SAGE: System for Accessible Guided Exploration of Health Information

Sabriya Maryam Alam, Haodi Zou, Reya Vir, and Niloufar Salehi

University of California, Berkeley
sabriya.alam@berkeley.edu, haodi.zou@berkeley.edu,
reyavir@berkeley.edu, nsalehi@berkeley.edu

Abstract

The Center for Disease Control estimates that 6 in 10 adults in the United States live with a chronic disease such as cancer, heart disease, or diabetes. Yet most patients lack sufficient access to comprehensible information and guidance for effective self-management of chronic conditions and remain unaware of gaps in their knowledge. To address this challenge, we introduce SAGE, a System for Accessible Guided Exploration of healthcare information. SAGE is an information system that leverages Large Language Models (LLMs) to help patients identify and fill gaps in their understanding through automated organization of healthcare information, generation of guiding questions, and retrieval of reliable and accurate answers to patient queries. While LLMs may be a powerful intervention for these tasks, they pose risks and lack reliability in such high-stakes settings. One approach to address these limitations is to augment LLMs with Knowledge Graphs (KGs) containing well-structured and pre-verified health information. Thus, SAGE demonstrates how LLMs and KGs can complement each other to aid in the construction and retrieval of structured knowledge. By integrating the flexibility and natural language capabilities of LLMs with the reliability of KGs, SAGE seeks to create a collaborative system that promotes knowledge discovery for informed decision-making and effective self-management of chronic conditions.

Introduction

Patients face barriers to effective information-seeking when finding, assessing, and understanding information that is accurate and relevant to their condition. This is crucial for those with chronic conditions, who must frequently make self-management decisions about their healthcare. Tailoring information to a patient's needs is a significant challenge, demanding considerable effort by healthcare providers. Patients with limited health literacy find it particularly challenging to navigate healthcare information and tend to ask significantly fewer questions than those with adequate health literacy, exacerbating health disparities (Menendez et al. 2017). In fact, a survey revealed that two-thirds of patients forget or omit questions during doctor appointments. Post-visit, patients may formulate new queries, but lack accessible means to find accurate and understandable information to address their concerns (Wolters-Kluwer 2023).

Addressing the pressing need for accessible interventions in patient education involves addressing 3 main challenges. Firstly, patients often overestimate their healthcare knowledge and fail to identify their knowledge gaps (Canady and Larzo 2023). Secondly, although asking questions enhances health knowledge, patients struggle to formulate questions for their healthcare provider (Katz et al. 2007). Lastly, even when provided access to healthcare information, patients find it challenging to comprehend medical complexity (Cai et al. 2023; Reifegerste and Hartleib 2016).

We introduce our System for Accessible Guided Exploration (SAGE), developed to address these challenges. Leveraging Large Language Models (LLMs), *SAGE is an information system that helps patients identify and fill gaps in their knowledge* by organizing health information for accessible navigation, recommending insightful questions for personal information needs, and delivering clear, accurate answers based on a patient's health literacy level. While LLMs may excel in such tasks, they have limitations that may cause them to be unreliable in high-stakes settings such as healthcare, including generating inaccurate or inconsistent information (Cascella et al. 2023). However, LLM performance can be enhanced by interfacing with other reliable knowledge models that can produce verifiable and predictable outputs, such as Knowledge Graphs (KGs) (Agrawal et al. 2023). SAGE aims to merge the benefits of reliability and verifiability of KGs with the flexibility offered by LLMs to customize a patient's information-seeking experience.

In practice, SAGE takes in complex healthcare information, such as discharge information from a physician, and leverages LLM capabilities to create a structured KG representing the information critical for patient understanding. This organized KG is augmented with LLM-generated suggested questions that can guide a patient to a deeper understanding of their healthcare condition and self-management practices. Finally, when a patient interfaces with SAGE, they traverse the KG and gain knowledge from accurate and pre-verified health information. This information is then presented in a manner that is tailored to their comprehension levels and health literacy, ensuring a personalized and accessible experience. Ultimately, SAGE strives to establish a versatile pedagogical framework that may be applied in diverse settings to facilitate clear and comprehensive communication and exploration of complex information.

Related Work

Health Information-Seeking Behaviors

Health information-seeking behavior describes how patients seek and manage information about their health. In the past several decades, patients have gained more access to diverse forms of healthcare information, from authoritative sources (e.g. WebMD, Mayo Clinic, NIH websites) as well as anecdotal sources (e.g. Google searches, Reddit posts, Quora forums, etc.). This has prompted a shift toward more patient-centered care. Unlike traditional models that emphasize patient compliance, modern models envision patients as active partners in managing their conditions (van der Eijk et al. 2013). This places greater responsibility on patients to take initiative to learn and act on their own, often without direct supervision from healthcare providers.

This model can empower patients and foster a sense of control and responsibility in managing their conditions. Benefits linked to acquiring healthcare information independently include increased control over one's illness and life, reduced pain, quicker recovery, enhanced participation in healthcare decisions, higher satisfaction with consultations, improved mental health, and better coping skills (St. Jean 2012). To make informed decisions leading to these benefits, patients require credible and accessible information within their comprehension abilities.

The practice of self-management is especially crucial when managing chronic health conditions that require sustained, long-term medical support over the course of months, years, or even a lifetime. Chronic care aims to manage symptoms, prevent complications, and improve quality of life through ongoing monitoring, treatment adjustments, and lifestyle modifications. As a result, these patients must continually seek information and make medical decisions on a day-to-day basis that directly impact their future health complications or recovery process (Longo et al. 2010). SAGE focuses on addressing the needs of this patient population.

