An Exploration On-demand Article Recommender System for Cancer Patients Information Provisioning

M. Mehdi Afsar, Trafford Crump, and Behrouz Far
University of Calgary
Calgary, AB, Canada
{mehdi.afsar, tcrump, far}@ucalgary.ca

Abstract

Information provision plays an important role in educating patients with serious illnesses, like cancer, to cope with their disease conditions and to actively participate in shared-decision making process. Recent studies suggest that there is a lack of appropriate educational resources for such patients, specifically prostate cancer patients. To address this issue, in this paper, a Knowledge-based Exploration on-demand article Recommender System (called KERS) is proposed that can provide evidence-based information for patients. Recognizing the fact that exploration is expensive when the user of the system is a human, the main idea in KERS is to minimize exploration while achieving the maximum long-term satisfaction. Therefore, using a knowledge-base developed by an expert in the field, KERS learns user interests as quickly as possible and then it exploits this knowledge to recommend the best articles. Furthermore, KERS needs no information from users beforehand and it learns them through interacting with users. The system will help patients make informed decisions, and at the same time, will reduce the burden on the healthcare providers. The results of experiments have confirmed the effectiveness of the proposed system compared to baseline methods.

1 Introduction

Prostate cancer is the most commonly diagnosed cancer among Canadian men; one in nine men in Canada are diagnosed with prostate cancer in their lifetime[1]. In 2020, there will be an estimated 23,300 Canadian men diagnosed with prostate cancer and 4,200 deaths due to this disease[2]. Coping with prostate cancer can be particularly difficult because the optimal treatment is not clear and treating the disease can have consequences for a man's quality of life (Middleton et al. 1995; Denmeade and Isaacs 2002), including urinary incontinence, sexual dysfunction, depression, and anxiety (Bowler et al. 2019). Accordingly, providing prostate cancer patients with appropriate educational resources can play an important role in treatment decision making. Nonetheless, there is a deficiency in the availability of high-quality educational resources for prostate cancer patients compared to other cancer patients (Kassianos, Raats, and Gage 2016; Bowler et al. 2019).

In cancer care, the provision of information has four purposes, to: 1) prepare patients for their treatment, 2) increase adherence to therapy, 3) increase their abilities to cope with the illness, and 4) to promote recovery (Husson, Mols, and Van de Poll-Franse 2011). Useful information is defined by providing patients with the kind of information they need to know, at the time in their disease progression that they need to know it, by the right source (Rutten et al. 2005). The benefits of information for cancer patients include increased satisfaction and involvement in decision making, improved ability in coping with stressful stages of disease, anxiety control, and enhanced communication with family members and clinicians (Rutten et al. 2005).

There is evidence that the satisfaction of prostate cancer patients has been improved with simple self-management educational resources, in the form of booklets (Bowler et al. 2019). However, many of these resources have been found to be too simple/impractical (Balakrishnan et al. 2019), unreliable (Alsyouf et al. 2019; Dee and Lee 2019), incomprehensible (Kim et al. 2019), and not personalized (Bowler et al. 2019). Therefore, it seems with more sophisticated educational resources, there is a chance to better improve the satisfaction, and perhaps the quality of life, of these patients. We hypothesize that a recommender system (RS) could be designed to provide prostate cancer patients with useful information as a basis for improving their educational resources.

RSs are software tools that help users find items of their interest. There are two main paradigms toward RSs: collaborative filtering (CF) and content-based filtering (CBF). In CF, which is the foundation of initial RSs (Ricci, Rokach, and Shapira 2011), the idea is to first find users similar to the main user, then recommend items liked by these similar users to the user. In CBF, items similar to the history of the user are recommended to the user. However, these methods are not effective when there is no database with enough user ratings. Knowledge-base RSs (KBRSs) are a third paradigm in which recommendations are generated using a knowledge-base (KB) developed by an expert in the domain. If appropriately developed, KBRSs can cover the problems of the two aforementioned methods, including new user/item (*cold start*), unreliable recommendations, and

[1]Prostate Cancer Canada

[2]Canadian Cancer Society

serendipity (Bouraga et al. 2014).

