

2022

Knowledge-infused Reinforcement Learning

Kaushik Roy

University of South Carolina - Columbia

Manas Gaur

University of Maryland Baltimore County

Qi Zhang

Amit Sheth

University of South Carolina - Columbia

Follow this and additional works at: https://scholarcommons.sc.edu/aii_fac_pub



Part of the [Computer Engineering Commons](#), and the [Electrical and Computer Engineering Commons](#)

Publication Info

Preprint version *Knowledge Graph Conference (KGC) 2022, 2022.*

© The Authors, 2023

This Paper is brought to you by the Artificial Intelligence Institute at Scholar Commons. It has been accepted for inclusion in Publications by an authorized administrator of Scholar Commons. For more information, please contact digres@mailbox.sc.edu.

Knowledge-infused Reinforcement Learning

Kaushik Roy¹ Manas Gaur² Qi Zhang¹ Amit Sheth¹
kaushikr@email.sc.edu manas@umbdc.edu qz5@cse.sc.edu amit@sc.edu

¹ Artificial Intelligence Institute South Carolina

² University of Maryland Baltimore County

Abstract

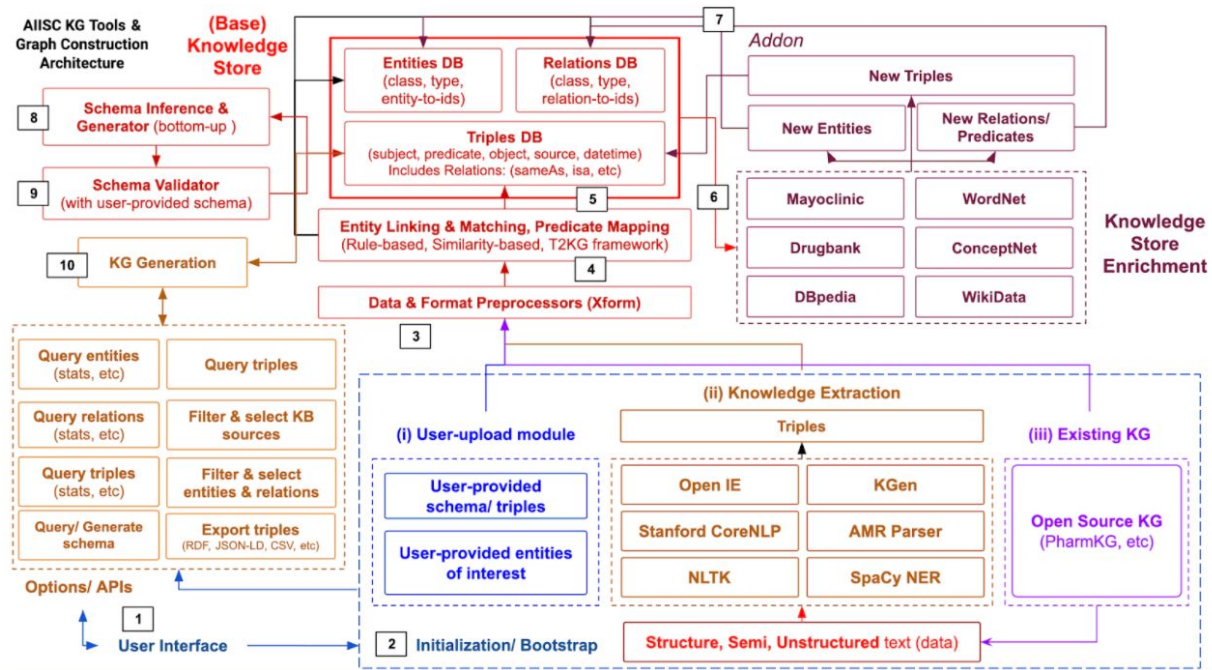
Virtual health agents (VHAs) have received considerable attention, but the early focus has been on collecting data, helping patients follow generic health guidelines, and providing reminders for clinical appointments. While presenting the collected data and frequency of visits to the clinician is useful, further context and personalization are needed for a VHA to interpret and understand what the data means in clinical terms. This has made their use in managing health limited. Such understanding enables patient empowerment and self-appraisal – i.e., aiding the patient in interpreting the data to understand the changes in the patient’s health conditions, and self-management – i.e., to help a patient better manage their health through better adherence to the clinician guidelines and clinician recommended care plan. Crisis conditions such as the current pandemic have further stressed our healthcare system and have made the need for such advanced support more attractive and in demand. Consider the rapid growth in mental health because the patients who already had mental health conditions worsen, and many develop such conditions due to the challenges arising from lockdown, isolation, and economic hardships. The severe lack of timely availability of clinical expertise to meet the rapidly growing demand provides the motivation for advancing this research in developing more advanced VHAs and evaluating it in the context of mental health management.

