

Enhancing Health Recommendations through Patient Metadata Integration: A Persona-Based Evaluation Approach

Khushboo Thaker^{1,*}, Peter Brusilovsky¹, Daqing He¹, Youjia Wang², Mohammad Hassany¹, Behnam Rahdari¹, Heidi Donovan² and Young Ji Lee²

¹School of Computing and Information, University of Pittsburgh, Pittsburgh, USA

²School of Nursing, University of Pittsburgh, Pittsburgh, USA

Abstract

Patients and caregivers managing chronic diseases frequently seek information to navigate symptoms and make informed decisions as their conditions evolve. Health recommender systems are crucial in delivering personalized and relevant content tailored to individual needs. This paper presents a novel approach to enhancing content-based health recommendations by integrating patient metadata and trajectories. We developed a system that leverages these elements and evaluated its effectiveness with the assistance of healthcare professionals, specifically nurses. To address diverse needs and varied patient scenarios, we employed personas that simulate different patient journeys, ensuring the system's robustness across multiple contexts. Our findings indicate that incorporating user profiles with trajectory-based features significantly enhances the recommender system's performance. Our model enhanced with both profile and trajectory, achieved a Mean Average Precision (MAP) of 0.7318, demonstrating a 17 point increase. Our future research aims to utilize personas during the cold-start recommendation phase, where patient metadata is limited and persona can work as digital twin supporting patients journey.

Keywords

Health Recommender Systems, Personalized Health Information, Persona based Evaluation, Patient Metadata Integration

1. Introduction

Individuals with chronic illnesses such as diabetes, cancer, and heart disease are often deeply involved in managing their health. They actively seek out information to make informed decisions and effectively manage their conditions [1]. Traditionally, reliable health information has been disseminated through patient education materials and health literacy workshops [2, 3]. However, these resources are typically designed for a general audience and often fail to meet the specific needs of individual patients [3, 4].

In pursuit of more tailored health information, many patients turn to online search platforms. Unfortunately, these platforms frequently fall short in delivering the specialized, personalized content that patients require [5, 6]. As a result, patients increasingly seek information in disease-specific online health communities (OHCs), where they can connect with others who share similar experiences [7]. While beneficial, OHCs also pose significant risks of spreading misinformation, lacking the necessary infrastructure to ensure patients access the most relevant and verified information [8].

Consumer health search engines and recommender systems offer a promising solution to these challenges by providing personalized and reliable health information. Recent research has shown that enhancing the user experience on consumer health search engines, for example, through query reformulation for specialized health terms, can significantly improve search outcomes [9, 10]. Similarly,

HealthRecSys'24: The 6th Workshop on Health Recommender Systems co-located with ACM RecSys 2024

*Corresponding author.

✉ k.thaker@pitt.edu (K. Thaker); peterb@pitt.edu (P. Brusilovsky); dah44@pitt.edu (D. He); yow14@pitt.edu (Y. Wang); moh70@pitt.edu (M. Hassany); ber58@pitt.edu (B. Rahdari); donovanh@pitt.edu (H. Donovan); leeyoung@pitt.edu (Y. J. Lee)

🌐 <http://kthaker.com> (K. Thaker)

🆔 0000-0003-3619-9376 (K. Thaker)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



Name	Maria -
Profession	Administrative Assistant – stopped working after cancer
Marital Status	Married
Children	2 Children
Maria's Age	58 (55 when diagnosed)
Current Disease Trajectory State	Taxol plus Avastin for 4th recurrence; life threatening reaction to carboplatin (survivorship for progression/recurrence)
Current Disease State	disease is progressing (CA-125 rising)
Maria's all Needs	drug side effect management, fatigue
<p>Maria is a 58 y.o. married woman with 2 adult daughters. She was diagnosed 3 years ago at age 55 and has had 4 recurrences since then. She has been on taxol plus avastin, but yesterday was taken off the Avastin for "too much protein in her urine". She is worried because the CA125 yesterday had taken a jump from 284 to 379. She also is desperate for something to help her fatigue. It is the one symptom she can't control. She is too tired to do normal house cleaning, crafts, going out to shop, much of anything she used to participated in. She continues to need more sleep and sometimes frequent naps. She gets depressed when she can't do what she wants to do. She does feel that she has a good quality of life....and still living after 3 1/2 years is a big deal to her!! She just wishes she could control and improve this fatigue.</p>	

Figure 1: Persona depicting a hypothetical patient Maria at disease progression stage

consumer health recommender systems can enhance user satisfaction by suggesting personalized content that incorporates user feedback [11, 12].

Despite these advancements, existing health recommender systems often focus on addressing users' immediate needs without sufficiently integrating user meta data like user profile and trajectory information. This limitation is particularly problematic in cold-start situations, where the system lacks adequate data to generate accurate recommendations. Additionally, current evaluation methods for these systems frequently overlook the diversity of user experiences, failing to account for the range of patient scenarios encountered in practice.

The primary motivation for this research is to address these limitations by developing a content-based recommendation system that integrates user profiles and trajectory information to enhance recommendation accuracy. Specifically, we aim to investigate how the inclusion of trajectory-based features can better address the diverse information needs of patients at different stages of their health journey. The key contributions of this paper are as follows:

- The design and implementation of a content-based health recommender system that effectively integrates user profiles and trajectory data, significantly improving recommendation accuracy.
- The introduction of a novel evaluation framework that employs personas, generated through clustering online health community data and refined via expert feedback, to simulate diverse patient scenarios.

