Process Knowledge-Infused AI: Toward User-Level Explainability, Interpretability, and Safety

Amit Sheth, Manas Gaur, Kaushik Roy, Revathy Venkataraman, and Vedant Khandelwal, University of South Carolina, Columbia, SC, 29201, USA

AI has seen wide adoption for automating tasks in several domains. However, AI's use in high-value, sensitive, or safety-critical applications such as self-management for personalized health or personalized nutrition has been challenging. These require that the AI system follows guidelines or well-defined processes set by experts, community, or standards. We characterize these as process knowledge (PK). For example, to diagnose the severity of depression, the AI system should incorporate PK that is part of the clinical decision-making process, such as the Patient Health Questionnaire (PHQ-9). Likewise, a nutritionist's knowledge and dietary guidelines are needed to create food plans for diabetic patients. Furthermore, the BlackBox nature of purely data-reliant statistical AI systems falls short in providing user-understandable explanations, such as what a clinician would need to ensure and document compliance with medical guidelines before relying on a recommendation. Using the examples of mental health and cooking recipes for diabetic patients, we show why, what, and how to incorporate PK along with domain knowledge in machine learning. We discuss methods for infusing PK and present performance evaluation metrics. Support for safety and user-level explainability of the PK-infused learning improves confidence and trust in the AI system.

Benchmarking datasets that assess the natural language understanding capabilities of large language models fall short in accelerating models to achieve user-level explainability, safety, uncertainty, and risk handling.[1,a] These challenges are associated with the limitations of AI in restricting its learning tasks to classification and generation, which are single shots. In comparison, real-world applications demand an orchestrated response going through a multistep process of learning the high-level needs of the user, then drilling down to specific needs, and subsequently yielding a structured response having a conceptual flow. For example, triaging patients in mental health requires clinical process knowledge manifested in a clinical questionnaire. Figure 1 illustrates a scenario where the agent maps user input to a sequence of yes or no questions to compile suicide risk severity. The agent can keep track of user-provided cues and ask appropriate follow-up questions through these ordered sets of questions. Upon receiving the required information to derive appropriate severity labels, the agent' outcome can be explained to MHPs for appropriate intervention. Similar but more complex applications include using Autism Diagnostic Observation Schedule to evaluate children with autism or using Montreal Cognitive Assessment score to measure the cognitive decline in poststroke Aphasia patients.[2] To train conversational agents for such functionality requires specialized datasets grounded in the knowledge that

[a][Online]. Available: https://tinyurl.com/KiL-MentalHealth-NLU

1089-7801 2022 IEEE
Digital Object Identifier 10.1109/MIC.2022.3182349
Date of current version 9 September 2022.

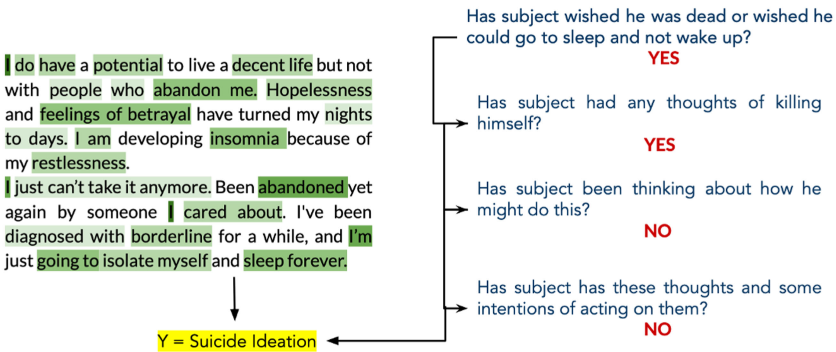

FIGURE 1. Illustration of a classification task that benefits from process knowledge. Here, an AI model using a process knowledge structure would consume the user's input, extract conceptual cues that can answer questions in process knowledge, and provide a classification label. The benefit of the process knowledge infusion in AI is the deterministic nature that it enforces in AI to achieve user-level explanations, handle uncertainty, and be safe. The figure illustrates this process in assessing suicide risk severity using a partial sequence of questions from the Columbia Suicide Severity Rating Scale (C-SSRS) (https://tinyurl.com/Posner-CSSRS). The highlighted text on the left is concept phrases that contribute to the yes/no in the C-SSRS questions. The graph on the right is process knowledge in C-SSRS.

enables AI systems to exploit the duality of data and knowledge for human-like decision making.[3,b] Furthermore, to develop agents that learn from such process knowledge-integrated datasets, we require interpretable and explainable learning mechanisms.[c] These learning mechanisms have been characterized under the umbrella of Knowledge-infused Learning (KiL).

