Using Social Robots and AI to Perform Genetic Risk Assessment for Cancer

Tyler Morris[1], Conor Brown[1], Seungwoo An[1], Jeremiah Augustine[1], Andrew Ward[1], Laura Enomoto[2,3], Erin Campbell[2,3], Xiaopeng Zhao[1]

[1]University of Tennessee, Knoxville
[2]University of Tennessee, UT Medical Center
[3]University of Tennessee, Graduate School of Medicine
tmorri35@vols.utk.edu, cbrow215@vols.utk.edu, san6@vols.utk.edu, jaugust4@vols.utk.edu, award@utk.edu, lenomoto@utmck.edu, ecampbell1@utmck.edu, xzhao9@utk.edu

Abstract

Genetic risk assessment (GRA) and genetic counseling have become integral to optimal patient care for patients with cancer. At present, there is a limited number of qualified healthcare providers who provide this service. To assist professionals in the GRA process, we have combined social robotics and retrieval-augmented generative artificial intelligence (RAG AI) to provide education related to hereditary cancer to be included in GRA sessions for individuals at risk. This GRA application pushes the boundary on previously available chatbots and AI systems by creating a novel and interactive experience enhanced by professionally verified information. In the future, we seek to improve the app as it is and obtain feedback from both GRA professionals and potential end-users that will be used to enhance the system and provide customized risk assessment. Overall, our GRA system takes the next step towards informing patients of their hereditary cancer risk and pertinent care options.

Introduction

Cancer has always been a highly feared illness, especially considering that the American Cancer Society expects roughly 611,720 deaths in 2024 alone (American Cancer Society 2024). Cancer is also a common disease, with an estimated 1 in 3 Americans diagnosed in their lifetime (American Cancer Society 2024). Public health initiatives recommend cancer screenings for common cancer types, including breast, colon, and prostate, among others (American Cancer Society 2024). Early detection can increase a patient's chance of survival. Cancer screenings may be recommended based on cancer incidence, personal and family history, and a person's genetic risk factors.

While cancer is common, only 5-10% of most malignancies have an underlying hereditary component. The practice of genetic counseling and genetic risk assessment (GRA) helps to identify individuals who have an underlying hereditary etiology for their cancer and/or may be at increased risk of developing cancer in their lifetime. Results can impact treatment and risk management options. The genetic risk assessment (GRA) and counseling process involves compiling a targeted oncology history, including demographics, personal history, family history of cancer, and other genetic risk factors (Baptista 2005, National Comprehensive Cancer Network 2023). After compiling this information, an appropriate medical provider would evaluate the person's risk of developing certain types of cancer, educate the patient, and discuss the risks and options for genetic testing (Baptista 2005). Individuals can then decide on testing, relevant potential treatment options, and cancer screening recommendations.

However, these sessions are time intensive, and a limited workforce of trained healthcare professionals provide this service. As genetic testing guidelines broaden to include an increasing number of eligible patients, health professionals are overwhelmed by the volume of referrals (Al-Hilli, 2023). In our work, we seek to alleviate professionals' burden by utilizing social robots and artificial intelligence. Our program, which we call the GRA app, uses the social robot Pepper and a GPT-4-based chatbot to conduct the educational portion of these GRA sessions, with plans to expand intake capability and risk assessment capability. Importantly, it utilizes an embedded knowledge base to generate content via a process called retrieval augments generation (RAG). Using this system, we can guarantee the accuracy of the generated content while the social robot facilitates an engaging session to ensure participant retention of information.

Previous Work

In previous research, simple AI chatbots have been used to enhance GRA sessions (Al-Hilli et al. 2023, Gortz et. Al. 2023, Chen et al. 2024, Siglen et al. 2023, Dummer et al.