1 Intellectual Contributions of this tutorial

This tutorial seeks to showcase AI strategies that provide medical context to patient data with the help of a knowledge graph. This supports personalization through a personalized knowledge graph that captures the patient’s personalized health management objectives within the context of the clinical guidelines and care plan. The continuous capture of this information through the analysis of patient-VHA interactions, and the strategy of creating engaging interactions (conversations) can further augment the personalized knowledge graph. These operations are required to support self-appraisal and self-management, and when necessary perform fail-safe tasks such as connecting the patient to a crisis help-line or professional help. The core innovation is the use of a novel knowledge-infused reinforcement learning method. The by-product of this approach leads to transparency in decision-making with the ability to offer a user understandable explanation.

2 Tutorial Focus

We combine the qualities of Knowledge infused Learning [1], Reinforcement Learning [2, 3], and advanced use of Knowledge Graphs [4, 5, 6, 7] (e.g., through the incorporation of the process knowledge) to develop a VHA which will have transparency in decision making due to the use of knowledge and generate explanations in terms that the user can understand. The trans-



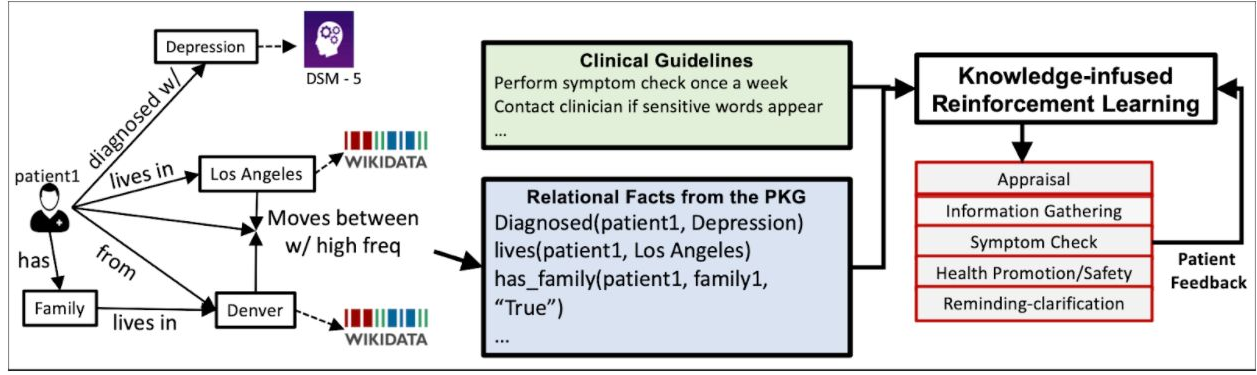


Figure 2: Illustration of Knowledge-infused Reinforcement Learning using knowledge graph and process knowledge in clinical guidelines to support patients in their goal management.

intelligence. Utilizing knowledge and data in deep learning models to enable learning from lower-level syntactic and lexical features from data through statistical (deep) learning as well as higher-level concepts from knowledge. Using knowledge also allows greater transparency in the decision-making, and enables explanations that users need for informed decision making.