> KiL is a class of neuro-symbolic AI techniques that utilize a variety of knowledge (lexical, linguistic, domain-specific, common-sense, process knowledge, and constraint-based) in different forms and abstractions into deep neural networks. It improves upon data-centric statistical learning to reduce training, reduce computing needs, and broaden coverage, resulting in improved performance, safety and model interpretation, and providing user-level explanations.

SHADES OF PROCESS KNOWLEDGE INFUSED LEARNING

KiL aligns with the third phase of DARPA to promote contextual adaptation in AI systems for user-level explanations. An AI system trained with knowledge infusion techniques provides forms of explanations by querying, traversing, and mapping the high-importance features to concepts in a knowledge graphs. Figure 2 illustrates the user-level explanations provided by an AI system infused with knowledge that highlights concept phrases in the input text. These concept phrases are used traverse a knowledge graph, which in Figure 2 is SNOMED-CT. Along with Figure 1 that explains process

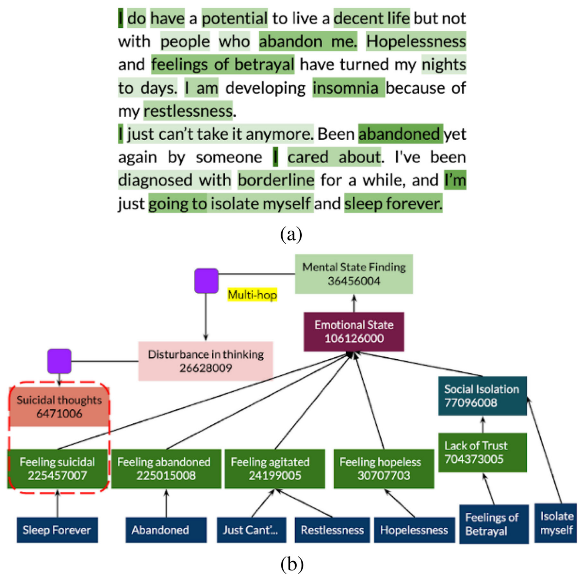

FIGURE 2. Illustration of user-level explainability using the important conceptual phrases identified by a deep learning model trained using the method in https://tinyurl.com/contextual-classify-reddit. Highlighted phrases in (a) are queried in SNOMED-CT, thus forming a contextual tree. The formation of this tree is stopped when a node is hit that has high similarity to either leaf nodes or one-hop parent nodes. The resulting tree is shown in (b). The numbers in the boxes are SNOMED-CT IDs.

[b] [Online]. Available: https://tinyurl.com/duality-data-knowledge
[c] [Online]. Available: https://tinyurl.com/petrinet-workflow

knowledge infusion contributing to reasonable path toward classification, Figure 2 provides additional user-level explanation. Within the three forms of knowledge-infusion under KiL (i.e., shallow, semideep, and deep[3]), process knowledge infusion develops a new and complementary set of methods, datasets, and evaluation methods under semideep and deep knowledge infusion.

Furthermore, any AI systems trained with such methods and over such datasets can also handle uncertainty and risk. They can establish the connection between the input and output, answering, "why such an outcome, given an input?" The AI systems are context sensitive rather than opinionated based on only the input data, i.e., a partial representation of the world. The structure and order provided by using process knowledge allows the end users a control over the AI system. Moreover, in-process knowledge in a particular domain and for a particular task (classification or generation), an AI system with a method that makes the model adaptable to process knowledge can make the system transferable across tasks. The subsequent sections will provide a concrete definition of process knowledge and its use in understanding and controlling AI models. With a focus on natural language generation (NLG), we will conceptually describe methods for infusing process knowledge into statistical AI systems. Thereafter, we provide use cases in the domain of mental health (continuing with Figures 1 and 2) and cooking.

PROCESS KNOWLEDGE AND ITS INFUSION INTO STATISTICAL AI

Process knowledge is an ordered set of information that maps to evidence-based guidelines or categories of conceptual understanding to experts in a domain. For instance, The American Academy of Family Physicians develops clinical practice guidelines (CPGs) that serve as a framework for clinical decisions and supporting best practices. CPG allows systematic assessment to optimize patient care. On the other hand, U.S. Departments of Agriculture and Health and Human Services develops Dietary Guidelines for Americans[d] that serves as a recommendation for meeting nutrient needs, promote health, and prevent disease. An AI system adapted to process knowledge can handle uncertainty in prediction, and the predicted outcomes are safe and user-level explainable. Furthermore, an AI system can consider process knowledge as metainformation to capture the sequential context necessary for carrying out a structured conversation. Also, it allows the developer of the AI system to probe the internal decision-making of AI systems using application-specific guidelines or specifications that inform the synchrony between the end-users thought process and the model's functioning.