Self-Management of Chronic Conditions - An Illustrative Example of Diabetic Foot Ulcers

Diabetes is one of the most highly prevalent chronic conditions affecting patients across the globe. As of 2021, 537 million adults worldwide live with this condition, and in 2019, diabetes was the direct cause of 1.5 million deaths (IDF 2021). As a condition that requires continual, high-quality clinical care, diabetes complications can encompass neuropathy, vision impairment, kidney disease, congestive heart failure, stroke, diabetic coma, and lower limb amputation.

Patients with diabetes have significant responsibility to self-manage their condition through routine monitoring of blood sugar, lifestyle modifications in dietary choices and daily habits, and mindfulness of developing symptoms. Taking appropriate measures early on can prevent the development of more serious complications. For example, diabetic patients must carefully monitor their feet for seemingly inconspicuous wounds that may develop into diabetic foot ulcers (Chin et al. 2019). This particular complication can be very serious. The mortality rate associated with diabetic foot ulcers ranges from 43% to 55% and up to 74% for patients with amputation. These numbers are dramatically higher than the mortality rate for many types of cancer (such as prostate, breast, and Hodgkin cancers) (Robbins et al. 2008).

Diabetic foot ulcers often arise from a combination of neuropathy, poor circulation, and immune system impairments. If not monitored and treated properly, they can become infected and rapidly progress to deep-seated tissue damage, and in severe cases, a diabetic foot ulcer may necessitate limb amputation . In fact, poorly managed diabetic foot ulcers are one of the most leading causes of lower extremity amputation in the United States. As a life-altering complication, 82% of all vascular-related lower extremity amputations in the United States coincide with diabetes (Molina and Faulk 2022). Well-informed patients are equipped to comprehend the nuanced interplay between diabetes and foot health, grasping the imperative need for meticulous foot care, regular monitoring, and prompt intervention at the earliest signs of complications (Chin et al. 2019). Patient education can serve as a proactive measure against these consequences.

Given the high stakes, understanding factors that drive or hinder patients' diabetes-related information-seeking is crucial. Information-seeking behavior, encompassing how individuals interact with and use information, is particularly impactful in chronic conditions like diabetes. For this impact to materialize, patients need easy access to personally relevant, understandable, reliable, timely, and actionable information about disease management (St. Jean 2017).

Unfortunately, studies reveal inadequate self-management behaviors among diabetic foot ulcer patients. Timely treatment is crucial for a positive prognosis, involving appropriate wound care, glycemic control, and lifestyle changes. However, many patients do not monitor their blood glucose, are not provided with proper foot care instructions, and do not receive clear guidance about how to identify complications. This underscores the broader issue: patients are ill-equipped to navigate complex health information effectively, leading to disengagement from healthful behaviors (Chin et al. 2019; Costa, Tregunno, and Camargo-Plazas 2022).

Barriers to Self-Management and Health Information-Seeking

Inadequate self-management is a preventable issue frequently stemming from knowledge gaps, which may be addressed through more effective information-seeking behavior. Information-seeking behavior encompasses the extent to which people are motivated to seek knowledge, the credibility and authority of the information sought, and how the information is used and found to be beneficial or harmful to broader health outcomes. This process can be significantly impeded by internal and external barriers that must be overcome to adequately address chronic health conditions (Costa, Tregunno, and Camargo-Plazas 2020).

External Barriers Patients typically obtain medical information from their healthcare provider, or from their own searches on the internet. Unfortunately, there are drawbacks to both in terms of convenience and reliability to each of these approaches (Wei, Du, and Zhang 2016). While talking to a medical professional may offer reliable information, many resource-constrained patients may not have the means to visit a doctor frequently or attend credible informational classes, due to geographic, logistical, or financial barriers.

In contrast, information on the internet is highly convenient and accessible for most patients, but the results of web searches can prove to be unreliable. Patients may struggle to discern what is accurate, trustworthy, relevant, and up-to-date information about their condition, and those with low digital and health literacy can face additional barriers to care as well. Even among patients who are adept at navigating the internet, symptoms of disease such as neuropathy and cognitive difficulties can make it challenging to actively seek and evaluate information when it is not centralized (Wei, Du, and Zhang 2016).

As a response to these challenges, telehealth services and electronic health communication have become increasingly prevalent in recent years as a means of delivering more consistent and convenient access to healthcare information. However, providing these services requires physicians to respond to high volumes of messages with quick turnaround, which is a significant burden on their already demanding workloads. Because of this, some major hospital systems across the country have implemented fees for electronic communication between patients and physicians. Patients are being charged for queries that require more than 5 minutes of a physician's time, and the cost can range from as little as \$3 to a charge of \$35 to \$100 (Holmgren et al. 2023). Such fees for online communication may also deter patients from asking follow-up questions, and may exacerbate existing health disparities. Thus, it is critical to develop interventions that alleviate the burden of answering questions from physicians themselves, while also allowing patients to have accessible information to manage their health.