2023, Pan et al. 2023, Soley et al. 2023, Soni et al. 2024, Roland-Vasquez et al. 2024, Wang et al. 2023). This primarily includes the use of chatbots on a website interface that can answer simple questions about genetics (Gortz et al. 2023, Siglen et al. 2023, Dummer et al. 2023, Soley et al. 2023), testing and treatment options (Dummer et al. 2023, Roland-Vasquez et al. 2024, Wang et al. 2023), and other general questions (Chen et al. 2024, Dummer et al. 2023, Pan et al. 2023, Roland-Vasquez et al. 2024) on different types of cancer. These chatbots are simple, as they can only answer specific questions and have a very small knowledge base. They are also limited to only one type of cancer, with pancreatic and breast being the most common. Nonetheless, these chatbots prove that AI can be used in GRA. A few studies have even been conducted to assess the feasibility of these robots, with tests being done to evaluate the accuracy (Roldan-Vasquez et al. 2024), user experience (Soni et al. 2024), and effectiveness (Al-Hilli et al. 2023) of the chatbots. Lastly, similar AI has been used in GRA settings. The two major applications of AI algorithms are genetic testing assistants (Wang et al. 2023) and risk assessment tools that use a variety of factors and patient information (Webster et al. 2024). These AI tools, however, are still under review and development.

Some have done various works with other AI chatbots and social robots. Some of these applications include robots to assist in dementia care (Yuan et al. 2023, Morris et al. 2024, Liao et al. 2023) and dementia diagnosis (Liu 2024). However, most closely related to this GRA project, we have developed a social robot-based, AI-enhanced system to inform novice dementia caregivers about any specific issues they may be having. This system uses OpenAI's GPT-4 to generate a presentation, question and answer session, and knowledge quiz based on the specified issue the caregiver is facing. This program, named RISE, also utilizes RAG to ensure only accurate and verified information is generated for each session. This system is currently in testing, but initial results already boast high levels of user acceptance and are considered reliable by dementia caregiving experts. These results come from personal feedback from our collaborators, who have been very pleased with the reliability and usability of the RISE system. Such a project proves the feasibility of using AI-enhanced robotics in healthcare applications to inform users about certain conditions with verified information.

Lastly, RAG AI is already being used in conjunction with other large language models (Chen et al. 2024, Wang et al. 2024, Atuhurra 2024, Kang et al. 2024, Yu et al. 2024, Salemi and Zamani 2024, Zhao et al. 2024). Given the technique's recent emergence, several studies have been done to investigate its capabilities (Chen et al. 2024), potential (Wang et al. 2024, Atuhurra 2024, Kang et al. 2024), and quality (Yu et al. 2024, Salemi and Zamani 2024, Zhao et al. 2024). While looking into its capabilities, one study sought to benchmark various large language models (LLMs) when combined with RAG (Chen et al. 2024). They found that all models improved thanks to RAG implementation, mostly due to reduced hallucinations or the generation of incorrect content (Chen et al. 2024). Furthermore, in studying the potential for RAG in various settings, studies have found applications in healthcare decision-making (Wang et al. 2024) and with robots for enriched interactions (Atuhurra 2024) and personalized usage (Kang et al. 2024). Such applications show that RAG can be implemented in various settings, including clinical ones. Further building off this, a few studies have gone through the process of verifying the reliability of RAG AI (Yu et al. 2024, Salemi and Zamani 2024, Zhao et al. 2024). These studies have found that, when combined with prebuilt LLM models, RAG AI-enhanced the model's ability to generate verified and correct information, greatly reducing the chances of hallucinations (Yu et al. 2024; Salemi and Zamani 2024). Furthermore, they found that the applications of RAG AI are boundless, given its relative simplicity and ability to be applied to any model (Zhao et al. 2024). This means such a process can be used in various settings, from personal use to healthcare to business. Overall, RAG AI has already proven capable and sophisticated enough for various settings, including healthcare.

RAG AI

The most crucial and technologically advanced part of the GRA application is using RAG AI. As previously mentioned, this system allows us to ensure quality responses by limiting the AI's knowledge base to certified texts. Our collaborators at our university's medical center include an oncology genetic counselor, surgeon, and nurse practitioner who each have experience in GRA and/or the treatment of colon cancer. They provided us with an array of books, websites, and medical literature related to hereditary cancer, genetic counseling, and genetic risk assessment. Content included expert guidelines from the National Comprehensive Cancer Network and the National Society of Genetic Counselors, among others.