The remainder of this paper is organized as follows: Section 2 details the system architecture and the methodology employed in this study. Section 3 presents the experimental results and discusses the findings. Section 4 offers implications, and potential areas for future research. Finally, Section 6 concludes the paper by summarizing the key findings and contributions.

2. Methodology

2.1. Creating Personas for Comprehensive System Evaluation

In designing a human-centered recommender system, it is critical to understand its target users, examine their information needs, and understand in which aspects the users might be different from each other. Specifically, in our case, where it is hard to evaluate the ongoing research on large patient population, which makes it essential to create a comprehensive user base to cover the broad spectrum of patient needs. For this we created ten Ovarian Cancer (OvCa) Patient and Caregiver personas by triangulating multiple sources, including online cancer forums, direct interviews with patients and caregivers, domain expert input, and clinical notes. The detailed process of persona creation is discussed in [13, 14]. To

Table 1

Clustering Features. All features were categorical.

Feature Name	Possible Values	Feature Name	Possible Values
User Type	Patient or Caregiver	Age	Above 50 or Below 50
Has Kids	Yes or No	Marital Status	Married, Separated, Divorced, Widow
Has Caregiver	Yes or No	Employed	Yes or No
Insurance	Yes or No	Cancer Stage	Early or Advanced
Disease Trajectory	Initial Treatment after Diagnosis, Survivorship after Initial Treatment (Remission), Survivorship during Treatment for Progression, Survivorship during Treatment for Recurrence, Survivorship after Recurrence/Progression, Advance Care Planning (End of Life)	Symptom Needs	Yes or No
Emotional Needs	Yes or No	Self-management Needs	Yes or No
Financial Needs	Yes or No	Treatment Needs	Yes or No
Sexuality Needs	Yes or No		

depict the appearance and details of a persona, we have included a comprehensive representation in Figure 1. The process include clustering of user base and than each cluster was enhanced by *OvCa* health professional to develop a general persona representing these cluster. Table 1 lists all the features utilized for clustering.

While creating personas it was revealed that the disease trajectory and information needs play key roles in differentiating groups of similar users. In contrast, demographics-related metadata, which serve as a key factor in grouping users in many existing recommender systems, were scarce and did not play a major role in forming the personas. This study provided us with important design implications for the development of personalized recommendation algorithms. First, the algorithms have to focus on adapting to both patients' disease trajectory and expressed information needs to recommend the most relevant documents. Since both factors could change over time, the system has to constantly monitor the users and keep their profiles up-to-date. Second, the presence of stable user clusters enables the system to perform efficient "cold-start" personalization at the start of the user's interaction with the system by matching the user to one of the clusters and offering a group-level adaptation. These implications affected the design of the system algorithms discussed in Section 2.3.

2.2. Data Source

The primary source of documents for this study is the HELPeR library ¹, a specialized repository meticulously organized into a topic hierarchy known as *OvCa TH* (Ovarian Cancer Topic Hierarchy) [15]. The *OvCa TH*, developed by a dedicated team of oncology nurses, categorizes documents into 17 broad categories and 90 specific sub-categories. This hierarchical structure ensures that documents are well-distributed collection across different categories. Such distribution is crucial for covering the diverse informational needs of *OvCa* patients and caregivers, ensuring that the library provides comprehensive and balanced content. The library consists of 2808 documents, a detail statistic about the library is provided in Table 3 and top 10 domains are listed in Table 2.

2.3. Recommendation Methods:

Our approach employs a content-based recommender system, which focuses on analyzing the content of documents to generate personalized recommendations. The input to the system includes explicit

¹<https://sites.pitt.edu/~dah44/helper/>

Table 2

Top 10 domains in HELPeR Library

Domain	Number of Articles
www.cancer.org	216
www.cancer.net	172
www.cancer.gov	155
www.mdanderson.org	61
www.mskcc.org	58
www.curetoday.com	56
www.ncbi.nlm.nih.gov	49
www.oncolink.org	38
clinicaltrials.gov	35
www.mayoclinic.org	31

Table 3

Summary Statistics of the HELPeR Library

Statistic	Value
Total Number of Documents	6,497
Manually Curated Documents	2,808
Crawler Collected Documents	3,689
Average Curated Documents per Broad Category	147
Median Curated Documents per Broad Category	109
Number of Unique Domains Crawled	141

user information, such as the patients’s profile (in our case Persona).

The output is a ranked list of documents that are tailored to the user’s specific needs.

2.3.1. Candidate Retrieval

In the initial phase of our recommendation pipeline, we focused on identifying a suitable candidate retrieval method. This step is crucial as it determines the pool of documents that will be further ranked and refined. The input query to this retrieval method is the need topics mentioned in the Patient Personas like nephropathy, financial assistance, cancer pain and hormonal therapy.