Internal Barriers Patients may also face internal barriers that impede their search for reliable and comprehensible medical information. There are physical, social, affective, and cognitive factors that can motivate, demotivate, or impede a patient's search for health information in light of a chronic condition such as diabetes. An initial diagnosis can be a strong motivating factor for some patients to learn more about their condition, while for other patients a diagnosis can be demotivating. Moreover, in medical conditions, apprehension is also a strong motivating factor. Patients fear for their future when navigating healthcare information, often discovering information about grave complications or mortality rates. They may fear that they will not be able to return to a sense of normalcy after their diagnosis as well (St. Jean 2017). In response, they may be seek as much information as possible to be able to actively take steps toward a return to normalcy, or be in denial and continue their lifestyle unchanged until further complications arise. The latter case is referred to as "information avoidance", and particularly for time-sensitive medical issues, information avoidance or uncertainty about how to navigate new, complex information can have dire consequences for patients (Guo, Si, and Sun 2023). Thus a patient's emotional state (including feelings of dread, fear, uncertainty, anxiety, or hope) can have a significant impact on their ability or motivation to seek and comprehend healthcare information.

This is particularly critical as the usefulness of medical information changes over time. For some conditions, being aware of symptoms proactively is more useful than learning about them after they occur, as by then it may be too late to take certain measures. This mismatch between when information is most needed and when it is actually provided, obtained, recognized as relevant, and acted upon can contribute to the development of serious medical complications and adds another layer of complexity to medical information seeking (St. Jean 2017). Moreover, patients may face anomalous states of knowledge, wherein they may be unaware that there is something they do not know, which they in fact need to know. Guiding patients through this process is key to helping them navigate their healthcare information most effectively.

Pedagogical Frameworks for Information Navigation

Understanding how people organize and navigate new information is crucial for assessing health information-seeking behaviors. Domain modeling is one strategy used to organize information by creating a structured conceptual representation of a subject area by defining concepts and relationships within it (Saini et al. 2022). Various theories also offer insight about how people explore and navigate information. Information foraging theory is an idea rooted in behavioral ecology, which draws an analogy between how animals forage for food and how people search for information. It suggests that individuals strategically allocate their efforts to maximize knowledge gains. This highlights the adaptive nature of information-seeking behavior, suggesting that users will adapt to the information landscape made available to them and seek paths that promise the most relevant and valuable information (Pirolli and Card 1999). Schema theory, originally proposed in 1976 by psychologist Jean Piaget, posits that people organize information into mental structures, which helps them comprehend new information by relating it to existing knowledge (Neumann and Kopcha 2018). SAGE draws from these theories to create organized information structures that align with people's natural information foraging instincts.

Self-Assessment of Knowledge and Understanding

Effective health information-seeking demands self-awareness regarding one's knowledge and limitations. This proves challenging as individuals often overestimate their capabilities due to cognitive biases, known as the Dunning-Kruger effect (Canady and Larzo 2023). Those with lower cognitive abilities struggle to accurately assess their own understanding and performance, leading to overestimation of their own capabilities and contentment with limited knowledge (Dunning 2011). False confidence means that patients do not feel perplexed or disoriented, and are less likely to seek out further information. Similar studies have also shown that 90% of respondents claim to have some knowledge of fully made-up words, when those words are interspersed with familiar words (Atir 2018). Similarly in medical contexts, patients may falsely believe they grasp complex terminology. This can be particularly detrimental for patients with low health literacy, who may lack the metacognition to assess their own knowledge and understand its impact on their self-management.

A remedy for inaccurate self-assessment lies in receiving constructive feedback, where individuals are gently made aware of their weaknesses and blind spots. Asking questions is an intuitive way to gain information, but most patients are reluctant to ask questions due to uncertainty about appropriate questions to ask, self-consciousness limited knowledge, or deference to healthcare providers (Katz et al. 2007).

Several studies document patients' reluctance to ask healthcare providers questions. Particularly, patients with limited health literacy ask significantly fewer questions than patients with adequate health literacy, especially regarding their treatment regimen. These are the patients who would arguably benefit most from detailed health education. Studies have also shown that low-literacy patients are less likely to use medical terminology, refer to medications by name, request additional services, or seek new information (Katz et al. 2007). Another survey found that two-thirds of patients leave a physician's office having forgotten to ask questions they intended to or without asking questions that came up during the visit (Wolters-Kluwer 2023).

There are limited strategies that physicians can implement to try to combat this and elicit questions from patients effectively. For example, physicians may ask open-ended questions like, "What questions do you have?" as opposed to a closed-ended inquiry such as, "Do you have any questions?" (Coleman, Salcido-Torres, and Cantone 2022). Even so, generating pertinent questions from scratch can still prove to be a challenging task for many patients. Alternative strategies include referring patients to Frequently Asked Questions (FAQ) documents or providing Question Prompt Lists that can inspire patients to formulate their own questions (Sansoni, Grootemaat, and Duncan 2015). Unfortunately, these pre-established questions are often too general and not tailored to individual patient knowledge, information needs, or literacy levels. They may include irrelevant or omit vital questions for the patient. Equipping patients to articulate thoughtful questions empowers them to make informed decisions regarding their self-management, and SAGE aims to address this by providing patients with questions customized to their information needs.

System Description

The findings from a longitudinal study on factors that motivate, demotivate, or impede health-related information seeking for diabetic patients suggest that the provision of health information should be well-structured (to minimize gaps in information), be tailored and interactive (to be personalized to be highly relevant and actionable for a particular patient),

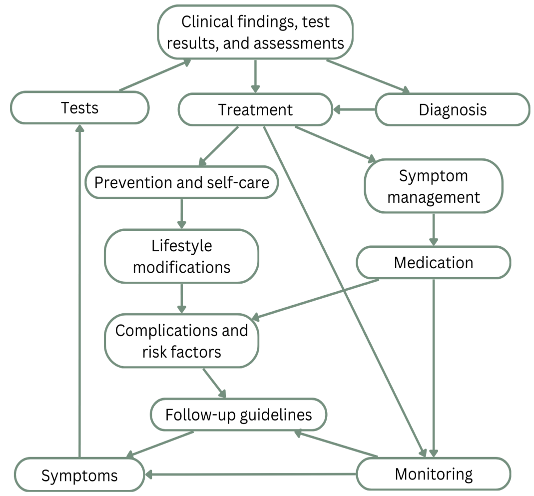

Figure 1: LLM-generated extracted knowledge graph based on key information in physician's note to patient

be digestible to avoid information overload, and be ongoing such that it meets the patient where they are and addresses the needs most pertinent to their current condition (St. Jean 2017). This lends itself naturally to several goals for a health information system, which SAGE aims to address: health information must be structured and organized in a format that is accessible and easily traversable, must address the personalized information-needs of a patient, and must provide reliable, comprehensible, and actionable information.