This unique form of knowledge differs from other forms of knowledge in the following manner: 1) knowledge graph: it is structured but not ordered. Knowledge graphs can support context capture but cannot enforce conceptual flow.[e] 2) Semantic lexicons: this is a flattened form of knowledge graph that makes deep language models context sensitive and add constraints but cannot enforce conceptual flow.[4] 3) Ontologies are curated schematic forms of knowledge graphs with classes, instances, and constraints. Thus, ontologies can provide stricter control over context and constraints. If defined, an ontology can enforce order in question generation using deep language models.[5] Process knowledge is represented differently for different applications. For instance, to assess the severity of suicide risk, the process knowledge used is C-SSRS, which is similar to a flow chart. On the other hand, the GAD-7-based process knowledge is used to assess anxiety severity, which has a flattened structure (see Figure 3). DASH diet-based process knowledge can be used to assess the dietary intake of hypertension patients and also recommend meals. These characteristic properties of process knowledge and its infusion into statistical AI would yield a new class of neuro-symbolic algorithms that would drive the question.

> What if we could use the annotator's labels and the process or guidelines used to label them and explicitly control the learning of a model to recover the guideline or process (instead of implicitly).

Such an algorithm would, by design, be explainable and emulate the human model of similarity between data points. For the task of classification, a process knowledge-infused AI system would solicit the use of interpretable machine learning algorithms (e.g., decision trees, random forest) that can enforce structure in decision making over traditional deep language model-based classification.[f]

In NLG, the biggest concern with deep generative language models is that they hallucinate when either asking questions or providing responses in a conversational setting. Along with the issue of hallucination, there have been extensive study about the inappropriate and unsafe risk behaviors of language models.[g] Efforts to pair these language models with passage retrievers and rankers

[d][Online]. Available: https://tinyurl.com/american-dietary

[e][Online]. Available: https://tinyurl.com/KI-summarization
[f][Online]. Available: https://tinyurl.com/PK-iL-suicide
[g][Online]. Available: https://tinyurl.com/LaMDA-dialog

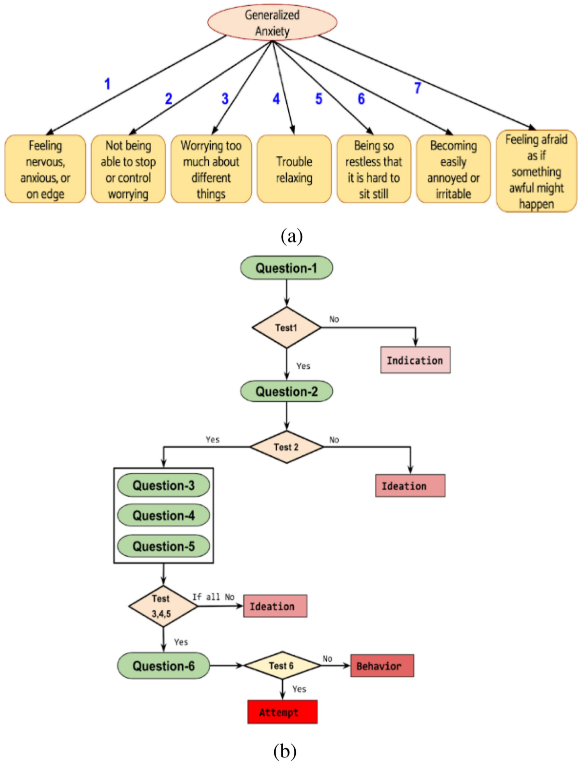

FIGURE 3. Illustration of process knowledge for different purposes. (a) GAD-7-based process knowledge with a flattened structure. (b) C-SSRS-based process knowledge with a flow chart structure.

have been proposed to control incoherent, irrelevant, and factually incorrect responses and questions; however, the order, like the one defined in process knowledge, is far from being realized.[6] Such process knowledge-based NLG is even more crucial in the field of healthcare NLP, where each response from the agent can have severe consequences. These concerns are further discussed with the help of two use cases: mental health and food domain.