Our primary goal was not to find the absolute best model, but to select a reliable candidate ranking method that would allow us to conduct further experiments effectively. Based on previous benchmarks [16] on biomedical retrieval tasks we evaluated five underlying retrieval models: BM25, a well-established probabilistic model, served as our baseline [17]. We also explored BERT-based models, including BERT [18] and BioBERT [19] and MedCPT [16], which are designed to capture contextual nuances in text, particularly in the biomedical domain [16] and ColBERT [20, 17, 16], which has shown to be a strong baseline and generalized to bio-medical domain.

To conduct all these experiments, we utilized Pyserini [21], which integrates PyTorch and FAISS-based indexing for dense retrieval [22] and Lucene engine for BM25 retrieval. For FAISS indexing we utilized default parameters with number of clusters to be 1 as our dataset is small and we are not interesting in optimizing the performance. Further investigation is required to identify the best-performing models, we leave these enhancements for future as discussed in Section 4.1. In our dataset patient articles can be smaller less than 100 words to huge more than 11000 words with mean $2069(\pm 4091)$ we have indexed the documents at the section level. In this study, we employed several re-ranking methods to refine the initial retrieval results based on both patient profiles and their trajectories.

2.4. Profile and Trajectory-Based Re-Ranking

In this study, we implemented two primary re-ranking approaches to refine the initial retrieval results by enhancing the relevance to the patient’s specific profile and trajectory:

Table 4
Disease Trajectories for Ovarian Cancer Patients

Trajectory	Description
End of Life	The patient is in an advanced stage of disease where treatment is no longer curative but palliative. Focus may shift to symptom management, including pain relief, managing symptom, and emotional support. Advanced care planning, hospice care, and addressing emotional needs such as anxiety, grief, and fear of death become priorities.
Survivorship After Treatment	The patient has completed initial treatment and is transitioning to post-treatment care. This stage involves managing long-term side effects. Emotional challenges may include fear of recurrence and adjusting to a new normal. Patients focus on recovery, rehabilitation, and monitoring for recurrence.
Survivorship	The patient is currently receiving initial treatment for a new diagnosis. Information needs center around understanding treatment options, managing side effects, and coping with the emotional stress of a new cancer diagnosis. Patients may need support in navigating treatment decisions and maintaining quality of life during active treatment.
Survivorship During Progression/Recurrence	The patient is dealing with disease recurrence or progression, requiring new treatment strategies. Focus shifts to managing new or worsening symptoms, exploring second-line treatments, and addressing the emotional impact of disease recurrence. Information on advanced therapies, clinical trials, and symptom management becomes critical.
Previvorship	The patient is at high risk for developing ovarian cancer due to genetic predispositions but has not yet been diagnosed. Information needs include preventive strategies, such as risk-reducing surgeries, regular screening, and lifestyle modifications. Emotional challenges may include anxiety over potential diagnosis and decision-making regarding preventive measures.
Survivorship After Treatment (Recurrence/Progression)	The patient has completed treatment for recurrent or progressive disease and is now managing life post-treatment. There is a focus on managing long-term side effects of the second treatment, ongoing surveillance, and addressing fears related to further progression. Emotional support may be needed for coping with uncertainty, and physical rehabilitation may be required to manage lingering symptoms from aggressive treatments.

- **Profile Re-Ranking (Keyphrase-Based) :**

We extracted keyphrases from the patient’s profile², which includes relevant attributes such as medical history, demographics, and treatment details. These keyphrases were then appended to the patient’s stated need to generate a refined rank list. This approach aims to align the retrieved documents more closely with the patient’s overall profile. We initially tried to use complete profile embeddings as well as all keyphrases but both of these methods performed poorly so we decided to restrict our keyphrases based on semantic types including drug, disease, therapeutic procedure, laboratory test.

- **Trajectory-Based Re-Ranking:** In this method, the patient’s trajectory, or their current stage in the progression of their disease, was integrated into the recommendation process to provide more personalized information. A patient’s disease trajectory refers to the stages they go through, from diagnosis to treatment and beyond, which significantly impact their information needs. For ovarian cancer patients, the type of information required can vary drastically depending on whether they are in early treatment, post-treatment survivorship, or dealing with a recurrence of the disease.

²Keyphrase Extraction was done using Scispacy with model en_core_sci_scibert. UMLS linking was utilized and semantic type of each keyphrase was extracted -<https://github.com/allenai/scispacy/tree/main>

To incorporate trajectory-based re-ranking, we created embeddings of the trajectory descriptions as listed in Table 4, which were used as query information to re-rank the top results. This approach ensured that documents related to the same information need but more specific to the patient’s trajectory were ranked higher. For example, if the patient was in the “Survivorship After Treatment” stage, their primary need might be managing side effects from chemotherapy. By embedding the description of this trajectory, the system would prioritize documents that address not only general chemotherapy side effects but those particularly relevant to patients who have completed treatment and are focusing on recovery and long-term management.

- **Combining Profile and Trajectory Re-Rankings:** We merged the rank lists generated from both the profile-based and trajectory-based re-rankings to produce a final combined rank list. This approach leverages the strengths of both methods to deliver a more comprehensive and contextually relevant set of results.