In this paper, we aim at tackling the information provisioning problem for prostate cancer patients through proposing a Knowledge-based, Exploration on-demand article Recommender System, called KERS. Because exploration is costly when recommending to human users, the main idea in KERS is to minimize the exploration while achieving the maximum satisfaction. To do this, we incorporate a new exploration vs. exploitation strategy into the KBRS technology. More specifically, KERS is composed of two phases: Exploration and Exploitation. User interests are learned through Exploration phase with the help of a KB. Then, this knowledge is used to recommend the best articles to the user in the Exploitation phase. The performance of KERS is validated through extensive experiments. In general, the contribution of this paper is threefold: 1) KERS is the first RS designed to cover the information needs of prostate cancer patients, 2) to our knowledge, this hybrid of KB and learning ability (balancing exploration vs. exploitation) is new, and 3) we present a simple, flexible user simulator to evaluate our method offline through a simulation study.

The remaining of this paper is organized as follows. Section 2 provides necessary background for this paper. Section 3 defines the problem and explains notations used. The proposed KERS is described in Section 4. The results of experiments are presented in Section 5 and the paper is concluded in Section 6.

2 Background

Since both KBRSs and bandits inspired us in developing KERS, we provide a quick background for them below.

Knowledge-based Recommender Systems (KBRSs)

KBRSs use the knowledge provided by a human expert to generate recommendations. Instead of using user ratings to figure out the taste of the user, they rely on deep knowledge about a topic and explicit user requirements to come up with the best recommendations (Felfernig et al. 2011). Thus, this type of RSs is a good match for applications where user ratings on items is scarce. Methods developed for KBRSs can be generally divided into *case-based* and *constraint-based* approaches (Felfernig et al. 2011). The idea in both methods is the same; the user poses his problem or user requirements are collected, repairs are handled when there are discrepancies, and recommendations are made and explained (Bouraga et al. 2014; Felfernig et al. 2011). The difference is in the way how these recommendations are generated; while case-based KBRSs use a similarity metric to match user's requirements with items in the KB (Burke 2000; Khan and Hoffmann 2003; Lee and Lee 2007; Chattopadhyay et al. 2013; Rosa et al. 2018), constraint-based methods rely on pre-defined and strong rules to match these two (Tsang 1993; Towle and Quinn 2000; Felfernig et al. 2011).

Bandits

The problem of balancing the exploration vs. exploitation is typically formulated as a multi-armed bandit (MAB) problem. In a MAB problem, the agent faces with a choice among k options and through selecting one choice, it receives a numerical reward (Sutton and Barto 2018). This scenario is similar to a situation where a gambler should decide to play which lever of a slot machine to gain more payoffs in a series of trials. Algorithms developed to solve the MAB problem can be generally divided into *context-free* and *contextual* bandit algorithms (Sutton and Barto 2018; Li et al. 2010). In the former, the agent knows nothing about the environment (items and users). On the other hand, in contextual bandits, the agent sees a feature vector of the environment alongside the history of each arm in order to select the best arm. Since we are inspired by context-free bandits in developing KERS, we provide more details about this type of bandits. ϵ-greedy is one of the simplest context-free bandits in which either the best arm with probability $1-\epsilon$ or a random arm with probability ϵ is selected (Sutton and Barto 2018). To address the unguided exploration by ϵ-greedy, upper confidence bound (UCB) (Auer, Cesa-Bianchi, and Fischer 2002) algorithms follow optimism in the face of uncertainty. More specifically, these algorithms usually compute a confidence bound for each arm (action) and then select the arm with the largest bound. In other words, the less confident they are about an arm, the more likely they select it. From this family, UCB1 algorithm first plays each arm once. Then, at time step t, it selects the arm that maximizes $Q_t(a) + \sqrt{2\log t/n_a}$, where $Q_t(a)$ is the value of arm a at time step t and n_a is the number of times that arm a has been played. We compare the performance of KERS with ϵ-greedy and UCB1 in Section 5.

In contrast, KERS is a hybrid method of KBRSs and MABs. On the one hand, it relies on the strength of KBs in terms of carefully classified information. On the other hand, it wisely modifies the unguided exploration in MABs and equips the KBRS with a new learning algorithm.