Our process of using RAG AI, as depicted below in Figure 1, is fairly simple. First, the knowledge documents are split into sets for easier embedding. Typically, this split is determined by the developers. In this case, we grouped information based on category, whether general knowledge, ethics, genetic testing, or treatment options. Each document is then processed using text preprocessing techniques, including tokenization, stemming, and stop-word removal, to ensure that only the most relevant information is retained. Once split, the data is vectorized and put into a vector storage system. There are several options, but we use Pinecone

embedding for this step, given its easy-to-use API and integration with OpenAI's GPT models. With the information now embedded, the chosen AI model (in our case OpenAI's GPT4) can be fed the necessary information to answer user questions. For example, when a prompt (anything from the generation of slides to specific user questions) is given to the AI model, the vector storage also undergoes a similarity search. For instance, if a user queries about the genetic risk factors for colon cancer, the system searches the vector space for the most semantically similar documents or sections. Any terms or key information related to the input prompt is given directly to the model. The model then generates a response to the prompt using only the information given by the vector storage. In doing so, the AI model is limited to the information we give it while still answering the user input to its utmost capability. By limiting the input information, the AI's outputs must be based on the vectorized knowledge. Therefore, we can ensure the outputs of the generative AI will only be based on the certified information from our GRA collaborators. As such, we create a generative AI system that can create unique sessions for each user while also maintaining accuracy, unlike other automated or generative systems.

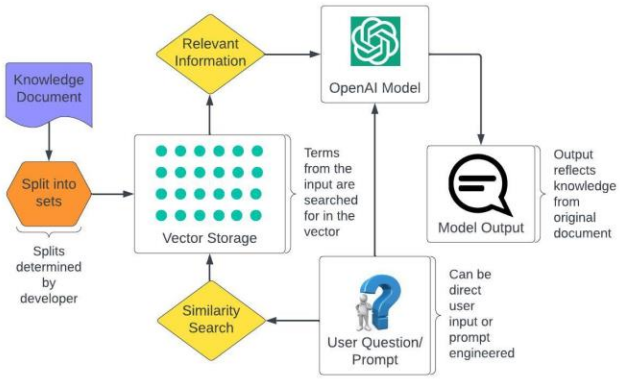

Figure 1: Diagram of RAG AI Flow

We chose to use OpenAI's model due to its extreme power and optimization, as well as its compatibility with various embedding models (OpenAI 2024). OpenAI's GPT-4 model has already proven itself to be more than capable of completing the tasks of the GRA project, primarily through its sophisticated and highly trained NLM base (OpenAI 2024). Furthermore, it is compatible with LangChain embeddings and with the social robot Pepper (which runs on Python 2.7) (OpenAI 2024). Lastly, OpenAI's models have been used in similar research in the past (Yuan et al. 2023, Liu 2024), giving us an edge in experience and making the development of the GRA app easier. Overall, the choice of the AI model comes down to robustness and familiarity, of which OpenAI has both.

Our use of RAG AI represents a significant advancement in the field of unifying representations for robot application development by integrating the strengths of both retrieval-based and generative models. By leveraging vast external knowledge bases, RAG AI enhances a robot's ability to access relevant information dynamically, which is essential for applications requiring real-time data and context awareness. This capability allows robots to respond more intelligently to user inquiries and adapt to diverse scenarios by retrieving pertinent knowledge while simultaneously generating coherent and contextually appropriate responses. RAG AI fosters a more seamless integration of diverse representations, such as text, images, and sensor data, into a unified framework. This unification not only streamlines the development process but also enables robots to operate more autonomously and effectively in complex environments. By improving information retrieval and natural language processing capabilities, RAG empowers robots to provide more accurate, context-sensitive interactions, enhancing user engagement and satisfaction in various applications, from healthcare to education.

Overall, using RAG AI in this project allows us to create a much more sophisticated system that only generates verified content. With this, we can use generative AI to create novel and customized sessions for each user while maintaining the knowledge level that traditional GRA sessions would provide. Given the verification of previous works' processes, it is safe to say that RAG AI is the cornerstone that propels this GRA application to new heights.

