedding light on its integral components, various approaches, and the benefits derived from such systems.

In this tutorial, we first introduce the foundational frameworks of Knowledge-enhanced Dialogue Systems (KEDS), establishing a solid understanding of how they augment dialogue systems. Following this, we explore the diverse methodologies employed to incorporate both internal and external knowledge sources, thereby enriching the conversational experience. We delve into internal knowledge sources embedded in the input text, such as topics, keywords, and internal graph structures, as discussed in (Ahmad et al., 2023; Mishra et al., 2022b; Firdaus et al., 2021a; Xie and Pu, 2021; Priya et al., 2023a). Concurrently, we investigate external knowledge acquisition from resources like uni-and-multi-modal knowledge bases, knowledge graphs, and grounded text such as persona information, Wikipedia information as elucidated in (Dinan et al., 2018; Zhou et al., 2018b; Firdaus et al., 2020f; Varshney and Singh, 2021; Ghazvininejad et al., 2018; Varshney et al., 2022a).

The discourse further extends to domain-specific applications, particularly in the healthcare sector. In the healthcare domain, having a thorough understanding of a person's medical history, mental state, symptoms, and treatment plan is crucial. Studies have indicated that the integration of extensive knowledge resources into healthcare dialogue systems presents multiple significant benefits. These include improving the system's understanding of medical terminology and concepts, equipping the system with the ability to reason and make inferences, grasping the emotional nuances within conversations, and discerning beneficial response patterns that contribute to emotional alleviation (Varshney et al., 2023b; Liang et al., 2021). Motivated by these insights, this tutorial session aims to explore various research endeavors that incorporate external knowledge into healthcare dialogue systems, thereby facilitating personalized and effective support (Shen et al., 2022; Deng et al., 2023; Varshney

LREC-COLING 2024 Tutorials, pages 80–87
20-25 May, 2024.

et al., 2022c, 2023b,c; Liu et al., 2021).

In the conclusion section, we highlight the shortcomings of conventional dialogue systems to provide a clearer pathway for newcomers to further research in KEDS systems.

2. Target Audience

We believe that the potential target audience could be the students at all levels (Doctorals, Masters, Bachelors), and anyone who is associated with healthcare, customer care, & related application areas, and researchers. We would assume an acquaintance with basic concepts about chatbots and neural networks, such as those included in most introductory Machine Learning (ML), Deep Learning (DL) and Natural Language Processing (NLP) courses. We expect an audience size of about 25-30 participants.

3. Outline

This tutorial is organized as follows:

- **Introduction (15 minutes)** We will briefly introduce dialogue systems, including the different types of dialogue systems and limitations of traditional dialogue systems (Chen et al., 2017). Afterward, we will discuss the notion of knowledge-enhanced response generation in dialogue systems and the different categories of knowledge sources, *viz.* internal knowledge and external knowledge. Precisely, we will delve into the concepts of (i) Internal knowledge sources embedded in the input text, including but not limited to topic, keyword, and internal graph structure (Xing et al., 2017; Xu et al., 2020; Li and Sun, 2018; Chen and Yang, 2023), and (ii) External knowledge acquisition, including but not limited to the multimodal information, persona, knowledge base, external knowledge graph, and grounded text (Firdaus et al., 2020b, 2022d; Dinan et al., 2018; Zhou et al., 2018b; Ghazvininejad et al., 2018).
- **Need and Challenges of Knowledge-enhanced Response Generation in Dialogues (15 minutes)**
  An effective dialogue system should be able to generate coherent, contextually relevant, user-centric, and informative responses. To achieve this, these systems require diverse information sources, including textual and structured data from external sources, user attributes (like sentiment, emotions, politeness, personal profile information - age, gender, persona, etc.), and contextual information (Wang et al., 2023a). Integrating the knowledge into the generated responses poses challenges concerning the retrieval or selection of pertinent knowledge and effective comprehension and utilization of the acquired knowledge to facilitate response generation (Wang et al., 2023b).
  In this section, we will discuss how the varied knowledge resources enhance response generation and improve the interpretability of dialogue systems by incorporating explicit semantics. Subsequently, we will address the challenges inherent in knowledge-enhanced response generation within dialogue systems.
- **Internal Knowledge-enhanced Response Generation in Dialogue Systems (60 minutes)**
  In this part of the tutorial, we aim to delineate the internal knowledge-enhanced response generation methods and applications. The information from internal knowledge sources helps enlighten and drive the generated responses to be informative and avoids generating universally relevant replies with little semantics. The internal knowledge can be obtained from topical information, keywords, and internal graph structures. We will point out the works that incorporate these knowledge sources for response generation.
  (i) *Response enhanced by Topic*: A dialogue system frequently employing responses such as "*I don't know*", "*Okay*" "*I see*" may appear repetitive and uninformative. While these off-topic replies are generally harmless for addressing various inquiries, they lack engagement and are likely to prematurely conclude conversations, significantly diminishing the overall user experience (Xing et al., 2017; Ahmad et al., 2023). Consequently, there is a pressing demand for on-topic response generation. This part of the tutorial delves into the works that have incorporated topical knowledge to guide the informative response generation (Xing et al., 2017; Xu et al., 2020).
  (ii) *Response enhanced by Keywords*: Recent research has incorporated personalized data into the dialogue generation process to enhance the quality of dialogue responses, particularly concerning emotional aspects, *viz.* emotion (Rashkin et al., 2019), sentiment (Chen and Nakamura, 2021), and politeness (Mishra et al., 2022b; Wang et al., 2020). We will discuss the works that attempt to integrate emotion (Zhou et al., 2018a; Firdaus et al., 2021a; Madasu et al., 2022; Majumder et al., 2022; Mishra et al., 2022c; Samad et al., 2022), sentiment (Firdaus et al., 2021b, 2022a), politeness (Golchha et al., 2019; Firdaus et al., 2020c; Mishra et al., 2022a; Firdaus et al., 2022a; Mishra et al., 2023a,c,b; Priya et al., 2023b), and intent (Xie and Pu, 2021) into the generated responses to make them personalized and engaging.

(iii) *Response enhanced by Internal Knowledge Graph*: Internal knowledge graphs are valuable for comprehending lengthy input sequences. They serve as intermediaries to consolidate or eliminate redundant data, resulting in a concise representation of the input document (Fan et al., 2019; Priya et al., 2023a). Furthermore, KG representations enable the creation of structured summaries and emphasize the connections between related concepts, particularly in cases where complex events associated with a single entity extend across multiple sentences (Huang et al., 2020). In this part of the tutorial, we will present works integrating an internal knowledge graph to enhance response generation capabilities (Liang et al., 2022; Firdaus et al., 2020e).