MENTAL HEALTH USE CASE

AI has contributed to the domains of drug research, customized medicine, and patient care monitoring and has the potential to aid physicians in making better diagnoses. However, when AI is used in health care, various dangers and problems might arise at the individual, macro, and technological levels (e.g., awareness, education, trust), as well as at the macrolevel (e.g., legislation and rules, risk of accidents due to AI faults) (e.g., usability, performance, data privacy, and security). In the context of mental healthcare, conversational agents are prone to unsafe generations that can harm the user or engage in a conversation involving escalation in the severity of medical conditions.

Figure 4 illustrates a pipeline wherein 1) the deep statistical language model pretrained on open domain corpus when tasked to converse with a user in a mental healthcare setting generates questions that it sees online. 2) Such questions are not what a MHP would ask. If we utilize a clinical guideline, in this case, C-SSRS, the model can measure the safety of the generated question before asking. 3) Figure 4 shows a process over the detailed process knowledge that an AI agent followed to control its question generation and ask medically correct questions. A recent study from Roy et al. details this approach using C-SSRS, and Gupta et al. detail this approach using GAD-7 and PHQ-9, which are clinical guidelines to check whether the user is a patient of an anxiety disorder (GAD-7) or clinical depression (PHQ-9).[7,8]

Process Knowledge as Constraints

Some more ways in which process knowledge can be infused to add constraints and improve NLG of the current AI methods are as follows.

- Textual entailment constraints (TEC) is a directional relationship between sentences in a response or questions. If the two sentences share semantic relations and logically agree, they are entailed. If the two sentences are synonymous based on the entities they contain, they are neutral. If the second sentence refutes the information in the first sentence, they are contradictory. Such constraints are manifestations of process knowledge in clinical practice. In machine-understandable form, we can model them as Rules containing Tags and Rank (see Figure 5).
- Rules (Tag and Rank): These rules can help structure the question generation process, which is random and unsafe in current state-of-the-art NLG models.[h] For instance, if the conditional probability function within an AI model, defined as $P(Q_{k+1}|Q_k)$ is augmented with a Tag containing the following labels: {Yes/No, Degree/Frequency, Causes, Treatment/Remedies} then the model can learn to follow a definite process:
  - if Q_k is Yes then Q_{k+1} is about Degree/Frequency;
  - if Q_k is Degree/Frequency then Q_{k+1} is about Causes;
  - if Q_k is Causes then Q_{k+1} is about Treatment/Remedies; and