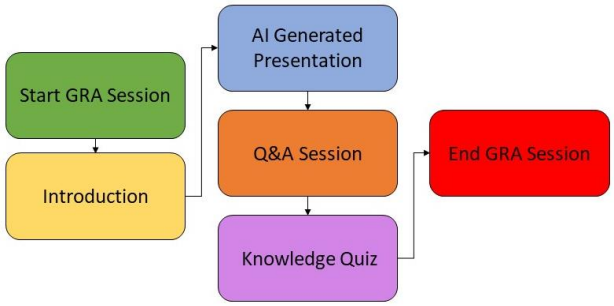

Figure 2: The GRA App Flow

The GRA App

The GRA app is divided into two parts: the social robot and the AI-enabled application. Both parts bring a unique aspect to the project, enabling it to go above and beyond previous work. We use the social robot Pepper, which is a humanoid with a non-moving face and 20 degrees of freedom. It also features a tablet interface and a microphone for voice control. Users can touch the screen or give voice commands to control the app. This robot has been used in various projects in similar research. Within these projects, participants have

been highly receptive to the robot, with many enjoying its looks, voice, and movements (Morris et al. 2024). They also enjoy interacting with the tablet interface since it is easy to control (Morris et al. 2024). We have created the GRA program as an app on Pepper, where information will be presented, users can ask questions, and the RAG AI will be used. The outline for this app is shown in Figure 2 and can be broken down into three parts: presentation, question and answer (Q&A), and knowledge quiz.

AI-Generated Presentation

The first part of the GRA app is a simple presentation of the user's diagnosis and potential genetic testing. This serves as a simple introduction to the session and directly provides relevant information to the user. The slides are generated using the previously described RAG AI system and displayed on the Pepper robot's tablet interface. The robot will narrate the slides, discussing each as the user scrolls through them. While speaking, the robot will make simple gestures and maintain eye contact with the user to foster engagement. The user can change slides using buttons on each corner of the screen and have the robot repeat a slide using another button on the bottom of the screen. Since only the text is AI-generated, it is difficult for images to be specific to each generation. However, we still use more generic images for each slide to go with the theme (general information, genetic testing, etc.). In future application iterations, we may also enhance the AI functionality to generate images to complement the text. Otherwise, we can include generic images/diagrams of the cancer or body.

Currently, our app focuses on colon cancer for simplicity as we study the efficacy of the GRA app. However, by expanding our reference library of sources, this can be expanded to any cancer type using RAG AI. We can also easily change the topic of each slide depending on user and/or expert preference. For example, we can generate slides if the healthcare providers want patients to learn about certain topics. Overall, the presentations are easily adaptable for future generations but are currently limited, given the current scope of the project.

AI-Powered Q&A

Mimicking traditional GRA sessions, individuals will have the opportunity to ask clarifying questions. The second part of the GRA app is the Q&A session, where users can ask the robot any question and receive an appropriate answer. The screen will look like a traditional messaging system, with user inputs on one side and robot responses on the other. Users can type questions via the touch screen or speak using the embedded microphone. Once they hit the "submit" button, the RAG AI will respond appropriately while using the embedded library of information. The robot will then speak this response, using simple gestures and eye contact to engage with the user.

In this session, users can ask any lingering questions that still need to be answered or clarified in the presentation section. This can range from questions on how genetic testing might work to asking for clarification on certain options. Regardless of the question, the AI will only generate an answer that uses information from the embedded knowledge base. This way, only correct answers are given. If a user's question cannot be answered using embedded knowledge, the robot will state that it cannot answer it and refer the user to a professional. Overall, the Q&A section will help ensure the patient is informed before speaking with their healthcare provider.

AI-Enabled Knowledge Quiz

The final part of the GRA app session is a simple knowledge quiz. Here, the RAG AI generates five multiple-choice questions based on the knowledge presented in the presentation section. The RAG AI will generate the correct answer, but the standard GPT-4 model will generate the three incorrect answers. This is because the RAG AI would generate three answers that are extremely similar to the correct one, making the question overtly difficult. Using the base model to generate incorrect answers, the review quiz is not too difficult and remains perfect for a simple review.