- **External Knowledge-enhanced Response Generation in Dialogue Systems (60 minutes)**
(i) *Persona Information*. Research focused on personas in dialogue systems requires that the agent adopts a specific character when engaging with users. This persona is closely linked to personality, which influences the emotional and personal aspects of users. In this section of the tutorial, we discuss studies that have employed persona-aware techniques to enhance the efficacy of response generation in dialogue systems (Firdaus et al., 2020f; Saha and Ananiadou, 2022; Firdaus et al., 2022d,b; Zhong et al., 2022). Findings from these studies suggest that persona information drives empathetic and personalized conversations more than non-empathetic ones.
(ii) *Multimodal Information*. Lately, the utilization of multimodal information has witnessed a surge in popularity in the field of dialogue systems. This approach is instrumental in comprehensively understanding users' emotional and mental states, as it leverages textual and non-textual attributes (Firdaus et al., 2023). In this part of the tutorial, we aim to discuss several notable studies in the literature that have harnessed multimodal data to enhance response generation within dialogue systems (Tavabi et al., 2019; Firdaus et al., 2020a, 2022c).
(iii) *External Knowledge Bases*. Knowledge-grounded systems utilize external resources such as Wikipedia documents to enhance response generation. (Dinan et al., 2018) released the first Wikipedia knowledge-grounded conversation dataset. (Varshney et al., 2023a) utilized the knowledge on various topics such as politics, and movies using the Topical Chat (Gopalakrishnan et al., 2019) and CMU_DoG (Zhou et al., 2018c) dataset to propose a knowledge-emotion enabled conversational model. (Lin et al., 2020) introduced a model that combined knowledge decoders with a pointer network to effectively handle out-of-vocabulary words. Experts suggest converting unstructured knowledge into organized knowledge graphs, composed of triplets (entity, relation, entity/item). Models, such as CCM, retrieve subgraphs from these graphs, especially using knowledge bases like ConceptNet (Speer and Havasi, 2012), and employ attention mechanisms to blend this knowledge into conversations (Zhou et al., 2018b). Concept Flow expands this by including extended subgraph ranges, integrating knowledge from two sources (Zhang et al., 2019). (Varshney et al., 2022a) utilizes both knowledge graphs and Wikipedia documents with a coreference-based knowledge graph augmenting method to improve factual accuracy in dialogue systems.
- **Knowledge-grounded Dialogue Systems in Healthcare (20 minutes)** In healthcare, background knowledge is vital in understanding an individual's medical history, mental condition, symptoms, and treatment plan. Research has shown that integrating comprehensive knowledge resources in the healthcare dialogue systems offers several key advantages, such as enhancing the system's grasp of medical concepts and terminology, empowering the system with reasoning and inference capabilities, comprehending emotional dynamics in conversations, and identifying useful response patterns leading to emotional relief (Varshney et al., 2022b; Liang et al., 2021). Driven by these considerations, in this tutorial session, we will discuss the studies that infuse external knowledge in healthcare dialogue systems for providing personalized and effective support (Shen et al., 2022; Deng et al., 2023; Varshney et al., 2022c, 2023b,c; Liu et al., 2021).
- **Hands-on Session (50 minutes)**
    1. Setting up a basic knowledge-enhanced dialogue system for healthcare domain (Varshney et al., 2023c,b).
    2. Integrating a sample knowledge base (e.g., Unified Medical Language System).
    3. Evaluating the performance of the dialogue using automated metrics such as BLEU, F1, and embedding-based metrics.
- **Conclusion and Future Perspectives (20 minutes)**
This tutorial explores notable studies on knowledge-enhanced dialogue generation, showcasing how leveraging diverse information sources can enhance dialogue model efficacy. Despite advancements, several challenges remain, highlighting exciting future re-

search avenues. We'll delve into four key research directions: (i) Knowledge Acquisition from Pre-trained Language Models: Pre-trained models harbor vast implicit knowledge without external memory reliance (Lewis et al., 2020), opening avenues for efficient knowledge extraction methods like knowledge distillation, data augmentation using pre-trained models as knowledge sources (Petroni et al., 2019), and prompting of language models (Li and Liang, 2021). (ii) Knowledge Acquisition from Limited Resources: In real-world scenarios, new domains often have scarce examples, necessitating rapid adaptation of knowledge-enhanced dialogue models via efficient meta-learning algorithms that minimize task-specific fine-tuning. (iii) Continuous Knowledge Acquisition: A noteworthy exploration is done in (Mazumder et al., 2018), where authors devised a knowledge acquisition engine for chatbots, enabling continuous learning from diverse information sources during interactions. (iv) Leveraging Emotional Knowledge through External Sources: Utilizing emotional knowledge bases like SenticNet aids in discerning user emotional states and background, thus generating emotionally coherent responses, crucial in healthcare and social good applications like persuasion and negotiation.

4. Proposed Length of the tutorial

Half-day (4h long including a coffee break (30m long))

5. Diversity Considerations

This tutorial on Knowledge-enhanced Dialogue Systems (KEDS) emphasizes inclusivity and diversity in three ways: (i) Enhancing Fairness: It educates on designing less biased, more inclusive dialogue systems, promoting equity in healthcare communication tools. (ii) Addressing Unique Needs: It's relevant to underrepresented groups like healthcare professionals and researchers from certain countries, offering tailored insights. (iii) Diverse Presenters: The presenters, originating from an underrepresented country, embody the commitment to diversity and inclusivity in computational linguistics.

6. Reading List

Extensive reading list is available at Reading List for Knowledge-enhanced Dialogue Systems.