Our work posits that LLMS can contribute to these goals, enhancing patients' information-seeking behaviors by applying pedagogical theories to complex health information. To operationalize this, we leverage the capabilities of LLMs to create structured representations of complex medical documents which align the system with users' natural information foraging instincts as they explore information and make connections to their existing knowledge. *SAGE is an information system that helps patients identify and fill gaps in their knowledge.* To this end, a key component is to suggest thoughtful questions to patients to help them recognize their unrecognized information needs, empowering them to learn more and make well-informed decisions about their self-management. SAGE demonstrates these capabilities by addressing the information needs of diabetic patients at high risk for diabetic foot ulcer complications.

Scenario Consider a case in which a patient is diagnosed with a diabetic foot ulcer, and they receive a message from their physician confirming the results from their last appointment. This note introduces many important concepts to the patient, including findings from their test results, their diagnosis, instructions for monitoring, and follow-up care. It can be difficult for a patient, especially one from a lower health literacy background, to parse this information and understand what steps to take to seek clarification.

SAGE creates a personalized KG based on key ideas in the doctor's note and augments it with follow-up questions that can prompt a patient to deepen their understanding. To

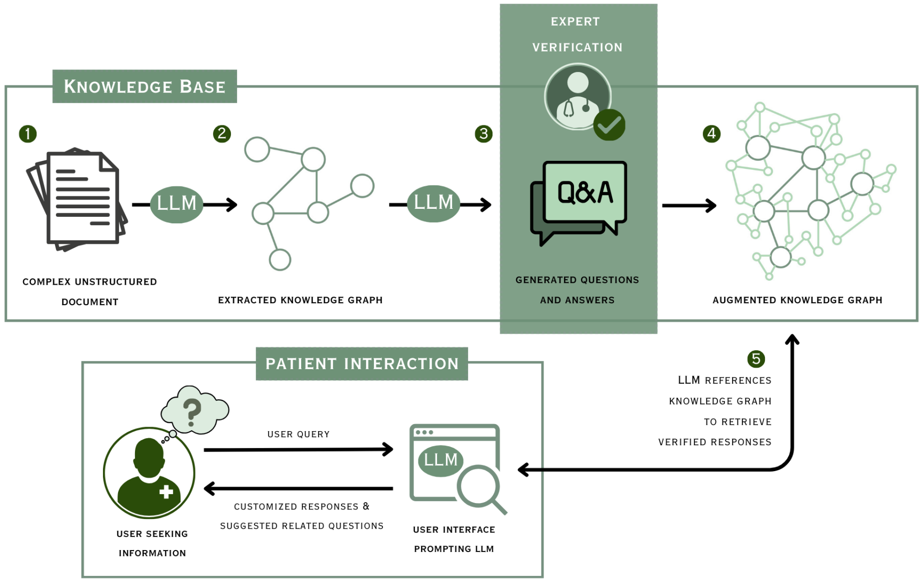

Figure 2: Overview of SAGE's Current Functionality

begin their exploration, the patient poses an initial query regarding the healthcare information they received. An LLM compares this input to the pre-generated and pre-verified questions in the question bank and retrieves the most similar question. The patient may choose to view the automated response to this suggested query, which has been pre-verified by healthcare professionals as reliable, or they may choose to send their original query directly to their physician.

After viewing the response, the patient can offer feedback regarding their understanding of the answer provided. If the answer was unclear to them, they are suggested another highly related question in the same information category to help them deepen their understanding. If they express that their understanding was clear, SAGE nudges them toward new unexplored areas of information. To help signpost patients through their information traversal, each suggested question is accompanied by a short explanation describing how it relates to previous questions asked by the patient, and why it is relevant to their healthcare needs. Related questions and answers are maintained in an ongoing log that the patient can refer to as needed. Through these mechanisms, patients choose their own path of traversing through the KG and gain information at a rate that is comfortable for them, with information provided that is most relevant to their needs. The patient may also request that the response text be modified for their comfort. Some patients may opt for simplified language to overcome language or health-literacy barriers while others prefer higher reading levels. SAGE is adaptable to these different patient preferences.

Connecting Information Aids in Navigation

Whereas a patient may struggle to create meaning out of a complex document such as a doctor's discharge instructions, a powerful LLM can thoroughly analyze such text and extract valuable information. An LLM can identify action items, provide concise summaries of content, and make medical terminology more accessible to patients. It can also assess the text's reading level, identify its central concepts, and categorize crucial information necessary for comprehensive understanding. SAGE aims to capture and represent this knowledge for easy retrieval and interpretation. According to information foraging and schema theory, this entails explicitly connecting themes and ideas in the original text, aiding patients in effective information navigation. One approach to achieve this is through the construction of a KG, which organizes information in a structure where nodes represent knowledge entities and edges between them represent the relationships between these entities. Each node and edge may also be associated with other properties, attributes, and metadata. This stores information in a format that is easily understandable by both humans and machines, and proving to be a useful tool for information organization and navigation (Rajabi and Kafaie 2022).