3 Problem Definition

We assume that KERS runs in a discrete time space, called *time steps*, $t = 1, 2, 3, ...T$, where T is a finite number. There are M users and N items (articles in our case) in the system and the problem is to recommend, at time step t, the best item to the *active user*. In every time step, the user is provided with k articles and the system receives the feedback from the user. Because articles in our KB ($\mathcal{K}$) are carefully developed and every one of them is about a specific topic, it is assumed that the title of articles is informative enough to convey the intent of every article. Thus, to not overwhelm the users, the system only shows the title of articles, which is no more than few words. On receiving every recommendation, the user provides a numerical *reward* (r). In this paper, r is considered to be the click of the user on the title of an article, which is equivalent to a numerical reward of 1 if clicked and 0 otherwise. That said, denoting a as an arm and A as the set of arms, the objective is

$$\max R_A(T), \tag{1}$$

where $R_A(T) = \sum_{t=1}^{T}\sum_{i=1}^{k} r_{t,a_i}$. That means, the objective at t is to recommend the best k arms that maximize the cumulative reward during $t = 1, ..., T$. In order to be consistent with bandit field's terminology, *arm* and *article* terms

Table 1: Notation Description

Notation	Meaning
$\mathcal{P}$	User Profile unit
$\mathcal{K}$	Knowledge-base
$\mathcal{I}_u$	User interests
N	Number of items
M	Number of users
T	Number of time steps
k	Number of articles to recommend in a single time step
C	Number of categories in the knowledge-base
r	Numerical reward
R	Cumulative reward in T time steps
a	An arm
A	The set of arms
η	The number of possible user interests
δ	The epoch for change in the user interest

are used interchangeably in this paper. Table 1 demonstrates the notation used in this paper.

4 Methodology

Rationale

In this subsection, we explain the rationale behind our work. In the MAB field, the agent's goal is to maximize the total reward in the long run so it needs exploration besides exploitation (Sutton and Barto 2018). It is undeniable that when the user of a system is a human, the exploration is very costly (Chen et al. 2019). This is mainly because the human is complex in nature and can become bored so quickly when recommendations are irrelevant. Therefore, our first and foremost goal is to use a method that needs minimum exploration, or better say, that explores whenever necessary while achieving the maximum satisfaction in the long-term. Secondly, it is clear that if the users *know* and *say* what specifically is of their interest, then the recommendation task becomes very straightforward. However, in many cases, the users are not clear about their interests (Beel and Langer 2015). Moreover, the research indicates that the users are too lazy to provide necessary feedback on recommendations received (Bai et al. 2019). Accordingly, we believe that a good recommendation method should either need no information from the users or be able to work with *implicit feedback*. Finally, in the healthcare field, it is of major importance to recommend accurate and reliable information to patients (Alsyouf et al. 2019; Dee and Lee 2019). Otherwise, it can have costly consequences. Typically, the best way to recommend the most reliable information to the users is to use KBs. Regarding these facts, we have proposed KERS — a hybrid method that utilizes an on-demand exploration scheme, described below.

Proposed KERS

Fig. 1 illustrates the architecture of KERS, which is composed of two phases: Exploration and Exploitation. In the Exploration phase, since the idea is to find user interests (i.e., topics or categories the user is interested to know about and indicated by $\mathcal{I}_u$), KERS recommends one random article from each category. When the user clicks on an article's title, the category of that article is added to $\mathcal{I}_u$. User Profile ($\mathcal{P}$)