[h][Online]. Available: https://tinyurl.com/adaptive-education

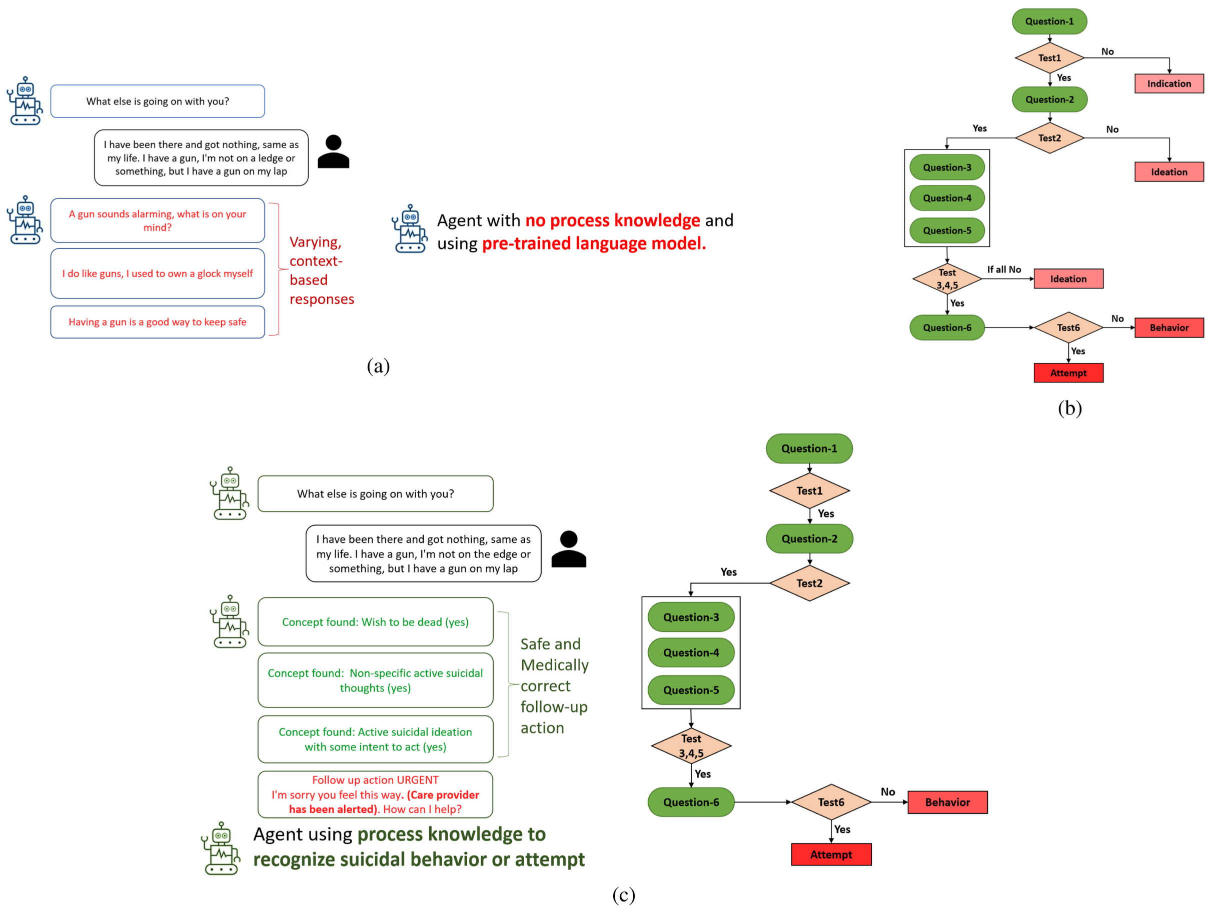

FIGURE 4. Illustration of safety in conversational artificial intelligence. It explains why process knowledge is needed to avoid unsafe conversations and make models interpretable and explainable. (a) Extreme Behavior Detection. (b) Process Knowledge to assess suicide risk severity. (c) Model utilize process knowledge for safe and medically correct follow up generation question.

– if $\hat{Q}_k$ is Treatment/Remedies then $\hat{Q}_{k+1}$ ask about Information on Other Side Effects.

Here, $\hat{Q}_{k+1}$ is the next generated question given $\hat{Q}_k$, a previous generated and accepted question. Further utility of constraints-based process knowledge infusion in AI is detailed in another application involving food recipe recommendation, wherein the constraints are defined based on allergens.

FOOD DOMAIN USE CASE

A conversational system to manage diet can help patients in various applications, such as hypertension and diabetes.[9] In most of the scenarios the interactions between user and system involve factual queries (e.g., Can you order a falafel for me? What are the sides offered with falafel? etc.). However, the challenge lies in how recommendations can be adapted to user preferences and context. It is still an open question.[10] Furthermore, how can recommendations be provided when users do not ask factual questions (e.g., Can you suggest some food that helps me control my calorie intake? I want to lose weight. What should I eat for lunch?). In case of Hypertension, patients need a nudge to switch toward healthy food habits. A nutrition management system can aid and assist them in this process. In such a scenario, when a user asks the following question to an agent: "Can you recommend dishes that are calorie efficient?" if the agent is augmented with the Internet, it would accurately respond to the following related questions (or people-also-ask questions): 1) "Are restaurants required to put calories on menus," 2) "Are calorie recommendations accurate," 3) "Should I eat less than my recommended calories?," and 4) "What food can you recommend?" Moreover, top-2 searches on google for the user query are 1) cut lots of calories and 2) how to lose weight eating more food, which are

GAD-7 Question (x)	Paraphrases (Y)	Process Knowledge (P) (Tag, Rank)
Feeling nervous, anxious, or on edge	Do you feel nervous anxious or on edge	(Yes/No,1)
	How likely are you to feel this way	(Degree/frequency,2)
	Any ideas on what may be causing this	(Causes,3)
	Have you tried any remedies to feel less nervous	(Remedies,4)
	Are you also feeling any other symptoms such as jitters or dread	(OSI, 5)
Not being able to stop or control worrying	Do you feel not able to stop or control worrying	(Yes/No,1)
	How likely are you to feel this way	(Degree/frequency,2)
	Any thoughts on what may be causing this	(Causes,3)
	Have you tried any remedies to stop worrying	(Remedies,4)
	Are you also feeling any other symptoms	(OSI, 5)

FIGURE 5. This is an example of how a process knowledge-integrated dataset is constructed in collaboration with MHPs. The leftmost column presents example questions MHPs asked. The MHPs provided Tag and Rank shown in the rightmost columns representing process knowledge. The middle column provides a series of questions gathered using Google SERP API (https://tinyurl.com/G-SERP-api) and Bing Search API (https://tinyurl.com/bing-search-api) logically ordered by MHPs.