Generating a KG typically demands substantial human effort and domain expertise, limiting scalability and flexibility across different applications (Carta et al. 2023; Meyer et al. 2023). However, LLMs can swiftly parse and organize information, and are capable of generating KGs when prompted with a specific structure, as showcased in our system's functionality. As shown in Figure 2 (steps 1 and 2), given a physician's note diagnosing a patient with a diabetic foot ulcer, SAGE not only identifies and categorizes key information elements, but also constructs relationships between them and explains their connections. An LLM-generated graph illustrating connections between information categories found in this doctor's note is shown in Figure 1.

Traditional KGs tend to be fairly rigid. Once populated with information, they are primarily used as a tool to retrieve and traverse information. However, in our application, the constructed KG is flexible and can be personalized to the specific information needs of a patient. Each node that contains a main idea of information can be augmented with personalized questions which can introduce new knowledge to the graph from outside the bounds of the original text.

Question Prompting Encourages Exploration

SAGE leverages an LLM to suggest new questions to patients to further explore and comprehend their healthcare information. Each category from the original document, connected through the generated KG, prompts 5 follow-up questions generated by the LLM, as shown in Figure 2 (steps 3 and 4). In the case we examined, this generated a bank of 60 patient-specific questions. To aid in navigation of healthcare information, these questions are organized in a web of related queries. This structure not only highlights intra-category connections between questions, but also introduces highly relevant inter-category questions which can expand patients' exploration beyond what their current understanding. An ongoing log of previously asked questions is also maintained, tracking patients' information-seeking behavior and traversal patterns based on the KG.

To elaborate, we organize and represent the relationships between all the questions and answers using a KG. We initialize the graph using each question and answer, setting relevant attributes including category of question. We encode each question and answer into its vector representation, to allow comparisons between them, and initialize edges to form connections between the nodes. An outgoing edge from the question to its answer carries a weight of 1, so the answer is always linked to the question. In each category, vector similarities are calculated between each answer and possible follow-up questions. An edge is added if the similarity is between 0.25 and 0.75 to avoid suggesting unrelated or redundant questions. A single outgoing edge is also added from each answer to the most similar question outside of its category. Based on this, patients are then provided with a generated question sequence by traversing the graph. Traversal starts from an initial question and follows question-answer edges by the highest semantic similarity, omitting previously visited nodes.

Reliable and Accessible Responses Ensure Patient Safety in Automated Q&A

Offering suggested questions to patients necessitates providing automated answers. LLMs have recently been used as conversational agents, in which users prompt the system with questions and receive responses. Automated question-answering could prove beneficial for patients seeking information about their healthcare condition, but accuracy in these high-stakes settings is paramount (Lee, Bubeck, and Petro 2023). A critical limitation of LLMs affecting their real-world applicability is their tendency to produce hallucinations, based on knowledge gaps in the model (Agrawal et al. 2023). Hallucinations may include incorrect content that deviates from the intended input and expected output, contradictions of prior outputs, or information that is inconsistent with factual knowledge. This can jeopardize reliability of LLMs in practical downstream tasks (Wen, Wang, and Sun 2023).

To address the issue of hallucinations, one emerging research direction is to enhance LLMs through the augmentation of external knowledge, using knowledge representation tools like KGs (Yang et al. 2023; Zhu et al. 2023). KGs have been leveraged to provide a constrained bank of verified or domain-specific information that an LLM can draw from (Sun et al. 2023). Whereas most LLMs lack transparency, graph-augmented LLMs have shown promise in providing explainable outputs, making them more user-friendly. Thus in our application, the LLM directly interfaces with a graph of pre-generated and pre-verified questions and answers, ensuring that generated questions are defined within constraints appropriate for safe and secure usage. Patients can access physician-endorsed information, without requiring one-on-one messaging with a physician.

Shown in Figure 2 (step 5), this process leverages the benefits of a KG, with constrained and verified information that prevents the LLM from hallucinating, while also enhancing the output with the capabilities of an LLM. After parsing the KG and retrieving the most relevant information, the LLM offers flexibility, customizing and refining the output to make it more applicable and understandable to the patient. For example, the output may be adapted to the patient's health literacy level. The national recommendation is that all patient education material be provided at a middle school reading level, however in practice this can be challenging for physicians to gauge and adapt to (Eltorai et al. 2014). An LLM can easily rephrase, define, and reword text to fit into an a patient's preferred reading level, whether at a middle school reading level or more advanced. Thus, SAGE employs a retrieve-rewrite-answer approach to provide accurate and comprehensible responses for patients, so they may navigate their healthcare information with confidence.

Discussion

Human-AI Complementarity to Enhance Information-Seeking Behaviors

Using Artificial Intelligence (AI) systems collaboratively to enhance human knowledge-seeking and understanding is an area of ongoing research. Human-AI partnership relies on the premise that the expertise of an AI system can complement the knowledge of a human, allowing the partnership to accomplish more than either entity could on their own (Holstein et al. 2023; Ren, Deng, and Joshi 2023). Studies have shown people are more likely to trust and rely on an AI system that demonstrates complementary expertise to their own knowledge, which is why an application of this nature which helps fill knowledge gaps may prove especially useful (Zhang, Lee, and Carter 2022).