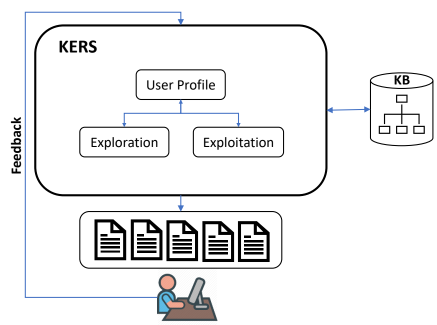

Figure 1: The proposed system architecture

unit (see Fig. 1) is responsible to keep the track of users and their interests by saving their IDs and $\mathcal{I}_u$. It is noteworthy to mention that if C is larger than k, Exploration might take a couple of time steps to complete. For instance, if $k = 5$ and $C = 8$, Exploration takes two time steps, e.g., t_1 and t_2, to complete. At t_1, five articles from first five categories are recommended. At t_2, three articles from the remaining three categories plus two articles from random categories are recommended (i.e., these two articles could be from any of eight categories). Note that, at t_2, we could only recommend three articles from the remaining three categories; however, we aim to preserve the consistency and recommend exactly k articles at each time step. In the Exploitation phase, KERS uses $\mathcal{P}$ and randomly recommends k articles from $\mathcal{I}_u$. Note that if $|\mathcal{I}_u| > 1$ for a user, KERS tries to fairly recommend k articles from all these categories. More specifically, if a user is interested in η topics, KERS recommends k/η articles from C_1, k/η from C_2, and so on. If k is not divisible by η, KERS uses a *rounding up* method; e.g., if $k = 5$, $\eta = 2$, and $C = 5$, it recommends three articles from C_1 and two articles from C_2. Overall, in the first time step(s), KERS is in the Exploration phase and recommends articles from all categories. Then, when Exploration completes, it switches to the Exploitation phase and remains there until $\mathcal{I}_u$ changes. KERS detects this change through receiving no reward in a time step. Once occurred, KERS switches back to Exploration to find out new $\mathcal{I}_u$. This process is repeated until time step T or when the user quits the system. The pseudo code of KERS is presented in Algorithm 1.

5 Experiments

In this section, the performance of KERS is validated through simulation study. First the experiments setup is described. Then, the results of experiments are presented.

Setup

Parameters and Baselines In order to investigate the effect of different values of M on the performance of the system, we have picked the following values: 1, 10, and 100. If $M > 1$, in every time step, one user is randomly selected as the *active user* by the system and he receives the recommendation. There are 10,000 time steps (T) and every experiment is the average over 20 runs. We pick $k = 5$ to avoid *choice overload* (Beierle et al. 2019) and use R as the performance metric. For performance comparison, ϵ-greedy

Algorithm 1: KERS algorithm

initialization ($t = 1, \mathcal{I}_u = \emptyset$, *phase* = Exploration, *exp_ctr = 0*)
while $t <= T$ **do**
  if *phase* = Exploration **then**
    recommend $\{a_1, \ldots, a_k\}$ from $\mathcal{K}$
    if $\exists i \in \{1, \ldots, k\}$, r_{a_i} = 1 **then**
      add C_{a_i} to $\mathcal{I}_u$
      if *user_id* $\notin \mathcal{P}$ **then**
        add *user_id* and $\mathcal{I}_u$ to $\mathcal{P}$
      end
    end
    exp_ctr += k
    if $\mathcal{I}_u \neq \emptyset$ **and** *exp_ctr* $> C$ **then**
      phase = Exploitation
      exp_ctr = 0
    end
  else
    if *user_id* $\notin \mathcal{P}$ **then**
      phase = Exploration
    else
      recommend $\{a_1, \ldots, a_k\}$ from $\mathcal{K}$ according to $\mathcal{I}_u$
      if $\nexists i \in \{1, \ldots, k\}$, r_{a_i} = 1 **then**
        $\mathcal{I}_u = \emptyset$
        phase = Exploration
      end
    end
  end
  $t = t + 1$
end

and UCB1 are implemented as the baseline algorithms. Although, for the sake of clarity, we only report the results for ϵ-greedy when $\epsilon = 0.1$, KERS outperforms ϵ-greedy irrespective of the value of ϵ.

Datasets Two datasets are used in our experiments: prostate cancer and BBC (Greene and Cunningham 2006). We have developed a dataset with 500 articles ($N = 500$) about prostate cancer, which are crawled from prostate cancer websites[3]. The reliability of the content of these websites have been validated by our experts and the level of information is understandable to a layperson. The articles are categorized into five groups: *pre-diagnosis, diagnosis, treatment, side effects,* and *recurrence issues*. Our previous work (Baverstock, Crump, and Carlson 2015) inspired us for this classification and, with the advice from our experts, we extended it to five categories, which better reflects the current informational needs of prostate cancer patients observed in clinical communications. To see the performance on a larger dataset, we have also used BBC dataset in our experiments. The articles in this dataset are categorized into five topics, i.e., *business, entertainment, politics, sport,* and *technology,* and there are 2,225 articles in the dataset (N =2,225).