To generate a final, more comprehensive ranking, we combined the results from the profile-based and trajectory-based re-rankings using Reciprocal Rank Fusion (RRF). RRF is a technique used to merge multiple ranking lists by assigning a score to each document based on its position in each list. The RRF score for a document d in a ranking list r is calculated as:

$$\text{RRF_score}(d) = \sum_{r=1}^n \frac{1}{k + \text{rank}_r(d)}$$

where k is a constant (typically set to 60), and $\text{rank}_r(d)$ is the rank of document d in ranking list r . For our experiments, we gave equal weight to all ranking lists, ensuring a balanced integration of both profile-based and trajectory-based relevance. This combined approach allows us to generate a final ranking that takes into account both the patient’s detailed profile and their trajectory, leading to a more personalized and contextually appropriate set of recommendations. We restricted RRF to top 10 documents from each rank-list.

2.5. The Study Design

The study was designed to evaluate the effectiveness of different document recommendation algorithms within the HELPeR system, based on feedback from expert caregivers. By simulating real-world decision-making processes, the study aimed to assess how well these algorithms supported the selection of relevant documents for ovarian cancer patients and their caregivers. The study involved several key steps, as illustrated in Figure 3, guiding participants through a structured sequence of tasks that included document annotation, feedback, and selection.

Step 1: Pre-Survey

The study began with a pre-survey, which participants were required to complete. The survey collected key demographic information and assessed participants’ caregiving backgrounds, as summarized in Table 5. To reduce the number of articles for review experts were asked to look for a single information need of the patient. This limitation of the study is discussed in Section 4.1

Step 2: Persona Familiarization

Following the pre-survey, participants were introduced to a randomly selected hypothetical patient persona as discussed in Section 2.1. Participants were given time to familiarize themselves with the persona and were asked to confirm their understanding of the patient’s context. This step ensured that participants felt comfortable and confident in making document selections on behalf of the persona.

Step 3: Initial Document Feedback (Subtask 1)

The document list was created by first selecting the top ‘ n ’ ($n=10$) documents generated by each of the algorithms discussed in Section 2.3. These top documents were then combined to form a single list of unique documents. To eliminate any positional bias [23], the combined list was randomly sorted. Participants were subsequently presented with this randomized list of unique documents for their review and feedback. Every time a participant reviewed an article, we ensured that the persona was accessible to them (Refer 2). This approach allowed participants to refer back to the persona as needed,

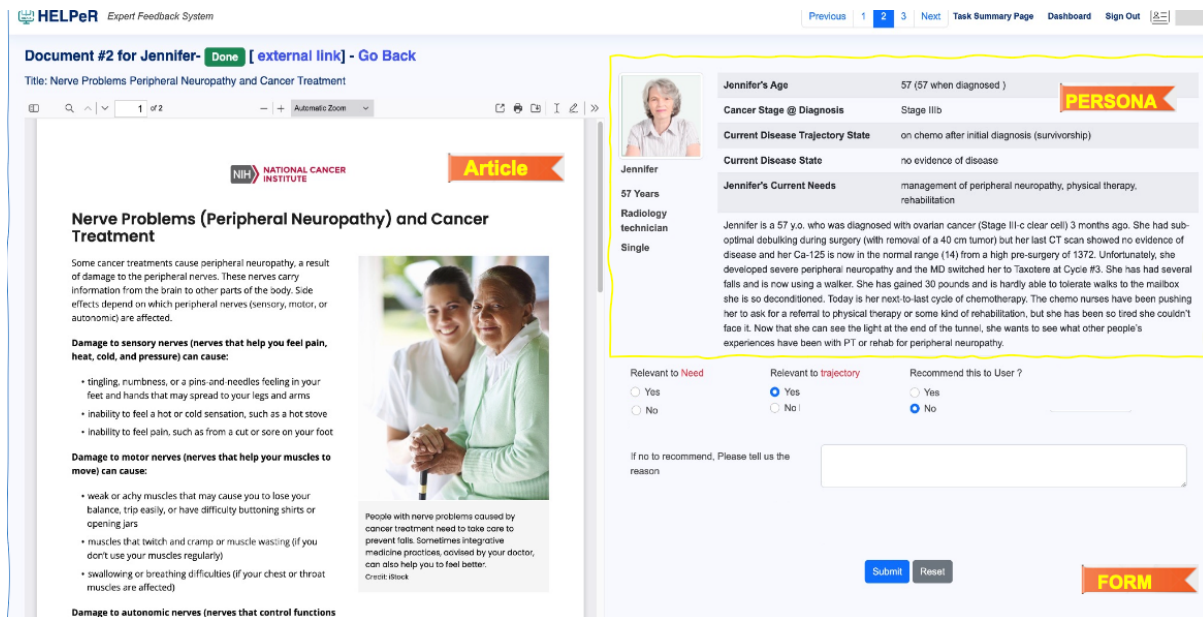


Figure 2: Feedback System used in the study. The left side displays the recommended article, while the right side shows the persona followed by the feedback form.

minimizing the risk of recall bias that could occur if they had to rely solely on their memory of the persona’s details. For each document, participants were required to provide detailed feedback [24]. This feedback focused on the relevance of the document to the persona’s needs and trajectory and the participant’s assessment of the document’s usefulness.

Step 4: Additional Document Feedback (Subtask 2)

Based on the feedback provided in Subtask 1, the HELPeR system generated a second set of documents. From the additional documents, the unique documents were selected to complement the initial set. These were shown to the participants for further feedback. This step was crucial in ensuring that the study captured feedback on a comprehensive set of relevant documents.