SAGE is exploring how LLMs could be powerful in helping to generate pedagogical tools for people navigating complex health information. These benefits may extend to other domains in which users must comprehend unfamiliar information. For example, helping students generate questions for their teachers during or after a lesson (Lee et al. 2022), encouraging children to ask more questions (Abdelghani et al. 2023), or generating educational questions to encourage and assess learning (Elkins et al. 2023; Kulshreshtha et al. 2022; Xu et al. 2022). Such a tool may also be applied in settings where people need to parse complex legal or technical documents. Within a medical setting, it can also be used to help a physician decide what questions are most important to ask their patients. On a broader scale, a similar system could also help find gaps in literature or unanswered questions in a series of documents. These are all high-impact cases in which complementary expertise can enhance and enrich a person's comprehension and knowledge as they seek information.

In our application, we also demonstrate how various knowledge models can be used to support, reinforce, and verify each other in practical applications. For example, LLMs and KGs can complement each other to aid in the construction and retrieval of structured knowledge. An LLM can be used for KG construction, and KGs can also be used to verify and justify the outputs of LLMs, especially when used in domain-specific applications where knowledge must be constrained and verified. While KGs themselves can be used in automated question-answering, using them in conjunction with LLMs allows for more customized presentation and utilization of information based on the specific information needs of the user.

Future Applications and Assessments of SAGE

SAGE is preparing to be piloted with patients at risk for or recently diagnosed with diabetic foot ulcers, and our forthcoming study will assess the impact of generated questions on patients' information-seeking behaviors and their engagement with health information. In our evaluation, we will elicit feedback from participants regarding how easily they are able to generate questions for their healthcare provider (efficiency), their overall understanding of their health condition and treatment (completeness), whether their information-seeking behavior led to adequate coverage of the KG (breadth of knowledge), and how well they comprehended subcategories of information (depth of information). Other metrics may assess the quality of the questions, such as the quantity or proportion of questions that patients say would not have occurred to them without suggestion (novelty) or if there is a low rate of suggesting known or irrelevant questions (utility). Additional assessments will gauge reliance on the system, trust in automated responses, and changes in patients' understanding of their condition.

Another system objective requiring assessment is SAGE's effectiveness in guiding patients toward unexplored topics in health and self-management. The goal is to present patients with relevant questions they may not have generated independently but recognize as valuable. To this end, examining the questions patients do generate independently can offer insights into their comprehension of their condition and highlight what is most pressing to them. Intuitively, patients are likely to ask questions about self-identified knowledge gaps where they possess some understanding, but can readily recognize that they need more information. They are unlikely to ask questions about concepts they understand well, so suggesting such questions would not be useful. However, they are also unlikely to ask questions about concepts they have no prior knowledge about. This information that does not even register on patients' awareness can in fact be the most crucial aspect of their own effective self-management. Distinguishing whether a patient refrains from inquiring about a specific topic due to already possessing comprehensive knowledge or simply because they are unaware of its relevance is a nuanced capability that SAGE may refine by analyzing factors such as the patient's medical history, health literacy, previous questions, and individual characteristics. By leveraging these insights, SAGE can tailor questions to be more precise and beneficial, aligning with the patient's existing understanding and addressing potential blind spots in their awareness, while emphasizing connections to existing knowledge.

Moreover, patients' information needs and acquisition strategies change as their cognition, affect, actions, and mood evolve over time. Information must be made available when it is most useful to the patient, and creating an adaptive information system that can account for these changes over time could prove beneficial. To this end, SAGE may ask a patient about their mood or mental health, assess their emotional state and mental readiness for certain types of information, and adapt questions and answers based on this feedback.

It is also important to note that health literacy itself is an evolving characteristic, and that a patient's health literacy may be different for their specific condition, rather than being captured by a generalized metric. Especially over time, a patient may acclimate to certain medical terminology or texts related to their condition, and this familiarity may translate into higher health literacy for their specific condition. An adaptable system should be able to assess these evolving dynamics and provide information that aligns with the patient's evolving comprehension.

As a patient acquires more knowledge, they are themselves weaving a web of information to build their understanding. If SAGE can also track this web of information, it may identify and populate gaps with new information. The system may periodically ask the patient if they learned anything new about their condition from external sources and incorporate this into future iterations. The system may correct misinformation or misunderstandings, or suggest questions related to new ideas brought forth by the patient. It may also predict what future complications the patient is at risk for, and proactively offer information to address these concerns. The patient and the information system will offer each other complementary information with the goal of filling knowledge gaps and creating more comprehensive understanding for the patient's health condition.

Acknowledgments

We extend our appreciation to our partners at UCSF Medicine—Dean Schillinger, Elaine Khoong, and Wagahta Semere—for their valuable insights. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 2146752. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

Abdelghani, R.; Wang, Y.-H.; Yuan, X.; Wang, T.; Lucas, P.; Sauzéon, H.; and Oudeyer, P.-Y. 2023. GPT-3-Driven Pedagogical Agents to Train Children's Curious Question-Asking Skills. *International Journal of Artificial Intelligence in Education*.

Agrawal, G.; Kumarage, T.; Alghami, Z.; and Liu, H. 2023. Can Knowledge Graphs Reduce Hallucinations in LLMs? : A Survey. ArXiv:2311.07914 [cs].

Atir, S. 2018. *Thinking about Self and Others in the Context of Knowledge and Expertise*. Ph.D. thesis, Cornell University. Doctoral dissertation.

Cai, P.; Yao, Z.; Liu, F.; Wang, D.; Reilly, M.; Zhou, H.; Li, L.; Cao, Y.; Kapoor, A.; Bajracharya, A.; Berlowitz, D.; and Yu, H. 2023. PaniniQA: Enhancing Patient Education Through Interactive Question Answering. ArXiv:2308.03253 [cs].