7. Presenters

1. Priyanshu Priya, Indian Institute of Technology Patna, India (priyanshu_2021cs26@iitp.ac.in; priyanshu528priya@gmail.com; LinkedIn)
2. Deeksha Varshney, Indian Institute of Technology Patna, India (deeksha_1821cs13@iitp.ac.in; deeksha.varshney2695@gmail.com; LinkedIn)
3. Mauajama Firdaus, University of Alberta, Canada (mauzama.03@gmail.com; LinkedIn)
4. Asif EKbal, Indian Institute of Technology Patna, India. (asif@iitp.ac.in; asif.ekabl@gmail.com); Webpage: http://www.iitp.ac.in/ asif/; LinkedIn.

8. Other Information

While we are dedicated to accommodate a flexible number of participants, we anticipate an audience of 25-30 people. Our estimate is based on the previous attendance at the tutorial delivered on the topic "Empathetic Conversational Artificial Intelligence Systems: Recent Advances and New Frontiers" was presented at the 32^{nd} International Joint Conference on Artificial Intelligence, held from 19-25 August, 2023 at Macao, S.A.R, China., as well as the outreach efforts we have undertaken to promote the tutorial.

We would appreciate access to standard audio-visual equipment, such as microphones, projectors, and screens, to guarantee the tutorial's success. Furthermore, a high-speed internet connection is essential to ensure a seamless hands-on session during the tutorial, and an interactive whiteboard might be useful during the presentation for explanatory reasons. This configuration will assist us in facilitating interesting and informative discussions.

9. Ethics Statement

Dialogue systems are becoming ubiquitous in daily applications like healthcare and customer care, necessitating ethical considerations in development and usage. Key considerations include: (i) Knowledge-enhanced dialogue systems can collect sensitive user information, including personal and health data. To safeguard users' privacy, the data used in the research presented here has been anonymized, and personal details have been protected; (ii) In the context of knowledge-enhanced dialogue systems, user-centric design is essential, ensuring that users have control over the conversation and information sharing. Respecting user autonomy, these systems should offer options to conclude the conversation or seek further assistance. The datasets created for various research topics covered in this tutorial have been crafted to preserve user autonomy.

10. Bibliographical References

Zishan Ahmad, Kshitij Mishra, Asif Ekbal, and Pushpak Bhattacharyya. 2023. RPTCS: A reinforced persona-aware topic-guiding conversational system. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3482–3494, Dubrovnik, Croatia. Association for Computational Linguistics.

Hongshen Chen, Xiaorui Liu, Dawei Yin, and Jiliang Tang. 2017. A survey on dialogue systems: Recent advances and new frontiers. *Acm Sigkdd Explorations Newsletter*, 19(2):25–35.

Jiaao Chen and Diyi Yang. 2023. Controllable conversation generation with conversation structures via diffusion models. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7238–7251.

Sinan Chen and Masahide Nakamura. 2021. Generating personalized dialogues based on conversation log summarization and sentiment analysis. In *The 23rd International Conference on Information Integration and Web Intelligence*, pages 217–222.

Yang Deng, Wenxuan Zhang, Yifei Yuan, and Wai Lam. 2023. Knowledge-enhanced mixed-initiative dialogue system for emotional support conversations. *arXiv preprint arXiv:2305.10172*.

Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2018. Wizard of wikipedia: Knowledge-powered conversational agents. In *International Conference on Learning Representations*.

Angela Fan, Claire Gardent, Chloé Braud, and Antoine Bordes. 2019. Using local knowledge graph construction to scale seq2seq models to multi-document inputs. *arXiv preprint arXiv:1910.08435*.

Mauajama Firdaus, Hardik Chauhan, Asif Ekbal, and Pushpak Bhattacharyya. 2020a. Emosen: Generating sentiment and emotion controlled responses in a multimodal dialogue system. *IEEE Transactions on Affective Computing*, 13(3):1555–1566.

Mauajama Firdaus, Hardik Chauhan, Asif Ekbal, and Pushpak Bhattacharyya. 2020b. Meisd: A multimodal multi-label emotion, intensity and sentiment dialogue dataset for emotion recognition and sentiment analysis in conversations. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4441–4453.

Mauajama Firdaus, Hardik Chauhan, Asif Ekbal, and Pushpak Bhattacharyya. 2021a. More the merrier: Towards multi-emotion and intensity controllable response generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 12821–12829.

Mauajama Firdaus, Asif Ekbal, and Pushpak Bhattacharyya. 2020c. Incorporating politeness across languages in customer care responses: Towards building a multi-lingual empathetic dialogue agent. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4172–4182.