3.1.1 Anatomy of KiL

Shallow Infusion: The knowledge is converted to an embedding vector and is concatenated to the data vector before being passed into the traditional deep learning pipeline. This method is fast and scalable and while it shows improvements over only using the data, the compression of knowledge in a vector-based embedding loses a lot of the semantic and relational level information.

Semi-Deep Infusion: This technique guides the parameter learning process of the deep learning model using the knowledge. Such a process ensures the parameters that govern data patterns are somewhat in alignment with the knowledge being infused. However, parametric models are layers and layers deep, representing different levels of abstraction, Therefore what kind of knowledge to guide the parameters of which specific layers remains an open question.

Deep Infusion: This technique guides the parameters of each layer of the deep learning model utilizing stratified layers of abstraction in the knowledge graph, thus providing the optimal level of knowledge infusion.

3.2 KiL for End User-level Explainability

Several methods to explain the outcomes of deep learning models have been proposed through the years, broadly classified as explainable AI (XAI). Similar to how a debugger understands the exceptions thrown by the program errors, XAI explanations can be thought of as a stack trace that the developer of the system can understand by not the end-user and application/domain expert. XAI explanations make sense to deep learning model debuggers and computer scientists that develop the models. However, to end-users like the clinicians, explanations that are meaningful at the application or domain levels (e.g., compliance with clinical guidelines) are required. As an example, visualization attention models to predict suicidality tells the computer scientist that the model and its components are functioning seemingly well. To the clinician, it still doesn't explain it through a line of reasoning that they can comprehend. For instance, considering the decision making based on Columbia Suicide Severity Rating Scale (C-SSRS), it is important for a model to explain its outcome based on the

flow of questions in C-SSRS. This would help clinicians assimilate model outcomes and help them make a conscious decision.

3.3 Knowledge-infused Reinforcement learning (KiRL)

Reinforcement Learning trains the VHA to understand the patient through trial and error based correction using directly the patient feedback. This adds a significant level of personalization based on the continuous user-feedback-based correction that is not present in supervised learning methods.

3.3.1 Why Reinforcement learning is an alternative way for Human-AI collaboration

Reinforcement Learning trains the VHA to understand the patient through trial and error based correction using directly the patient feedback. This adds a significant level of personalization based on the user-feedback-based correction that is not present in supervised learning methods.

3.3.2 KiRL

We combine the qualities of the aforementioned Knowledge infused Learning, Reinforcement Learning, to develop a VHA which will have transparency in decision making due to the use of knowledge and generate explanations in terms that the user can understand i.e. User level explanations as an upgrade from the traditional XAI based explanations. The transparency of the system will also allow component-wise understanding of the model which the model developers can benefit from as well. Finally, explicitly using the user's feedback to reiterate and revise the user-level personalized knowledge graph, the model and quality of explanations can deliver a VHA that clinicians and patients can use in practice to help their healthcare needs.

4 Length of the Tutorial

We plan a 1-hour lecture-style tutorial with 1 break (5-10 mins for Questions and Answers). Following it, we plan to have 30 minutes of hands-on practice. (Total: 1 hour 30 minutes)

5 Presenter's biographies and Related Papers

Kaushik Roy is a Ph.D. student at the Artificial Intelligence Institute South Carolina (AIISC), Qi Zhang is an assistant professor at the AIISC, Manas Gaur is an assistant professor at the University of Maryland Baltimore County, and Amit Sheth is the director at AIISC. The authors have published several related papers at premier conferences such as AAAI [9, 10, 11, 12, 13, 14], IEEE [15, 16, 17], ACL [18, 19], KGC [20], NeurIPS [4], and ICLR [21]. The author's have also given several related talks and tutorials [22, 23, 24, 25, 26].

6 Expected background and prerequisite of audience

The tutorial would be a mix of lecture-style and hands-on in the python programming language. The audience is expected to have a basic understanding of deep/machine learning, natural language processing, and semantic technologies (e.g., linked open data). We aim to guide attendees through a high-level tour of the most recent approaches proposed by researchers. Also, we expect basic familiarity with social media platforms such as Twitter and Reddit. We expect participants to bring their laptops with all the required tools installed. Details on tools needed and the background material will be provided upon acceptance of the tutorial proposal. We expect that by the end of the tutorial, the attendees will understand the use of knowledge graphs to enhance the performance (quality of results), utility, interpretability, and explainability of deep learning and be prepared to

apply knowledge-infused deep learning to real-world applications.