Step 5: Document Selection for Patient

In this step, participants were explicitly asked to assume the role of a nurse responsible for the patient. They were instructed to select up to three articles from the documents provided that they deemed most appropriate for the patient to read. Participants were required to justify their selections by indicating specific qualities that made these documents and how to perceive the understanding of this document for the presented patient. (Refer Table 5)

Step 6: Overall Feedback

Finally, participants were asked to provide overall feedback on the study. This included their impressions of the persona they interacted with, the quality and relevance of the documents encountered, and their thoughts on the HELPeR system. The design of this study was carefully structured to simulate a realistic caregiving scenario, allowing us to gather valuable insights into the usability and effectiveness of different recommendation algorithms within the HELPeR system, as judged by expert caregivers.

2.6. Participants

The study focused on assessing HELPeR from the perspective of expert caregivers. A total of **eleven** nurses from the School of Nursing were recruited for the study. These participants were selected based on their expertise and experience, with all participants possessing a minimum of **1 years** of nursing experience and all of them had care-giving experience. Additionally, each participant confirmed their understanding of ovarian cancer and their ability to provide care for ovarian cancer patients.

This study serves as a preliminary investigation, preceding a larger study involving patients and caregivers. The choice to use personas in the study was influenced by the fact that nurses typically

Table 5
Study Questionnaire

Question Group	Questions
Pre-Survey	<ol style="list-style-type: none"> 1. Age 2. Education level 3. Employment status 4. Relationship to patients (caregivers) 5. Professional Experience with general caregiving 6. Experience with oncology caregiving in years 7. Experience with ovarian cancer caregiving in years
Subtask 1 & 2: Document Feedback	<ol style="list-style-type: none"> 1. Relevance to Need 2. Relevance to Trajectory 3. Would you recommend this document?
Subtask 3: Document Selection for Patient	<ol style="list-style-type: none"> 1. Which articles would you give to the patient? 2. What specific qualities of the document led you to select it? 3. Agreement with the statement: "The document is easy for the patient to understand"
Overall Feedback	<ol style="list-style-type: none"> 1. Agreement with the statement: "This system would be beneficial for the care of my loved ones" 2. Agreement with the statement: "I would recommend this system to my loved ones"

Table 6
Participant Demographics and Experience

Category	Participants
Age Groups	
20-30 years	5
31-40 years	4
41+ years	2
Nursing Experience	
1-2 years	5
3-10 years	3
10+ years	3
Oncology Experience	
With Oncology Experience	8
Without Oncology Experience	3
Gyno-Oncology Experience	5

have limited information about a patient during a clinical visit, making the use of a persona a realistic representation of a real-world scenario.

To recruit participants, we employed a convenience sampling method due to the limited availability of qualified candidates. However, we ensured that the selected participants were representative of a diverse range of ages. Specifically, the group consisted of **eleven** participants aged between **20** and **44** years. The detailed statistics of the participants are provided in Table 6 Overall, this group of expert caregivers provided a representative sample for the study, allowing us to gather high quality gold-standard to evaluate recommender algorithms.

3. Experimental Results

3.1. Baseline Candidate Retrieval Performance

BM25 shows moderate performance, achieving a MAP of 0.4253 and an MRR of 0.2667 for patient need, but its effectiveness declines for patient trajectory (MAP: 0.2271, MRR: 0.2292) and expert recommendations (MAP: 0.1850, MRR: 0.1339). Among dense retrieval models, BioBERT outperforms BERT, especially in expert recommendation alignment (BioBERT MAP: 0.2289, MRR: 0.1339 vs. BERT MAP: 0.1736, MRR: 0.1250). However, both models fall short in trajectory relevance compared to sparse retrieval. As show in previous work [16], Fine-tuned models, ColBERT and MedCPT, lead in