Mauajama Firdaus, Asif Ekbal, and Pushpak Bhattacharyya. 2022a. Polise: Reinforcing politeness using user sentiment for customer care response generation. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6165–6175.

Mauajama Firdaus, Umang Jain, Asif Ekbal, and Pushpak Bhattacharyya. 2021b. Seprg: sentiment aware emotion controlled personalized response generation. In *Proceedings of the 14th International Conference on Natural Language Generation*, pages 353–363.

Mauajama Firdaus, Arunav Shandilya, Asif Ekbal, and Pushpak Bhattacharyya. 2022b. Being polite: Modeling politeness variation in a personalized dialog agent. *IEEE Transactions on Computational Social Systems*.

Mauajama Firdaus, Gopendra Vikram Singh, Asif Ekbal, and Pushpak Bhattacharyya. 2023. Affect-gcn: a multimodal graph convolutional network for multi-emotion with intensity recognition and sentiment analysis in dialogues. *Multimedia Tools and Applications*, pages 1–22.

Mauajama Firdaus, Nidhi Thakur, and Asif Ekbal. 2020d. MultiDM-GCN: Aspect-guided response generation in multi-domain multi-modal dialogue system using graph convolutional network. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2318–2328, Online. Association for Computational Linguistics.

Mauajama Firdaus, Nidhi Thakur, and Asif Ekbal. 2020e. Multidm-gcn: Aspect-guided response generation in multi-domain multi-modal dialogue system using graph convolutional network. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2318–2328.

Mauajama Firdaus, Nidhi Thakur, and Asif Ekbal. 2021c. Aspect-aware response generation for multimodal dialogue system. *ACM Transactions* *on Intelligent Systems and Technology (TIST)*, 12(2):1–33.

Mauajama Firdaus, Nidhi Thakur, and Asif Ekbal. 2022c. Sentiment guided aspect conditioned dialogue generation in a multimodal system. In *Advances in Information Retrieval: 44th European Conference on IR Research, ECIR 2022, Stavanger, Norway, April 10–14, 2022, Proceedings, Part I*, pages 199–214. Springer.

Mauajama Firdaus, Naveen Thangavelu, Asif Ekba, and Pushpak Bhattacharyya. 2020f. Persona aware response generation with emotions. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.

Mauajama Firdaus, Naveen Thangavelu, Asif Ekbal, and Pushpak Bhattacharyya. 2022d. I enjoy writing and playing, do you: A personalized and emotion grounded dialogue agent using generative adversarial network. *IEEE Transactions on Affective Computing*.

Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. A knowledge-grounded neural conversation model. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

Hitesh Golchha, Mauajama Firdaus, Asif Ekbal, and Pushpak Bhattacharyya. 2019. Courteously yours: Inducing courteous behavior in customer care responses using reinforced pointer generator network. 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)*, pages 851–860.

Karthik Gopalakrishnan, Behnam Hedayatnia, Qinglang Chen, Anna Gottardi, Sanjeev Kwatra, Anu Venkatesh, Raefer Gabriel, Dilek Hakkani-Tür, and Amazon Alexa AI. 2019. Topical-chat: Towards knowledge-grounded open-domain conversations. In *INTERSPEECH*, pages 1891–1895.

Luyang Huang, Lingfei Wu, and Lu Wang. 2020. Knowledge graph-augmented abstractive summarization with semantic-driven cloze reward. *arXiv preprint arXiv:2005.01159*.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.

Jingyuan Li and Xiao Sun. 2018. A syntactically constrained bidirectional-asynchronous approach for emotional conversation generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 678–683.

Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*.

Yunlong Liang, Fandong Meng, Ying Zhang, Yufeng Chen, Jinan Xu, and Jie Zhou. 2021. Infusing multi-source knowledge with heterogeneous graph neural network for emotional conversation generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13343–13352.

Yunlong Liang, Fandong Meng, Ying Zhang, Yufeng Chen, Jinan Xu, and Jie Zhou. 2022. Emotional conversation generation with heterogeneous graph neural network. *Artificial Intelligence*, 308:103714.

Xiexiong Lin, Weiyu Jian, Jianshan He, Taifeng Wang, and Wei Chu. 2020. Generating informative conversational response using recurrent knowledge-interaction and knowledge-copy. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 41–52.