This quiz aims to test the user's knowledge and understanding. It is meant to serve as a brief review of the information to ensure comprehension. If desired, the user can repeat the presentation or Q&A session if there are topics they need help comprehending. Ultimately, the review quiz serves as a simple test of knowledge for the user to be confident in their understanding.

Discussion

Contributions and Implications

The development of the GRA app represents a significant advancement in AI-powered genetic risk assessment. RAG AI sets a new bar in using chatbot systems, enabling high-quality outputs not corrupted by misinformation. Additionally, since this algorithm can be easily updated via new embeddings, the RAG AI can swiftly account for and use new techniques and information as the field progresses. This includes changes in cancer treatment, genetic testing, and new forms/methods of GRA and genetic counseling. Furthermore, our use of RAG AI also allows for more customizable sessions. Just as it can be updated with new information from the research field, the system can also be updated for

each user to account for personal information and medical history. This includes family histories, lifestyles, and behaviors that can affect their health. Compared to most AI-based chatbots, this system is the easiest to personalize since the embeddings can be done quickly and automatically. Additionally, given the generative nature of the presentations, the personalized information can easily be added which each session, making the sessions more engaging and useful for each individual user. Overall, the GRA app sets new bounds in how generative AI can be used in the medical field, as it showcases a secure, verifiable, easily personalized, and updatable system while generating unique sessions for each user.

The GRA app also furthers the field of social robotics. By using the Pepper robot, we show another application for the interactive robot in a new sub-specialty in healthcare. Here, the use of the robot makes the session far more interactive than a web- or tablet-based interface. As such, the session is more engaging and could be more impactful for users. While the use of the robot is not as unique, integrating RAG AI is yet another research milestone. Our success in integrating this model with a social robot shows that these robots can be custom-tailored with information, making them more personalized for each user. As mentioned, this can be updated easily, making the system readily adaptable. Overall, the GRA app is a novel development in social robotics.

To summarize, the system has the potential to scale cancer genetic counseling services and make them more accessible to a wider population by leveraging RAG AI and social robotics. The personalized health information provided by the GRA App empowers patients to make informed decisions about their health and enhances patient-provider communication, fostering shared decision-making.

Larger Impacts

From a healthcare perspective, the GRA app could alleviate some burdens on genetic counselors and other healthcare professionals. By automating initial risk assessments and providing detailed information to patients, healthcare providers can focus their time and resources on more complex cases and treatment planning. This efficiency could lead to more effective use of healthcare resources, improved patient outcomes, and potentially reduced healthcare costs.

On an individual level, the GRA app empowers users with knowledge about their genetic risks and healthcare options. Informed patients are better equipped to make proactive decisions about their health, such as undergoing appropriate screenings, adopting preventive measures, or seeking timely medical advice. This empowerment can lead to a more engaged and health-conscious population, improving public health outcomes. Moreover, the GRA app's personalized approach ensures that users receive information specifically relevant to their unique genetic risk profile. This level of personalization enhances the user experience and ensures that the advice and recommendations provided are practical and actionable for each individual.

Early interventions are typically less expensive and more effective than treatments required at later stages of the disease (Luvhengo et al. 2023). Consequently, the GRA app could contribute to significant cost savings for individuals and healthcare systems.

Limitations and Challenges

Although the GRA app does set new ground regarding RAG AI and social robot development in healthcare, it has limitations and challenges moving forward. The major limitation is its limited scope. To ensure adequate case volume and efficacy for our initial study, the GRA app is currently only designed for individuals with colon cancer. At present, the app is not well suited to provide GRA sessions to persons at risk of other cancer types. Furthermore, the limited scope of sessions makes it difficult to get all of the information to the users. The lack of personalized data input by the patient does limit the personalized education available to users. As such, information provided through the robot, while verified by experts, is very generic. Full risk assessment and whether a patient qualifies for hereditary cancer testing outside of their colon cancer diagnosis is not available through the GRA app at this time.

In addition to the limitations of the app and AI, the robot also has a few limitations. Although we have not yet conducted user testing, we fear the robot may have issues regarding user interactions. First, the tablet may be small and hard to read. Although the button layout and texts are as large as possible, concerns about a bland design or confusing layout may arise. As such, we are constantly looking to improve the UI, with many changes being suggested by collaborators and users.