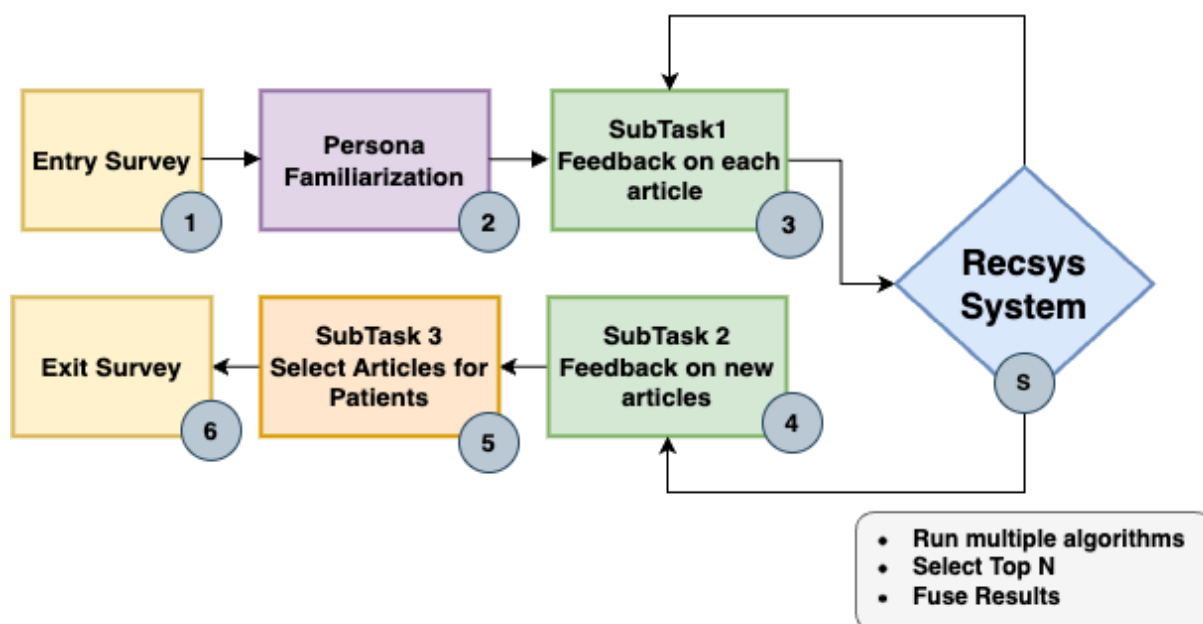


Figure 3: Flowchart of the User Study Design.

performance. MedCPT, in particular, excels across all criteria, with the highest MAP and MRR values (e.g., MAP: 0.5523 for need, 0.2815 for trajectory, and 0.4123 for expert recommendations), highlighting the benefits of fine-tuning for retrieval tasks.

The results demonstrate that fine-tuned dense retrieval models, especially MedCPT, offer superior accuracy in retrieving information relevant to patient needs, trajectories, and expert recommendations, surpassing both traditional sparse and standard dense retrieval approaches. For our further experiments, we will focus on BM25, ColBERT (general domain), and MedCPT (medical articles), as they are fine-tuned models and have demonstrated superior performance already.

3.2. Effect of integrating Patient Metadata

The performance of the re-ranking strategies across the BM25, ColBERT, and MedCPT models is summarized in the Table 8. MedCPT consistently demonstrates superior performance across all criteria, particularly in the combined profile and trajectory re-ranking approach, where it achieves a MAP of 0.7318 for Need, 0.6018 for Trajectory, and 0.6314 for Expert Recommendations. ColBERT also performs well, especially in the Trajectory-based re-ranking, indicating its effectiveness in integrating patient-specific trajectory information. In contrast, while BM25 shows reasonable performance, it falls behind the more advanced models, highlighting the significant benefits of incorporating both profile and trajectory data in enhancing retrieval relevance.

4. Implications

The use of personas, which are generic representations of patients at specific trajectory stages, offers a promising approach to addressing the cold-start problem in health recommender systems. These systems often struggle with limited patient data due to privacy concerns or patient reluctance to share information. Additionally, small population sizes in certain health contexts make collaborative filtering data scarce. By leveraging personas formed through representative patient stereotypes, health recommender systems can potentially provide personalized recommendations even in the absence of direct patient data, offering a viable solution for enhancing system effectiveness in these challenging scenarios.

Table 7

Performance of Candidate Retrieval Models on Biomedical Retrieval Tasks

Model	Need		Trajectory		Expert Rec.	
	MAP	MRR	MAP	MRR	MAP	MRR
Sparse Retrieval						
BM25	0.4253	0.2667	0.2271	0.2292	0.1850	0.1339
BERT	0.3820	0.2381	0.1426	0.1667	0.1736	0.1250
Dense Retrieval						
BioBERT	0.4195	0.2500	0.1926	0.2167	0.2289	0.1339
Dense Retrieval (fine-tuned for Retrieval)						
ColBERT	0.5476	0.2753	0.2458	0.2167	0.3173	0.2250
MedCPT	0.5523	0.2893	0.2815	0.2380	0.4123	0.2667

Table 8

Performance of Re-ranking based retrievals

Model	Need		Trajectory		Expert Rec.	
	MAP	MRR	MAP	MRR	MAP	MRR
Profile based re-ranking						
BM25	0.5371	0.3750	0.2678	0.2292	0.2645	0.1714
ColBERT	0.5397	0.2927	0.2458	0.2167	0.3863	0.2500
MedCPT	0.5611	0.2975	0.2381	0.2381	0.2381	0.2381
Trajectory based re-ranking						
BM25	0.6396	0.5018	0.5016	0.4917	0.4085	0.3518
ColBERT	0.6919	0.6676	0.5250	0.6636	0.5125	0.3321
MedCPT	0.7034	0.5327	0.5413	0.5191	0.6141	0.5163
Combine Profile and Trajectory re-ranking						
BM25	0.6431	0.5218	0.5127	0.5137	0.5251	0.3621
ColBERT	0.7092	0.6863	0.5750	0.6713	0.5365	0.3420
MedCPT	0.7318	0.6135	0.6018	0.5321	0.6314	0.5491

4.1. Limitations

While hybrid retrieval and fine-tuning approaches have shown potential for improving retrieval effectiveness [17], optimizing the retrieval pipeline was not the primary focus of this work. We utilized a basic zero-shot retrieval pipeline, which, while sufficient for our study’s objectives, may not fully leverage the potential of more advanced methods. In future work, we plan to explore and incorporate more sophisticated retrieval pipelines, including hybrid methods and fine-tuning techniques, to enhance the accuracy and relevance of health recommendations.