Wenge Liu, Jianheng Tang, Xiaodan Liang, and Qingling Cai. 2021. Heterogeneous graph reasoning for knowledge-grounded medical dialogue system. *Neurocomputing*, 442:260–268.

Avinash Madasu, Mauajama Firdaus, and Asif Eqbal. 2022. A unified framework for emotion identification and generation in dialogues. *arXiv preprint arXiv:2205.15513*.

Navonil Majumder, Deepanway Ghosal, Devamanyu Hazarika, Alexander Gelbukh, Rada Mihalcea, and Soujanya Poria. 2022. Exemplars-guided empathetic response generation controlled by the elements of human communication. *IEEE Access*, 10:77176–77190.

Sahisnu Mazumder, Nianzu Ma, and Bing Liu. 2018. Towards a continuous knowledge learning engine for chatbots. *arXiv preprint arXiv:1802.06024*.

Kshitij Mishra, Mauajama Firdaus, and Asif Ekbal. 2022a. Please be polite: Towards building a politeness adaptive dialogue system for goal-oriented conversations. *Neurocomputing*, 494:242–254.

Kshitij Mishra, Mauajama Firdaus, and Asif Ekbal. 2022b. Predicting politeness variations in goal-oriented conversations. *IEEE Transactions on Computational Social Systems*.

Kshitij Mishra, Mauajama Firdaus, and Asif Ekbal. 2023a. Genpads: Reinforcing politeness in an end-to-end dialogue system. *Plos one*, 18(1):e0278323.

Kshitij Mishra, Priyanshu Priya, and Asif Ekbal. 2023b. Help me heal: A reinforced polite and empathetic mental health and legal counseling dialogue system for crime victims. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 14408–14416.

Kshitij Mishra, Priyanshu Priya, and Asif Ekbal. 2023c. Pal to lend a helping hand: Towards building an emotion adaptive polite and empathetic counseling conversational agent. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12254–12271.

Kshitij Mishra, Azlaan Mustafa Samad, Palak Totala, and Asif Ekbal. 2022c. Pepds: A polite and empathetic persuasive dialogue system for charity donation. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 424–440.

Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. 2019. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066*.

Priyanshu Priya, Mauajama Firdaus, and Asif Ekbal. 2023a. A multi-task learning framework for politeness and emotion detection in dialogues for mental health counselling and legal aid. *Expert Systems with Applications*, 224:120025.

Priyanshu Priya, Kshitij Mishra, Palak Totala, and Asif Ekbal. 2023b. Partner: A persuasive mental health and legal counselling dialogue system for women and children crime victims. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, pages 6183–6191. International Joint Conferences on Artificial Intelligence Organization. AI for Good.

Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. Towards empathetic open-domain conversation models: A new benchmark and dataset. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5370–5381, Florence, Italy. Association for Computational Linguistics.

Tulika Saha and Sophia Ananiadou. 2022. Emotion-aware and intent-controlled empathetic response generation using hierarchical transformer network. In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.

Azlaan Mustafa Samad, Kshitij Mishra, Mauajama Firdaus, and Asif Ekbal. 2022. Empathetic persuasion: Reinforcing empathy and persuasiveness in dialogue systems. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 844–856.

Siqi Shen, Verónica Pérez-Rosas, Charles Welch, Soujanya Poria, and Rada Mihalcea. 2022. Knowledge enhanced reflection generation for counseling dialogues. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3096–3107.

Robyn Speer and Catherine Havasi. 2012. Representing general relational knowledge in ConceptNet 5. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 3679–3686, Istanbul, Turkey. European Language Resources Association (ELRA).

Leili Tavabi, Kalin Stefanov, Setareh Nasihati Gilani, David Traum, and Mohammad Soleymani. 2019. Multimodal learning for identifying opportunities for empathetic responses. In *2019 International Conference on Multimodal Interaction*, pages 95–104.

Deeksha Varshney, Asif Ekbal, Ganesh Prasad Nagaraja, Mrigank Tiwari, Abhijith Athreya Mysore Gopinath, and Pushpak Bhattacharyya. 2020. Natural language generation using transformer network in an open-domain setting. In *Natural Language Processing and Information Systems: 25th International Conference on Applications of Natural Language to Information Systems, NLDB 2020, Saarbrücken, Germany, June 24–26, 2020, Proceedings 25*, pages 82–93. Springer.