User Simulator According to (Zheng et al. 2018), the user's interest is dynamic and changes over time. This change is sporadic and depends on many factors, including personality, mood, and time. Accordingly, it is almost infeasible to exactly model a user. Instead, in this paper, we present a simple user simulator and use it to evaluate our method. Basically, the user interest has two main elements: the number of topics a user might be interested to know about (η) and the time or duration the user is interested about them (δ). For example, a patient may be interested in the topics of diagnosis and treatment, as well as he might become bored after some time (e.g., 10 minutes) and leave the platform. Upon returning to use the platform again, he might become interested in side effects and recurrence issues topics. For simplicity, we assume that these two factors are fixed in our users. In other words, if $\eta = 2$ and $\delta = 20$ time steps for a user, he stays interested in two topics until $t = T$, but he *might* change these two topics every 20 time steps. We let $\eta \in [1, 2, 3]$ and $\delta \in [10, 20, \ \ldots \ , 100]$. That means, one user may change his interest every 10 time steps and another one every 20 time steps. With this setting, a user can stay focused on merely one topic all the time. The system assigns η and δ to the users randomly. Note that users can have similar behaviors; it is quite possible that two users have similar η and δ.

Results

Prostate cancer dataset Figs. 2(a) to 2(i) show R for the three algorithms when the values of M and η vary from 1 to 100 and 1 to 3, respectively. In general, KERS outperforms the two methods in all settings. This superiority is more tangible when M and η are smaller. More specifically, when $M = 10$ and $\eta = 1$, KERS outperforms the baseline methods by 100% (Fig. 2(d)). The reason of this superiority is that KERS finds $\mathcal{I}_u$ as quickly as possible and exploits this knowledge in subsequent time steps. KERS returns to the Exploration phase only if it sees a change in $\mathcal{I}_u$. While the baseline methods perform greedily in most of the time, their exploration is not wise enough. UCB1 performs relatively better than ϵ-greedy when η is smaller (see Figs. 2(a) and 2(b)). As M increases, predicting $\mathcal{I}_u$ becomes more difficult and the performance of the three methods diminishes accordingly. Whereas KERS achieves almost the perfect R (i.e., 50,000 after 10,000 time steps) when $M = 1$ and $\eta \in [1, 2, 3]$, this number becomes about 15,000 when $M = 100$ (Figs. 2(g), 2(h), and 2(i)). This is mainly because when there are more users in the system and their interest changes after a while, more exploration is required to find the current interest of users and exploitation becomes ineffective. Another observation is that when η increases, R is better in all methods. This makes sense as when the users are interested in multiple topics simultaneously, it is easier to predict their interests compared to when they are interested in only one topic.

BBC dataset Figs. 2(j) to 2(r) show R for the three algorithms applied to the BBC dataset. Again, KERS outperforms the baseline methods in all settings and the same observations about change in M and η and their effect on the performance of the algorithms are seen here. The only difference is that UCB1 performs the worst among the three algorithms. The reason is that this method should explore

[3]Prostate Cancer Canada, Canadian Cancer Society, Canadian Urological Association, American Urological Association, and Wikipedia