Recent advancements in large language models (LLMs) have demonstrated significant potential for enhancing personalization in recommender systems through zero-shot and few-shot prompting techniques [25]. These techniques have been particularly effective for tasks such as query expansion and user profile integration, allowing for more nuanced and context-aware recommendations [26, 27]. However, in this work, we did not explore these LLM-based approaches due to constraints on time and computational resources. While our current methods are effective for the study’s scope, integrating LLM techniques could potentially lead to more refined and personalized recommendations. In future research, we intend to investigate the application of LLMs to our recommendation framework, leveraging their capabilities to improve both query processing and user profile adaptation.

One of the primary limitations of this study is the relatively small sample size, with only 11 experts participating in the evaluation. This limited number of participants is insufficient to establish statistical significance for the results. Consequently, the findings should be interpreted with caution. However, it is important to note that this study serves as a preliminary investigation. Future research will involve a more comprehensive study with a larger sample size, including actual patients and caregivers. This expanded dataset will enable a more robust analysis and provide deeper insights into the effectiveness

and applicability of the proposed recommendation system.

5. Conclusion

In this study, we discussed and demonstrated the significance of integrating persona meta-data into health recommender systems. By leveraging personas—generic representations of patients at specific trajectory stages—we showed how these meta-data points can effectively address challenges in personalization, particularly in cold-start scenarios where patient data may be limited due to privacy concerns or patient reluctance. Our findings suggest that even minimal persona meta-data can enhance the relevance of recommendations, providing a strong foundation for further exploration in this area.

While we primarily focused on two specific meta-data points in this study, we recognize the potential of other persona attributes to further improve recommender systems. Our future research will explore how additional persona meta-data can be integrated to enrich recommendations and support more nuanced personalization.

Additionally, we collected qualitative feedback from nurse experts, which revealed a preference for selecting documents that are easy for patients to understand. This insight aligns with our ongoing research into knowledge adaptation [28], where we aim to tailor recommendations based on the patient's level of understanding and comprehension.

Acknowledgments

This work was supported by awards from the National Library of Medicine (NLM) of the National Institutes of Health (NIH) (Award Number: R01-LM013038). The content is solely the responsibility of the authors and NIH or NLM had no role in the design or conduct of the study; analysis, or interpretation of data; or preparation or review of the manuscript.

References

- [1] A. E. Anker, A. M. Reinhart, T. H. Feeley, Health information seeking: A review of measures and methods, *Patient Education and Counseling* 82 (2011) 346–354. URL: <https://www.sciencedirect.com/science/article/pii/S0738399110007470>. doi:<https://doi.org/10.1016/j.pec.2010.12.008>, methodology in Health Communication Research.
- [2] J. Susic, Nihseniorhealth classes for senior citizens at a public library in louisiana, *Journal of Consumer Health on the Internet* 13 (2009) 417–419.
- [3] A. O'Connor, D. Stacey, V. Entwistle, H. Llewellyn-Thomas, D. Rovner, M. Holmes-Rovner, V. Tait, J. Tetroe, V. Fiset, M. Barry, J. Jones, Decision aids for people facing health treatment or screening decisions., *The Cochrane database of systematic reviews* 2 (2003) CD001431.
- [4] E. L. Carter, G. Nunlee-Bland, C. Callender, A patient-centric, provider-assisted diabetes telehealth self-management intervention for urban minorities., *Perspectives in health information management* 8 (2011) 1b.
- [5] K. Lee, K. Hoti, J. Hughes, L. Emmerton, Consumer use of “dr google”: A survey on health information-seeking behaviors and navigational needs, *Journal of Medical Internet Research* 17 (2015).
- [6] K. Thaker, Y. Chi, S. Birkhoff, D. He, P. Brusilovsky, H. Donovan, Y. Lee, Exploring resource sharing behaviors for finding relevant health resources: An analysis of an online ovarian cancer community, *JMIR Cancer (Preprint) Accepted* 8 (2021) e33110.
- [7] S. Kanthawala, A. Vermeesch, B. Given, J. Huh, Answers to health questions: Internet search results versus online health community responses, *Journal of Medical Internet Research* 18 (2016).
- [8] A. C. Johnston, J. L. Worrell, P. Gangi, M. Wasko, Online health communities: An assessment of the influence of participation on patient empowerment outcomes, *Inf. Technol. People* 26 (2013) 213–235.