Deeksha Varshney, Asif Ekbal, Mrigank Tiwari, and Ganesh Prasad Nagaraja. 2023a. Emokbgan: Emotion controlled response generation using generative adversarial network for knowledge grounded conversation. *PloS one*, 18(2):e0280458.

Deeksha Varshney, Akshara Prabhakar, and Asif Ekbal. 2022a. Commonsense and named entity aware knowledge grounded dialogue generation. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1322–1335, Seattle, United

States. Association for Computational Linguistics.

Deeksha Varshney, Akshara Prabhakar, and Asif Ekbal. 2022b. Commonsense and named entity aware knowledge grounded dialogue generation. *arXiv preprint arXiv:2205.13928*.

Deeksha Varshney and Asif Ekbal Anushkha Singh. 2021. Knowledge grounded multimodal dialog generation in task-oriented settings. In *Proceedings of the 35th Pacific Asia Conference on Language, Information and Computation*, pages 425–435.

Deeksha Varshney, Aizan Zafar, Niranshu Kumar Behera, and Asif Ekbal. 2023b. Knowledge graph assisted end-to-end medical dialog generation. *Artificial Intelligence in Medicine*, 139:102535.

Deeksha Varshney, Aizan Zafar, Niranshu Kumar Behera, and Asif Ekbal. 2023c. Knowledge grounded medical dialogue generation using augmented graphs. *Scientific Reports*, 13(1):3310.

Deeksha Varshney, Aizan Zafar, Niranshu Kumar Behra, and Asif Ekbal. 2022c. Cdialog: A multi-turn covid-19 conversation dataset for entity-aware dialog generation. *arXiv preprint arXiv:2212.06049*.

Ming Wang, Bo Ning, and Bin Zhao. 2023a. A review of knowledge-grounded dialogue systems. In *2023 8th International Conference on Image, Vision and Computing (ICIVC)*, pages 819–824. IEEE.

Xintao Wang, Qianwen Yang, Yongting Qiu, Jiaqing Liang, Qianyu He, Zhouhong Gu, Yanghua Xiao, and Wei Wang. 2023b. Knowledgpt: Enhancing large language models with retrieval and storage access on knowledge bases. *arXiv preprint arXiv:2308.11761*.

Yi-Chia Wang, Alexandros Papangelis, Runze Wang, Zhaleh Feizollahi, Gokhan Tur, and Robert Kraut. 2020. Can you be more social? injecting politeness and positivity into task-oriented conversational agents. *arXiv preprint arXiv:2012.14653*.

Yubo Xie and Pearl Pu. 2021. Empathetic dialog generation with fine-grained intents. In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 133–147.

Chen Xing, Wei Wu, Yu Wu, Jie Liu, Yalou Huang, Ming Zhou, and Wei-Ying Ma. 2017. Topic aware neural response generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.

Minghong Xu, Piji Li, Haoran Yang, Pengjie Ren, Zhaochun Ren, Zhumin Chen, and Jun Ma. 2020. A neural topical expansion framework for unstructured persona-oriented dialogue generation. *arXiv preprint arXiv:2002.02153*.

Houyu Zhang, Zhenghao Liu, Chenyan Xiong, and Zhiyuan Liu. 2019. Grounded conversation generation as guided traverses in commonsense knowledge graphs. *arXiv preprint arXiv:1911.02707*.

Hanxun Zhong, Zhicheng Dou, Yutao Zhu, Hongjin Qian, and Ji-Rong Wen. 2022. Less is more: Learning to refine dialogue history for personalized dialogue generation. *arXiv preprint arXiv:2204.08128*.

Hao Zhou, Minlie Huang, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. 2018a. Emotional chatting machine: Emotional conversation generation with internal and external memory. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018b. Commonsense knowledge aware conversation generation with graph attention. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 4623–4629.

Kangyan Zhou, Shrimai Prabhumoye, and Alan W Black. 2018c. A dataset for document grounded conversations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 708–713, Brussels, Belgium. Association for Computational Linguistics.