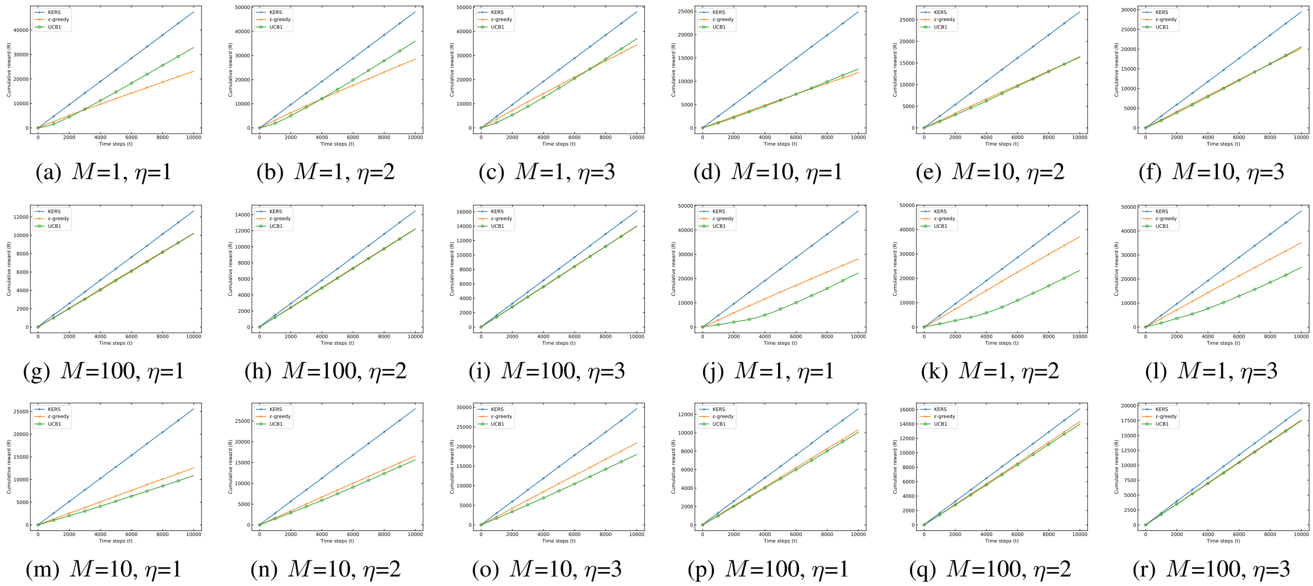

(a) M=1, η=1 (b) M=1, η=2 (c) M=1, η=3 (d) M=10, η=1 (e) M=10, η=2 (f) M=10, η=3

(g) M=100, η=1 (h) M=100, η=2 (i) M=100, η=3 (j) M=1, η=1 (k) M=1, η=2 (l) M=1, η=3

(m) M=10, η=1 (n) M=10, η=2 (o) M=10, η=3 (p) M=100, η=1 (q) M=100, η=2 (r) M=100, η=3

Figure 2: The cumulative reward for the three algorithms on prostate cancer dataset ((a)-(i)) and BBC dataset ((j)-(r))

and play all arms first and then starts to balance exploration vs. exploitation. Because there are more arms to explore in this dataset and users interests also change over time, UCB1 performs the worst.

Finally, Fig. 3 illustrates the number of times each arm in the BBC dataset is played by each algorithm. To be clear, we have shown the results when $M = 1$, $\delta = T$ (i.e, the user interest does not change over time), and $\eta = 1$ (the user is interested in the topic of technology). As depicted, both KERS and UCB1 fairly recommend all articles in this topic. On the other hand, ϵ-greedy performs quite poorly in this term as it greedily recommends a few articles all the time. For example, it recommends one article in this topic more than 9,000 times. This clearly shows while ϵ-greedy performs relatively better in this dataset in terms of R compared to UCB1, it is almost impractical for recommending articles to a human.

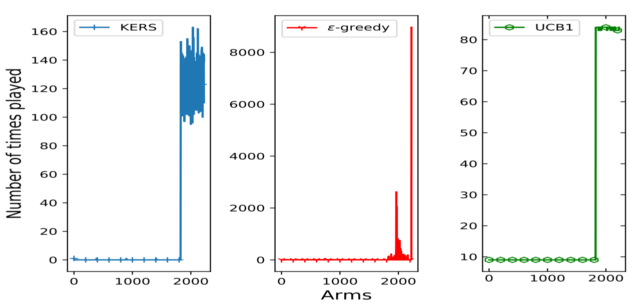

Figure 3: The number of times each arm played by algorithms ($M = 1$ and $\eta = 1$)

Discussion

KERS borrows the advantages of the two methods, i.e., KBRSs and MABs, as well as it covers their problems. In other words, while it is simple and practical, it alleviates the exploration vs. exploitation problem. At the same time, it is content (articles' body text) and dataset size independent and can provide reliable recommendations. However, as all knowledge-based systems, one might argue that the performance of KERS depends on the accuracy of the KB used and its performance degrades if the KB is not carefully developed. This dependence on the KB may also have a detrimental effect on the exploration time in case the number of categories in the KB is very large, although it is rare.

Patients diagnosed with prostate cancer may be asked by their clinician to share in decisions regarding treatment, given the inherent trade-offs across the treatment options associated with adverse quality of life outcomes. To make informed decisions, these patients may have informational needs that cannot be fulfilled with conventional resources. The personalization supplied by KERS could be used to provide prostate cancer patients with more useful information compared to existing educational resources (e.g., patient pamphlet). This, in turn, could help better inform patients about prostate cancer, the treatments that are available to them, and better inform their expectations over the duration of their care pathway.