not relevant to the user query. There are two fundamental problems here: 1) The AI system behind these recommendations is confused about whether "calorie efficiency" is positive or negative. 2) The AI system fails to bridge the gap between dishes and calorie efficiency. Furthermore, a response to such a question is dependent on the time of the day: breakfast, lunch, or dinner. A process knowledge-based conversational agent would generate the following information-seeking questions: 1) Do you have any preference in cuisine? 2) Do you want to know about low-calorie food in this cuisine for breakfast/lunch/dinner? 3) Do you want me to book reservations for restaurants that have this cuisine? 4) Do you want me to save your preferences? If the answer to 2) is no, then an alternate path in process knowledge is triggered. Here, process knowledge is the procedure for recommending and ordering food. Moreover, the agent can benefit from the 2015–2020 dietary guidelines for Americans to emphasize overall healthy eating patterns supported by five food groups: fruits, vegetables, grains, protein foods, and dairy.[i]

Similarly, for type-I Diabetes, patients need to monitor carbohydrate (CHO) intake for their insulin dosage; hence, the source of CHO determines whether a given food item is advisable. CHO count due to the fibers present in vegetables and fruits are considered healthy, whereas CHO from added sugars, white rice, and pasta are considered unhealthy.[j] Existing models advise meals based on the daily value of the CHO limit of an individual. In this case, the CHO count of a recipe derived from added sugar will be recommended by the agent if it is within the daily CHO limit of an individual. This can have severe effects on an individual's health. By infusing the process knowledge of diabetic dietary guidelines into the learning process, the agent can learn to advise appropriate meals and generate explanations to enhance interpretability and safety (see Figure 6).

Along with abovementioned two scenarios, the nutrients content of the food change based on the adverse effects of the cooking actions on the final cooked food item, such as nutrition loss or the introduction of harmful elements. To add, the dietary restrictions for each chronic disease have respective guidelines. Hence, in this scenario, two kinds of process knowledge, adverse effects of cooking actions combined with ingredients and dietary guidelines for chronic conditions, are involved in generating explanations, improving interpretability and safety of food recommendation agents.

Process Knowledge as Constraints

In addition to specific dietary guidelines for chronic conditions, cooking actions produce adverse effects. For example, the ingredients for potato fries involve potatoes, oil, salt, pepper, and other seasonings. These are advisable ingredients as per dietary guidelines for diabetes. However, the cooking action is deep frying, which produces trans-unsaturated fatty acids. The trans-unsaturated fatty acids are not advisable for any chronic diseases and the general population.[k,l]

Similarly, grilling a slice of meat can introduce carcinogenic agents[11] due to the animal fat dripping onto direct heat. However, grilling vegetables and fruits do not produce carcinogenic agents.[m] The

[i][Online]. Available: https://tinyurl.com/Dietary-Guidelines
[j][Online]. Available: https://tinyurl.com/Mayo-Diabetes-plan
[k][Online]. Available: https://tinyurl.com/trans-fats
[l][Online]. Available: https://tinyurl.com/trans-fat-cholestrol

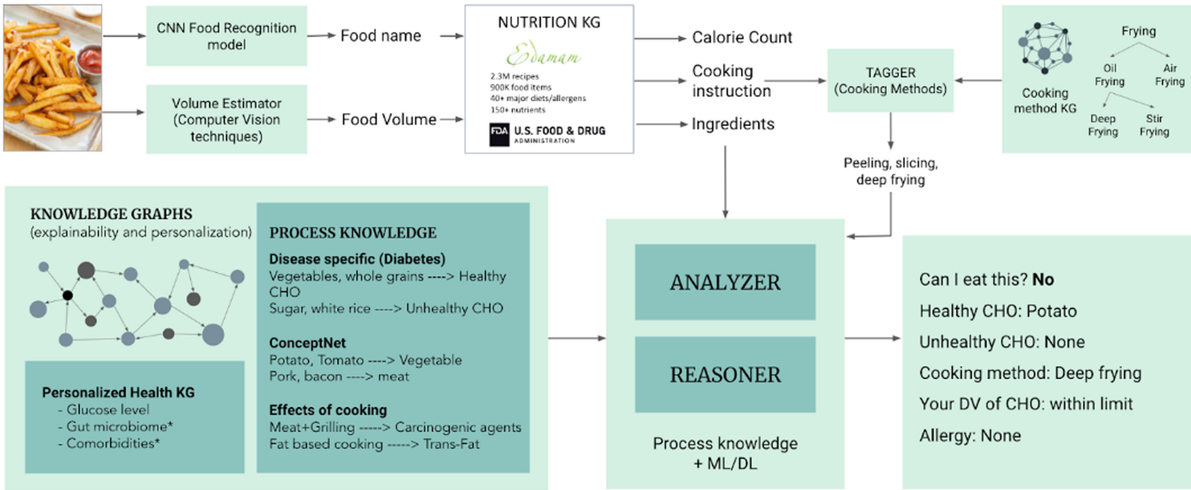

FIGURE 6. Illustration of the infusion of process knowledge in food recommendation agents.