Acknowledgements

This work was supported in part by the National Science Foundation under Grant 2133842, “EAGER: Advancing Neuro-symbolic AI with Deep Knowledge-infused Learning” [27, 1, 28]

References

- [1] Amit Sheth, Manas Gaur, Kaushik Roy, and Keyur Faldu. Knowledge-intensive language understanding for explainable ai. *IEEE Internet Computing*, 25(5):19–24, 2021.
- [2] Kaushik Roy, Qi Zhang, Manas Gaur, and Amit Sheth. Knowledge infused policy gradients with upper confidence bound for relational bandits. In *Machine Learning and Knowledge Discovery in Databases. Research Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part I 21*, pages 35–50. Springer, 2021.
- [3] Srijita Das, Sriraam Natarajan, Kaushik Roy, Ronald Parr, and Kristian Kersting. Fitted q-learning for relational domains. *arXiv preprint arXiv:2006.05595*, 2020.
- [4] Vipula Rawte, Megha Chakraborty, Kaushik Roy, Manas Gaur, Keyur Faldu, Prashant Kikani, Hemang Akbari, and Amit Sheth. Tdlr: Top (semantic)-down (syntactic) language representation. *UMBC Faculty Collection*, 2022.
- [5] Ugur Kursuncu, Manas Gaur, and Amit Sheth. Knowledge infused learning (k-il): Towards deep incorporation of knowledge in deep learning. *arXiv preprint arXiv:1912.00512*, 2019.
- [6] Kaushik Roy, Manas Gaur, Qi Zhang, and Amit Sheth. Process knowledge-infused learning for suicidality assessment on social media. *arXiv preprint arXiv:2204.12560*, 2022.
- [7] Keyur Faldu, Amit Sheth, Prashant Kikani, and Hemang Akbari. Ki-bert: Infusing knowledge context for better language and domain understanding. *arXiv preprint arXiv:2104.08145*, 2021.
- [8] Kaushik Roy, Manas Gaur, Vipula Rawte, Ashwin Kalyan, and Amit Sheth. Proknow: Process knowledge for safety constrained and explainable question generation for mental health diagnostic assistance. *UMBC Faculty Collection*, 2022.
- [9] Kaushik Roy, Yuxin Zi, Manas Gaur, Jinendra Malekar, Qi Zhang, Vignesh Narayanan, and Amit Sheth. Process knowledge-infused learning for clinician-friendly explanations. *arXiv preprint arXiv:2306.09824*, 2023.
- [10] Kaushik Roy, Vedant Khandelwal, Raxit Goswami, Nathan Dolbir, Jinendra Malekar, and Amit Sheth. Demo alleviate: Demonstrating artificial intelligence enabled virtual assistance for telehealth: The mental health case. *arXiv preprint arXiv:2304.00025*, 2023.
- [11] Kaushik Roy, Usha Lokala, Vedant Khandelwal, and Amit Sheth. ” is depression related to cannabis?”: A knowledge-infused model for entity and relation extraction with limited supervision. *arXiv preprint arXiv:2102.01222*, 2021.
- [12] Kaushik Roy, Qi Zhang, Manas Gaur, and Amit Sheth. Knowledge infused policy gradients for adaptive pandemic control. *arXiv preprint arXiv:2102.06245*, 2021.
- [13] Parth Asawa, Manas Gaur, Kaushik Roy, and Amit Sheth. Covid-19 in spain and india: comparing policy implications by analyzing epidemiological and social media data. *arXiv preprint arXiv:2010.14628*, 2020.
- [14] Kaushik Roy, Yuxin Zi, Vignesh Narayanan, Manas Gaur, and Amit Sheth. Ksat: Knowledge-infused self attention transformer-integrating multiple domain-specific contexts. *arXiv preprint arXiv:2210.04307*, 2022.