- [9] Jimmy, G. Zuccon, B. Koopman, Choices in knowledge-base retrieval for consumer health search, in: G. Pasi, B. Piwowarski, L. Azzopardi, A. Hanbury (Eds.), *Advances in Information Retrieval*, Springer International Publishing, Cham, 2018, pp. 72–85.
- [10] D. He, Z. Wang, K. Thaker, N. Zou, Translation and expansion: Enabling laypeople access to the covid-19 academic collection, *Data and Information Management* 4 (2020) 177–190.
- [11] R. D. Croon, L. V. Houdt, N. N. Htun, G. Štiglic, V. V. Abeele, K. Verbert, Health recommender systems: Systematic review, *Journal of Medical Internet Research* 23 (2021).
- [12] M. Etemadi, S. B. Abkenar, A. Ahmadzadeh, M. H. Kashani, P. Asghari, M. Akbari, E. Mahdipour, A systematic review of healthcare recommender systems: Open issues, challenges, and techniques, *Expert Syst. Appl.* 213 (2023) 118823.
- [13] Y. Wang, K. Thaker, V. Hui, P. Brusilovsky, D. He, H. Donovan, Y. J. Lee, Utilizing digital twin to create personas representing ovarian cancer patients and their families, in: *Innovation in Applied Nursing Informatics*, IOS Press, 2024, pp. 754–756.
- [14] K. Thaker, V. Hui, Z. Luo, Y. Wang, B. Rahadari, P. Brusilovsky, D. He, H. Donovan, Y. J. Lee, Walk a mile in their shoes: Using personas to obtain health professional perspectives for an ovarian cancer patient-focused recommender system, in: *Proceedings of the AMIA Annual Symposium*, AMIA, New Orleans, USA, 2023.
- [15] B. Rahdari, P. Brusilovsky, D. He, K. M. Thaker, Z. Luo, Y. J. Lee, Helper: An interactive recommender system for ovarian cancer patients and caregivers, in: *Proceedings of the 16th ACM Conference on Recommender Systems, RecSys '22*, Association for Computing Machinery, New York, NY, USA, 2022, p. 644–647.
- [16] Q. Jin, W. Kim, Q. Chen, D. C. Comeau, L. Yeganova, W. J. Wilbur, Z. Lu, MedCPT: Contrastive Pre-trained Transformers with large-scale PubMed search logs for zero-shot biomedical information retrieval, *Bioinformatics* 39 (2023).
- [17] N. Thakur, N. Reimers, A. Rücklé, A. Srivastava, I. Gurevych, BEIR: A heterogeneous benchmark for zero-shot evaluation of information retrieval models, in: *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [18] X. Jiao, Y. Yin, L. Shang, X. Jiang, X. Chen, L. Li, F. Wang, Q. Liu, TinyBERT: Distilling BERT for natural language understanding, in: T. Cohn, Y. He, Y. Liu (Eds.), *Findings of the Association for Computational Linguistics: EMNLP 2020*, Association for Computational Linguistics, Online, 2020, pp. 4163–4174. URL: <https://aclanthology.org/2020.findings-emnlp.372>. doi:10.18653/v1/2020.findings-emnlp.372.
- [19] J. Lee, W. Yoon, S. Kim, D. Kim, S. Kim, C. H. So, J. Kang, Biobert: a pre-trained biomedical language representation model for biomedical text mining, *Bioinformatics (Oxford, England)* 36 (2020) 1234–1240.
- [20] K. Santhanam, O. Khattab, J. Saad-Falcon, C. Potts, M. Zaharia, ColBERTv2: Effective and efficient retrieval via lightweight late interaction, in: M. Carpuat, M.-C. de Marneffe, I. V. Meza Ruiz (Eds.), *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Association for Computational Linguistics, Seattle, United States, 2022, pp. 3715–3734. URL: <https://aclanthology.org/2022.naacl-main.272>. doi:10.18653/v1/2022.naacl-main.272.
- [21] J. Lin, X. Ma, S.-C. Lin, J.-H. Yang, R. Pradeep, R. Nogueira, Pyserini: A python toolkit for reproducible information retrieval research with sparse and dense representations, in: *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21*, Association for Computing Machinery, New York, NY, USA, 2021, p. 2356–2362.
- [22] J. Johnson, M. Douze, H. Jégou, Billion-scale similarity search with GPUs, *IEEE Transactions on Big Data* 7 (2019) 535–547.
- [23] Y. Yue, R. Patel, H. Roehrig, Beyond position bias: examining result attractiveness as a source of presentation bias in clickthrough data, in: *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, Association for Computing Machinery, New York, NY, USA, 2010, p. 1011–1018.
- [24] M. Prince, 9 - epidemiology, in: P. Wright, J. Stern, M. Phelan (Eds.), *Core Psychiatry (Third*

Edition), third edition ed., W.B. Saunders, Oxford, 2012, pp. 115–129.

- [25] P. Liu, L. Zhang, J. A. Gulla, Pre-train, prompt, and recommendation: A comprehensive survey of language modeling paradigm adaptations in recommender systems, *Transactions of the Association for Computational Linguistics* 11 (2023) 1553–1571.
- [26] A. Acharya, B. Singh, N. Onoe, Llm based generation of item-description for recommendation system, in: *Proceedings of the 17th ACM Conference on Recommender Systems, RecSys '23*, Association for Computing Machinery, New York, NY, USA, 2023, p. 1204–1207.
- [27] Y. Wu, R. Xie, Y. Zhu, F. Zhuang, X. Zhang, L. Lin, Q. He, Personalized prompt for sequential recommendation, *IEEE Transactions on Knowledge and Data Engineering* (2024).
- [28] K. Thaker, Ka-recsys: Knowledge appropriate patient focused recommendation technologies, in: *Proceedings of the 16th ACM Conference on Recommender Systems, RecSys '22*, Association for Computing Machinery, New York, NY, USA, 2022, p. 720–721. URL: <https://doi.org/10.1145/3523227.3547422>. doi:10.1145/3523227.3547422.