process knowledge of cooking actions can aid the agent in learning general adverse effects due to specific combinations of cooking actions and ingredients. This process knowledge can aid the agent in learning to nudge any user toward healthy eating habits irrespective of the dietary guidelines for various chronic diseases. An agent learned by infusing the two kinds of process knowledge will be able to generate explanations, be interpretable, and thereby improve the safety aspect of meal advice.

NEED FOR NEW EVALUATION METRICS

The precision of AI is not always a good indicator of clinical effectiveness. The area under the receiver operating characteristic curve, another frequent metric, is not always the ideal indicator for clinical application. Such AI measures may be complex for physicians to comprehend or may not be clinically relevant. Furthermore, AI models have been assessed using a range of indices, including the F1 score, accuracy, and false-positive rate, which are indicators of distinct elements of AI's analytical ability. Understanding how complicated AI works necessitates a level of technical understanding not commonly seen among physicians.

AI models with process knowledge infusion require specialized metrics for evaluating the performance concerning safety and uncertainty, and risk handling. For instance, stanford natural language inference, multigenre natural language inference, and others similar datasets can be used to create a learned evaluation metric to assess safety in generation by comparing the generated hypothesis with a premise.[n] In essence, safety and uncertainty and risk handling would require human evaluation, which is a mandate; these metrics are also equally important as they either involve: 1) annotators' agreements/disagreements; 2) knowledge source; and 3) train deep language models on datasets that have data samples ordered by some relationships.[12,13]

a) Average number of unsafe matches: This represents the average number of matches across all model-generated questions against a set consisting of utterances, lexical content, or ontology concepts used to describe harmful communication. Such a measure provides a range of means to impose safety checks that can be extracted from unstructured, semistructured, and structured sources and domain experts. For example, named entities in the generated content could match against harmful concepts in a knowledge base or in a lexicon set containing harmful phrases (unigrams, bigrams, and trigrams).
b) Perceived risk measure: This is an annotator-in-the-loop metric to judge the model's stability in light of agreement and disagreement between the annotators, a notion of uncertainty and safety. It is composed of two components: 1) Penalty: A ratio of the count of misclassified samples over the count of those samples where the annotators

[m][Online]. Available: https://tinyurl.com/healthy-grill

[n][Online]. Available: https://tinyurl.com/NLI-datasets

disagree with each other. 2) Benefit: A ratio of the count of samples where the model's predicted label agrees with some annotators (ignoring the disagreement between them) over the total number of annotators. Such a metric is efficient for controlling unsafe predictions as opposed to using statistical loss functions that quantify uncertainty in predictions and overwhelm the experts in the loop with reannotations.[14]

c) Semantic relations and logical agreement measures: These are trained metrics constructed using the RoBERTa model, a deep language model trained independently on sentence similarity and natural language inference GLUE tasks. These metrics have been introduced in a recent study by Gaur et al. that unites meta-information-guided passage retrievers and TEC for inducing logical ordering in the generations and preventing retrieval-augmented language models from hallucinations.[15] Semantic relation is a metric that counts the number of generations semantically similar to a user query over the total number of generations. The logical agreement score records the count when the current generated question entails a previously generated question. The score takes the sum of such counts and divides them by the number of generations.

SUMMARY AND FUTURE DIRECTIONS

Real-world interactions between the users are not a single shot activity but rather a chain of exchanges involving procedural questions and responses; at a macroscopic level and the microscopic level, it comprises entities and actions that keep changing during a task-oriented conversation. This phenomenon can be well understood and controlled through a process of knowledge that represents a human's mental model of conversation. In this article, through example use cases in mental healthcare and food, we explained the notion of process knowledge that naturally concerns consistency, explainability, and interpretability in AI's decision-making process. To the best of our knowledge, this article projects its role in pushing statistical AI to be safe, less uncertain, and risky in its classification and NLG tasks. With process knowledge, the AI model can support reasoning, which is essential to develop trust in stakeholders using the application in various downstream tasks. We showed various existing process knowledge and the methods that data-driven AI models can use. As a future direction, we envision the utility of process knowledge in personalization, which is essential in developing interventional plans for patients with other mental health disorders (e.g., autism, aphasia) and developing food plans for patients with specific dietary needs.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation under Grant 2133842, "EAGER: Advancing Neuro-symbolic AI with Deep Knowledge-infused Learning."