Furthermore, we are concerned with the speech-to-text function used in the Q&A session. We currently use the embedded function with the robot, but there are concerns with different accents. Variations in speech and accents may be difficult for the system to pick up on, leading to frustrations as the microphone does not hear them properly. This will have to be tested with end-users, but we seek to use more advanced systems in the future. Nonetheless, despite these limitations and challenges, we seek to have the GRA app push the boundaries in the field of AI and robots in healthcare and health informatics.

As with any technology used in healthcare, ethical considerations are paramount. Integrating AI and social robots in genetic counseling raises several ethical questions. One major concern is data privacy and security. The GRA app handles sensitive personal information, including genetic

data and family health history. Ensuring this information is stored securely and responsibly is crucial to maintaining patient trust. The most efficient way to address this is to implement whatever data encryption and protection the facility uses. Most hospitals and medical facilities use highly secure data protection algorithms to protect patient information, and since the GRA app would work on their network it would be easy to implement such protections. Another ethical consideration is the potential for bias in AI-generated content. Although RAG AI uses verified sources to generate responses, it is essential to continually monitor and update these sources to prevent disseminating outdated or biased information. Additionally, the AI system must be transparent about its limitations. If the system cannot provide a specific answer, it should communicate this to the user and suggest consulting a human professional. There is also the question of accessibility and inclusivity. The GRA app must be designed to be accessible to users from diverse backgrounds, including those with different levels of digital literacy. As previously discussed, the speech-to-text function should be capable of understanding various accents and dialects to ensure that all users can interact with the system effectively.

Future Work

As this project progresses, we seek to test and further improve the GRA app to accommodate professional standards and user desires better. This process will begin with two small-scale studies. First, our collaborators and other professionals in the GRA field will evaluate the system's quality and reliability. They will be invited to participate in sessions and provide feedback on the UI, session content, and any other critiques they have. We will measure quality and reliability by comparing the AI-generated content to that provided by GRA experts, ensuring all generated content is accurate. Moving on from this, we also wish to host workshops and study sessions with potential end-users. In these sessions, people who have been through or recommended GRA will be invited to participate in a robot-guided session. There, they can provide feedback on the system's various aspects. This can include the UI and its layout, the app's ease of use, or even how useful the generated content is. This will focus more on user acceptability and app feasibility, as the content presented has already been verified. While these studies are not fully planned out at this point in time, we will put together a more comprehensive plan as the app is completed and we identify areas that need to be studied more intensely. Altogether, these two studies will provide insight into which aspects of the system need improvement and which aspects need to be kept the same.

We also want to make technological improvements to the app eventually. For example, we want to incorporate more personalized aspects for each user, including family history, personal health and habits, and other cancer risk factors. In this, we can make the GRA app more personalized and tailored to each user, resulting in a more impactful session. We also want to branch out to other cancer types. We are currently focusing on colon cancer, but with more embedded information, we can tailor the session to various other illnesses depending on the user.

Additionally, we want to use more sophisticated systems for speech-to-text to address concerns about the robot properly understanding different accents. Our current projection is to use an AI-enhanced system, like OpenAI's Whisper, that is trained to understand a variety of accents. This would likely improve user acceptance of the system, as users would be less frustrated with the speech functionality and interactions. Lastly, we will also incorporate the changes suggested in the studies we plan to run. This might be changes to the UI, the robot, or the general app design. The goal of these changes is to improve user acceptance, which will, in turn, make the app more useful and get more use from end-users. Overall, we hope to continue improving and changing this app to serve its end goal of helping people be more informed about their health and healthcare options.

Conclusion

We seek to have our app be one of the best novel tools to assist in GRA. Using novel systems like RAG AI and social robotics, this system can provide information to users reliably and effectively. Our use of RAG AI enables the system to generate new content for each session while maintaining validity through verified sources. Furthermore, using social robots creates a more engaging session, resulting in a more impactful session for users than a web-based service. And while we still have room for improvement, we have already set goals to make the GRA app better for users in the future. Using novel technology and systems, our GRA app sets a new bar in automated health sessions to assist professionals better and benefit patients.