Canady, B. E.; and Larzo, M. 2023. Overconfidence in Managing Health Concerns: The Dunning–Kruger Effect and Health Literacy. *Journal of Clinical Psychology in Medical Settings*, 30(2): 460–468.

Carta, S.; Giuliani, A.; Piano, L.; Podda, A. S.; Pompianu, L.; and Tiddia, S. G. 2023. Iterative Zero-Shot LLM Prompting for Knowledge Graph Construction. ArXiv:2307.01128 [cs].

Cascella, M.; Montomoli, J.; Bellini, V.; and Bignami, E. 2023. Evaluating the Feasibility of ChatGPT in Healthcare: An Analysis of Multiple Clinical and Research Scenarios. *Journal of Medical Systems*, 47(1): 33.

Chin, Y.-F.; Huang, T.-T.; Hsu, B. R.-S.; Weng, L.-C.; and Wang, C.-C. 2019. Factors associated with foot ulcer self-management behaviours among hospitalised patients with diabetes. *Journal of Clinical Nursing*, 28(11-12): 2253–2264. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jocn.14822.

Coleman, C.; Salcido-Torres, F.; and Cantone, R. E. 2022. "What Questions Do You Have?": Teaching Medical Students to Use an Open-Ended Phrase for Eliciting Patients' Questions. *HLRP: Health Literacy Research and Practice*, 6(1): e12–e16.

Costa, I. G.; Tregunno, D.; and Camargo-Plazas, P. 2020. I Cannot Afford Off-loading Boots: Perceptions of Socioeconomic Factors Influencing Engagement in Self-management of Diabetic Foot Ulcer. *Advances in Nursing Science*, 43(4): 322.

Costa, I. G.; Tregunno, D.; and Camargo-Plazas, P. 2022. Patients' Perceptions of Reasons Contributing to Delay in Seeking Help at the Onset of a Diabetic Foot Ulcer : A Grounded Theory Study. *Journal of Wound, Ostomy and Continence Nursing*, 49(5): 481–487.

Dunning, D. 2011. Chapter five - The Dunning–Kruger Effect: On Being Ignorant of One's Own Ignorance. In Olson, J. M.; and Zanna, M. P., eds., *Advances in Experimental Social Psychology*, volume 44, 247–296. Academic Press.

Elkins, S.; Kochmar, E.; Serban, I.; and Cheung, J. C. K. 2023. How Useful Are Educational Questions Generated by Large Language Models? In Wang, N.; Rebolledo-Mendez, G.; Dimitrova, V.; Matsuda, N.; and Santos, O. C., eds., *Artificial Intelligence in Education. Posters and Late Breaking Results, Workshops and Tutorials, Industry and Innovation Tracks, Practitioners, Doctoral Consortium and Blue Sky*, Communications in Computer and Information Science, 536–542. Cham: Springer Nature Switzerland. ISBN 978-3-031-36336-8.

Eltorai, A. E. M.; Ghanian, S.; Adams, C. A.; Born, C. T.; and Daniels, A. H. 2014. Readability of Patient Education Materials on the American Association for Surgery of Trauma Website. *Archives of Trauma Research*, 3(2): e18161.

Guo, C.; Si, L.; and Sun, Y. 2023. Research on the Process and Influencing Factors of Online Diabetes Information Users' Avoidance Behavior: A Qualitative Study. *Behavioral Sciences*, 13(3): 267.

Holmgren, A. J.; Byron, M. E.; Grouse, C. K.; and Adler-Milstein, J. 2023. Association Between Billing Patient Portal Messages as e-Visits and Patient Messaging Volume. *JAMA*, 329(4): 339–342.

Holstein, K.; De-Arteaga, M.; Tumati, L.; and Cheng, Y. 2023. Toward Supporting Perceptual Complementarity in Human-AI Collaboration via Reflection on Unobservables. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW1): 152:1–152:20.

IDF. 2021. *IDF Diabetes Atlas*. Brussels, Belgium: International Diabetes Federation, 10th edition.

Katz, M. G.; Jacobson, T. A.; Veledar, E.; and Kripalani, S. 2007. Patient Literacy and Question-asking Behavior During the Medical Encounter: A Mixed-methods Analysis. *Journal of General Internal Medicine*, 22(6): 782–786.

Kulshreshtha, D.; Shayan, M.; Belfer, R.; Reddy, S.; Serban, I. V.; and Kochmar, E. 2022. Few-shot Question Generation for Personalized Feedback in Intelligent Tutoring Systems. ArXiv:2206.04187 [cs].

Lee, P.; Bubeck, S.; and Petro, J. 2023. Benefits, Limits, and Risks of GPT-4 as an AI Chatbot for Medicine. *New England Journal of Medicine*, 388(13): 1233–1239. Publisher: Massachusetts Medical Society _eprint: https://doi.org/10.1056/NEJMsr2214184.

Lee, Y.; Chung, J. J. Y.; Kim, T. S.; Song, J. Y.; and Kim, J. 2022. Promptiverse: Scalable Generation of Scaffolding Prompts Through Human-AI Hybrid Knowledge Graph Annotation. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22. New York, NY, USA: Association for Computing Machinery. ISBN 9781450391573.

Longo, D. R.; Schubert, S. L.; Wright, B. A.; LeMaster, J.; Williams, C. D.; and Clore, J. N. 2010. Health Information Seeking, Receipt, and Use in Diabetes Self-Management. *The Annals of Family Medicine*, 8(4): 334–340.