REFERENCES

1. A. Sheth, M. Gaur, K. Roy, and K. Faldu, "Knowledge-intensive language understanding for explainable AI," IEEE Internet Comput., vol. 25, no. 5, pp. 19–24, Sep./Oct. 2021.
2. R. D. Newman-Norlund et al., "The aging brain cohort (abc) repository: The university of south Carolina's multimodal lifespan database for studying the relationship between the brain, cognition, genetics and behavior in healthy aging," Neuroimage: Rep., vol. 1, no. 1, 2021, Art. no. 100008.
3. A. Sheth, M. Gaur, U. Kursuncu, and R. Wickramarachchi, "Shades of knowledge-infused learning for enhancing deep learning," IEEE Internet Comput., vol. 23, no. 6, pp. 54–63, 2019.
4. G. Libben, "From lexicon to flexicon: The principles of morphological transcendence and lexical superstates in the characterization of words in the mind," Front. Artif. Intell., vol. 4, 2021, Art. no. 788430.
5. K. Stasaski and M. A. Hearst, "Multiple choice question generation utilizing an ontology," in Proc. 12th Workshop Innov. NLP Building Educ. Appl., 2017, pp. 303–312.
6. M. Glass, G. Rossiello, M. F. M. Chowdhury, A. R. Naik, P. Cai, and A. Gliozzo, "Re2G: Retrieve, Rerank, Generate," in Proc. UTC-5 Conf. - NAACL, 2022.
7. K. Roy, M. Gaur, Q. Zhang, and A. Sheth, "Process knowledge-infused learning for suicidality assessment on social media," 2022, arXiv:2204.12560.
8. S. Gupta et al., "Learning to automate follow-up question generation using process knowledge for depression triage on reddit posts," 2022, arXiv:2205.13884.
9. S. Yagcioglu, A. Erdem, E. Erdem, and N. Ikizler-Cinbis, "Recipeqa: A challenge dataset for multimodal comprehension of cooking recipes," 2018, arXiv:1809.00812.

10. F. Pecune, L. Callebert, and S. Marsella, “Designing persuasive food conversational recommender systems with nudging and socially-aware conversational strategies,” Front. Robot. AI, vol. 8, 2022, Art. no. 733835, doi: 10.3389/frobt.
11. A. J. Cross and R. Sinha, “Meat-related mutagens/carcinogens in the etiology of colorectal cancer,” Environ. Mol. Mutagenesis, vol. 44, no. 1, pp. 44–55, 2004.
12. A. Williams, N. Nangia, and S. R. Bowman, “A broad-coverage challenge corpus for sentence understanding through inference,” 2017, arXiv:1704.05426.
13. O.-M. Camburu, T. Rocktäschel, T. Lukasiewicz, and P. Blunsom, “e-snli: Natural language inference with natural language explanations,” in Proc. Adv. Neural Inf. Process. Syst., vol. 31, 2018.
14. R. Sawhney, A. T. Neerkaje, and M. Gaur, “A risk-averse mechanism for suicidality assessment on social media,” in Proc. Assoc. Comput. Linguistics, 2022, pp. 628–635.
15. M. Gaur, K. Gunaratna, V. Srinivasan, and H. Jin, “ISEEQ: Information seeking question generation using dynamic meta-information retrieval and knowledge graphs,” 2021, arXiv:2112.07622.

AMIT SHETH is the founding director of the AI Institute of South Carolina (AIISC: aiisc.ai), Columbia, SC, 29208, USA, and the NCR chair and professor of computer science and engineering. Contact him at amit@sc.edu.

MANAS GAUR recently completed Ph.D. degree and is joining the University of Maryland, Baltimore County, MD, 21250, USA, as an assistant professor with the Department of Computer Science. Contact him at mgaur@email.sc.edu.

KAUSHIK ROY is a Ph.D. student advised by Prof. Sheth at AI Institute, University of South Carolina, Columbia, SC, 29208, USA. Contact him at kaushikr@email.sc.edu.

REVATHY VENKATARAMAN is a Ph.D. student advised by Prof. Sheth at AI Institute, University of South Carolina, Columbia, SC, 29208, USA. Contact her at revathy@email.sc.edu.

VEDANT KHANDELWAL is a Ph.D. student advised by Prof. Sheth at AI Institute, University of South Carolina, Columbia, SC, 29208, USA. Contact him at vedant@email.sc.edu.