Acknowledgments

This research has been partially supported by a seed grant from the AI Tennessee Initiative at the University of Tennessee Knoxville and The University of Tennessee Medical Center Cancer Institute Cancer Research Endowment Gift Fund.

References

Al-Hilli, Z.; Noss, R.; Dickard, J.; Wei, W.; Chichura, A.; Wu, V.; Renicker, K.; Pederson, H. J.; and Eng, C. 2023. A Randomized Trial Comparing the Effectiveness of Pre-test

Genetic Counseling Using an Artificial Intelligence Automated Chatbot and Traditional In-person Genetic Counseling in Women Newly Diagnosed with Breast Cancer. In *Annals of Surgical Oncology*. 5990-5996.

American Cancer Society. 2024. *Cancer Facts & Figures 2024*. Atlanta: American Cancer Society.

Atuhurra, J. 2024. Enrich Robots with Updated Knowledge in the Wild via Large Language Models. In Proceedings of ACM Conference 2024. Washington, D.C.

Baptista, P.V. 2005. Principles in genetic risk assessment. In *Therapeutics and Clinical Risk Management.* 1(1), 15-20. Caparica: Dove Medical Press.

Chen, D.; Parsa, R.; Hope, A.; Hannon, B.; Mak, E.; Eng, L.; Liu, F.; Fallah-Rad, N.; Heesters, A. M.; and Raman, S. 2024. Physician and Artificial Intelligence Chatbot Responses to Cancer Questions From Social Media. In *JAMA Oncology*. 10(7). 956-960.

Chen, J.; Lin, H.; Han, X.; and Sun, L. 2024. Benchmarking Large Language Models in Retrieval-Augmented Generation. In the Proceedings of the 38th AAAI Conference on Artificial Intelligence. 17754-17762.

Dummer, R.; Gyorki, D. E.; Hyngstrom, J. R.; Ning, M.; Lawrence, T.; and Ross, M. I. 2023. Use of Artificial Intelligence Chatbots for Cancer Treatment Information. In *JAMA Oncology*. 9(10). 1459-1462.

Felding, S. A.; Koh W. Q.; Teupen S.; Budak, K. B.; Uribe, F. L.; and Roes, M. 2023. A Scoping Review Using the Almere Model to Understand Factors Facilitating and Hindering the Acceptance of Social Robots in Nursing Homes. In *International Journal of Social Robotics*. 15. 1115-1153. https://doi.org/10.1007/s12369-023-01012-1.

Gortz, M.; Baumgartner, K.; Schmid, T.; Muschko, M.; Woessner, P.; Gerlach, A.; Byczkowski, M.; Sultmann, H.; Duesing, S.; and Hohenfellner, M. 2023. An artificial intelligence-based chatbot for prostate cancer education: Design and patient evaluation study. In *Digital Health.* 9. 1-11.

Kang, H.; Moussa, M. B.; and Magnenat-Thalmann, N. 2024. Nadine: an LLM-driven intelligent social robot with affective capabilities and human-like memory. arXiv:2405.20189v1

Liao, Y.; Jao, Y.; Boltz, M.; Adekeye, O. T.; Berish, D.; Yuan, F.; and Zhao, X. 2023. Use of a Humanoid Robot in Supporting Dementia Care: A Qualitative Analysis. In *Geriatrics*. 9. 1-14.

Liu, Z. 2024. Uniting Cognitive Models and AI: Early Alzheimer's Screening through Language Analysis. PhD Dissertation, Department of Mechanical, Aerospace, and Biomedical Engineering, The University of Tennessee, Knoxville, TN.

Luvhengo, T.; Molefi, T.; Demetriou, D.; Hull, R.; Dlamini, Z. 2023. Use of Artificial Intelligence in Implementing Mainstream Precision Medicine to Improve Traditional Symptom-driven Practice of Medicine: Allowing Early Interventions and Tailoring better-personalised Cancer Treatments. In *Artificial Intelligence and Precision Oncology*. 49-72. https://doi.org/10.1007/978-3-031-21506-3_3

Morris, T.; Walker, S.; Holder, L.V.; Vaughan, E.; Petrovsky, D. V.; Swaminathan, S.; and Zhao, X. 2024. In Proceedings of CHASE 2024. Wilmington: ACM/IEEE.