- [15] Revathy Venkataramanan, Kaushik Roy, Kanak Raj, Renjith Prasad, Yuxin Zi, Vignesh Narayanan, and Amit Sheth. Cook-gen: Robust generative modeling of cooking actions from recipes. *arXiv preprint arXiv:2306.01805*, 2023.
- [16] Kaushik Roy, Tarun Garg, Vedant Palit, Yuxin Zi, Vignesh Narayanan, and Amit Sheth. Knowledge graph guided semantic evaluation of language models for user trust. *arXiv preprint arXiv:2305.04989*, 2023.
- [17] Manas Gaur, Kaushik Roy, Aditya Sharma, Biprav Srivastava, and Amit Sheth. “who can help me?”: Knowledge infused matching of support seekers and support providers during covid-19 on reddit. In *2021 IEEE 9th International Conference on Healthcare Informatics (ICHI)*, pages 265–269. IEEE, 2021.
- [18] Adam Tsakalidis, Jenny Chim, Iman Munire Bilal, Ayah Zirikly, Dana Atzil-Slonim, Federico Nanni, Philip Resnik, Manas Gaur, Kaushik Roy, Becky Inkster, et al. Overview of the clpsych 2022 shared task: Capturing moments of change in longitudinal user posts. In *Proceedings of the Eighth Workshop on Computational Linguistics and Clinical Psychology*, pages 184–198, 2022.
- [19] Shrey Gupta, Anmol Agarwal, Manas Gaur, Kaushik Roy, Vignesh Narayanan, Ponnuram Kumaraguru, and Amit Sheth. Learning to automate follow-up question generation using process knowledge for depression triage on reddit posts. *arXiv preprint arXiv:2205.13884*, 2022.
- [20] Nathan Dolbir, Triyasha Dastidar, and Kaushik Roy. Nlp is not enough—contextualization of user input in chatbots. *arXiv preprint arXiv:2105.06511*, 2021.
- [21] Usha Lokala, Francois Lamy, Triyasha Ghosh Dastidar, Kaushik Roy, Raminta Daniulaityte, Srinivasan Parthasarathy, and Amit Sheth. edark-trends: Harnessing social media trends in substance use disorders for opioid listings on cryptomarket. *arXiv preprint arXiv:2103.15764*, 2021.
- [22] Usha Lokala. Lecture-style tutorial: Causal ai for web and health care.
- [23] Valeria Fionda, Olaf Hartig, Reyhaneh Abdolazimi, Sihem Amer-Yahia, Hongzhi Chen, Xiao Chen, Peng Cui, Jeffrey Dalton, Xin Luna Dong, Lisette Espin-Noboa, et al. Tutorials at the web conference 2023. In *Companion Proceedings of the ACM Web Conference 2023*, pages 648–658, 2023.
- [24] Kaushik Roy, Usha Lokala, Manas Gaur, and Amit P Sheth. Tutorial: Neuro-symbolic ai for mental healthcare. In *Proceedings of the Second International Conference on AI-ML Systems*, pages 1–3, 2022.
- [25] Usha Lokala, Manas Gaur, Kaushik Roy, and Amit Sheth. Knowledge infused natural language understanding for public health, epidemiology, and substance use. 2021.
- [26] Amit Sheth, Kaushik Roy, Manas Gaur, and Usha Lokala. Tutorial on knowledge in-wisdom out-explainable data for ai in cyber social threats and public health. 2021.
- [27] Amit Sheth, Kaushik Roy, and Manas Gaur. Neurosymbolic ai-why, what, and how. *arXiv preprint arXiv:2305.00813*, 2023.
- [28] Amit Sheth, Manas Gaur, Kaushik Roy, Revathy Venkataraman, and Vedant Khandelwal. Process knowledge-infused ai: Toward user-level explainability, interpretability, and safety. *IEEE Internet Computing*, 26(5):76–84, 2022.