Menendez, M. E.; van Hoorn, B. T.; Mackert, M.; Donovan, E. E.; Chen, N. C.; and Ring, D. 2017. Patients With Limited Health Literacy Ask Fewer Questions During Office Visits With Hand Surgeons. *Clinical Orthopaedics and Related Research*, 475(5): 1291–1297.

Meyer, L.-P.; Stadler, C.; Frey, J.; Radtke, N.; Junghanns, K.; Meissner, R.; Dziwis, G.; Bulert, K.; and Martin, M. 2023. LLM-assisted Knowledge Graph Engineering: Experiments with ChatGPT. ArXiv:2307.06917 [cs].

Molina, C. S.; and Faulk, J. 2022. Lower extremity amputation. In *StatPearls [Internet]*. StatPearls Publishing.

Neumann, K. L.; and Kopcha, T. J. 2018. The Use of Schema Theory in Learning, Design, and Technology. *TechTrends*, 62(5): 429–431.

Pirolli, P.; and Card, S. 1999. Information foraging. *Psychological Review*, 106(4): 643–675. Place: US Publisher: American Psychological Association.

Rajabi, E.; and Kafaie, S. 2022. Knowledge graphs and explainable ai in healthcare. *Information*, 13(10): 459.

Reifegerste, D.; and Hartleib, S. 2016. Hypoglycemia-related information seeking among informal caregivers of type 2 diabetes patients: Implications for health education. *Journal of Clinical & Translational Endocrinology*, 4: 7–12.

Ren, Y.; Deng, X. N.; and Joshi, K. 2023. Unpacking Human and AI Complementarity: Insights from Recent Works. *ACM SIGMIS Database: the DATABASE for Advances in Information Systems*, 54(3): 6–10.

Robbins, J. M.; Strauss, G.; Aron, D.; Long, J.; Kuba, J.; and Kaplan, Y. 2008. Mortality rates and diabetic foot ulcers: is it time to communicate mortality risk to patients with diabetic foot ulceration? *Journal of the American Podiatric Medical Association*, 98(6): 489–493.

Saini, R.; Mussbacher, G.; Guo, J. L. C.; and Kienzle, J. 2022. Automated, interactive, and traceable domain modelling empowered by artificial intelligence. *Software and Systems Modeling*, 21(3): 1015–1045.

Sansoni, J. E.; Grootemaat, P.; and Duncan, C. 2015. Question Prompt Lists in health consultations: A review. *Patient Education and Counseling*, 98(12): 1454–1464.

St. Jean, B. 2012. "I just don't know what I don't know!": A longitudinal investigation of the perceived usefulness of information to people with type 2 diabetes. *Proceedings of the American Society for Information Science and Technology*, 49(1): 1–10.

St. Jean, B. 2017. Factors motivating, demotivating, or impeding information seeking and use by people with type 2 diabetes: A call to work toward preventing, identifying, and addressing incognizance. *Journal of the Association for Information Science and Technology*, 68(2): 309–320. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/asi.23652.

Sun, J.; Xu, C.; Tang, L.; Wang, S.; Lin, C.; Gong, Y.; Ni, L. M.; Shum, H.-Y.; and Guo, J. 2023. Think-on-Graph: Deep and Responsible Reasoning of Large Language Model on Knowledge Graph. ArXiv:2307.07697 [cs].

van der Eijk, M.; Nijhuis, F. A.; Faber, M. J.; and Bloem, B. R. 2013. Moving from physician-centered care towards patient-centered care for Parkinson's disease patients. *Parkinsonism & related disorders*, 19(11): 923–927.

Wei, X.; Du, Z.; and Zhang, S. 2016. Analysis of barriers to health information seeking and utilizing in patients with diabetes. *Cross-Cultural Communication*, 12(2): 48–51.

Wen, Y.; Wang, Z.; and Sun, J. 2023. MindMap: Knowledge Graph Prompting Sparks Graph of Thoughts in Large Language Models. ArXiv:2308.09729 [cs].

Wolters-Kluwer. 2023. Wolters Kluwer survey reveals two-thirds of patients still have questions after healthcare visits.

Xu, Y.; Wang, D.; Yu, M.; Ritchie, D.; Yao, B.; Wu, T.; Zhang, Z.; Li, T. J.-J.; Bradford, N.; Sun, B.; Hoang, T. B.; Sang, Y.; Hou, Y.; Ma, X.; Yang, D.; Peng, N.; Yu, Z.; and Warschauer, M. 2022. Fantastic Questions and Where to Find Them: FairytaleQA – An Authentic Dataset for Narrative Comprehension. ArXiv:2203.13947 [cs].

Yang, L.; Chen, H.; Li, Z.; Ding, X.; and Wu, X. 2023. ChatGPT is not Enough: Enhancing Large Language Models with Knowledge Graphs for Fact-aware Language Modeling. ArXiv:2306.11489 [cs].

Zhang, Q.; Lee, M. L.; and Carter, S. 2022. You Complete Me: Human-AI Teams and Complementary Expertise. In *CHI Conference on Human Factors in Computing Systems*, 1–28. New Orleans LA USA: ACM. ISBN 978-1-4503-9157-3.

Zhu, Y.; Wang, X.; Chen, J.; Qiao, S.; Ou, Y.; Yao, Y.; Deng, S.; Chen, H.; and Zhang, N. 2023. LLMs for Knowledge Graph Construction and Reasoning: Recent Capabilities and Future Opportunities. ArXiv:2305.13168 [cs].