National Comprehensive Cancer Network. 2023. Genetic/Familial High-Risk Assessment: Colorectal. In *NCCN Clinical Practice Guidelines in Oncology*.

OpenAI. 2024. GPT-4 Technical Report. arXiv:2303.08774v6.

Pan, A.; Musheyev, D.; Bockelman, D.; Loeb, S.; and Kabarriti, A. E. 2023. Assessment of Artificial Intelligence Chatbot Responses to Top Searched Queries About Cancer. In *JAMA Oncology*. 9(10). 1437-1440.

Roland-Vasquez, E.; Mitri, S.; Bhasin, S.; Bharani, T.; Capasso, K.; Haslinger, M.; Sharma, R.; and James, T. D. 2024 Reliability of artificial intelligence chatbot responses to frequently asked questions in breast surgical oncology. In *Journal of Surgical Oncology*. 1-16.

Salemi, A., and Zamani, H. 2024. Evaluating Retrieval Quality in Retrieval-Augmented Generation. In Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval. Washington, DC. Association for Computing Machinery. doi.org/10.1145/3626772.3657957

Siglen, E.; Vetti, H. H.; Augestad, M.; Steen, V. M.; Lunde, A.; and Bjorvatn, C.; 2023. Evaluation of the Rosa Chatbot Providing Genetic Information to Patients at Risk of Hereditary Breast and Ovarian Cancer: Qualitative Interview Study. In *Journal of Medical Internet Research.* 25.

Soley, N.; Klein, A.; Taylor, C. O.; Nguyen, M.; Eqachiw, G.; Shah, H.; and Bodurtha, J. 2023 Feasibility of the Genetic Information Assistant Chatbot to Provide Genetic Education and Study Genetic Test Adoption Among Pancreatic Cancer Patients at Johns Hopkins Hospital. In Proceedings of the AMIA Joint Summits on Translational Science. 497-504

Soni, H; Morrison, H.; Vasilev, D.; Ong, T.; Wilczewski, H.; Allen, C.; Hughes-Halbert, C.; Ritchie, J. B.; Narma, A.; Ivanova, J.; Bunnel, B. E.; and Welch, B. M. 2024. User experience of a family history chatbot: A quantitative analysis. In *Health Informatics Journal*. 1-16

Wang, C.; Lu, H.; Bowen, D. J.; and Xuan, Z. 2023. Implementing digital systems to facilitate genetic testing for hereditary cancer syndromes: An observational study of 4 clinical workflows. In *Genetics in Medicine*. 25.

Wang, C.; Ong, J.; Wang, C.; Ong, H.; Cheng, R.; and Ong, D. 2024. Potential for GPT Technology to Optimize Future Clinical Decision-Making Using Retrieval-Augmented Generation In *Annals of Biomedical Engineering*. 52. 1115-1118.

Webster, E. M.; Perez, L.; Ahsan, M. D.; Levi, S.; Chandler, S.; Thomas, C.; Babagemi, K.; Sharaf, R. N.; and Frey, M. K. 2024. Integration and usability of a digital cancer risk stratification tool to optimize identification of patients at risk for hereditary cancers: A pilot study. In *Gynecologic Oncology*.

Yu, H.; Gan, A.; Zhang, K.; Tong, S.; Liu, Q.; and Liu, Z. 2024. Evaluation of Retrieval-Augmented Generation: A Survey. arXiv:2405.07437v2

Yuan, F.; Boltz, M.; Bilal, D.; Jao, Y.; Crane, M.; Duzan, J.; Bahour, A.; and Zhao, X. 2023. Cognitive Exercise for Persons With Alzheimer's Disease and Related Dementia Using a Social Robot. In *IEEE Transactions on Robotics*. 39(4). 3332-3346.

Zhao, P.; Zhang, H.; Yu, Q.; Wang, Z.; Geng, Y.; Fu, F.; Yang, L.; Zhang, W.; Jiang, J.; and Cui, B. 2024. Retrieval-Augmented Generation for AI-Generated Content: A Survey. arXiv:2402.19473v6