6 Conclusion and Future Work

In this paper, we have proposed KERS — a knowledge-based exploration on-demand RS algorithm for cancer patients information provisioning. Since exploration is expensive when the user of the system is a human, the main objective of KERS is to achieve the maximum long-term satisfaction through minimum exploration. The results of experiments have confirmed the effectiveness of KERS compared to baseline algorithms.

This paper is a pilot study to initially test the performance of KERS offline and in a small scale. In the future, we will investigate the performance of KERS in an online study, i.e., using real prostate cancer patients, with a large dataset and various numbers of categories.

Acknowledgements

This paper was partially supported by NSERC (Natural Sciences and Engineering Research Council of Canada).

References

Alsyouf, M.; Stokes, P.; Hur, D.; Amasyali, A.; Ruckle, H.; and Hu, B. 2019. 'fake news' in urology: evaluating the accuracy of articles shared on social media in genitourinary malignancies. *BJU international* 124(4):701–706.

Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47(2-3):235–256.

Bai, X.; Wang, M.; Lee, I.; Yang, Z.; Kong, X.; and Xia, F. 2019. Scientific paper recommendation: A survey. *IEEE Access* 7:9324–9339.

Balakrishnan, A. S.; Nguyen, H. G.; Shinohara, K.; Yeung, R. A.; Carroll, P. R.; and Odisho, A. Y. 2019. A mobile health intervention for prostate biopsy patients reduces appointment cancellations: Cohort study. *Journal of medical Internet research* 21(6):e14094.

Baverstock, R. J.; Crump, R. T.; and Carlson, K. V. 2015. Patient educational technologies and their use by patients diagnosed with localized prostate cancer. *BMC health services research* 15(1):433.

Beel, J., and Langer, S. 2015. A comparison of offline evaluations, online evaluations, and user studies in the context of research-paper recommender systems. In *International conference on theory and practice of digital libraries*, 153–168.

Beierle, F.; Aizawa, A.; Collins, A.; and Beel, J. 2019. Choice overload and recommendation effectiveness in related-article recommendations. *International Journal on Digital Libraries* 1–16.

Bouraga, S.; Jureta, I.; Faulkner, S.; and Herssens, C. 2014. Knowledge-based recommendation systems: a survey. *International Journal of Intelligent Information Technologies (IJIIT)* 10(2):1–19.

Bowler, M.; Dehek, R.; Thomas, E.; Ngo, K.; and Grose, L. 2019. Evaluating the impact of post-treatment self-management guidelines for prostate cancer survivors. *Journal of medical imaging and radiation sciences* 50(3):398–407.

Burke, R. 2000. Knowledge-based recommender systems. *Encyclopedia of library and information systems* 69(Supplement 32):175–186.

Chattopadhyay, S.; Banerjee, S.; Rabhi, F. A.; and Acharya, U. R. 2013. A case-based reasoning system for complex medical diagnosis. *Expert Systems* 30(1):12–20.

Chen, M.; Beutel, A.; Covington, P.; Jain, S.; Belletti, F.; and Chi, E. H. 2019. Top-k off-policy correction for a reinforce recommender system. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, 456–464.

Dee, E. C., and Lee, G. 2019. Adverse effects of radiotherapy and chemotherapy for common malignancies: what is the quality of information patients are finding online? *Journal of Cancer Education* 1–11.

Denmeade, S. R., and Isaacs, J. T. 2002. A history of prostate cancer treatment. *Nature Reviews Cancer* 2(5):389.

Felfernig, A.; Friedrich, G.; Jannach, D.; and Zanker, M. 2011. Developing constraint-based recommenders. In *Recommender systems handbook*. 187–215.

Greene, D., and Cunningham, P. 2006. Practical solutions to the problem of diagonal dominance in kernel document clustering. In *Proceedings of the 23rd international conference on Machine learning*, 377–384.

Husson, O.; Mols, F.; and Van de Poll-Franse, L. 2011. The relation between information provision and health-related quality of life, anxiety and depression among cancer survivors: a systematic review. *Annals of Oncology* 22(4):761–772.

Kassianos, A. P.; Raats, M. M.; and Gage, H. 2016. An exploratory study on the information needs of prostate cancer patients and their partners. *Health psychology research* 4(1).

Khan, A. S., and Hoffmann, A. 2003. Building a case-based diet recommendation system without a knowledge engineer. *Artificial Intelligence in Medicine* 27(2):155–179.

Kim, C.; Prabhu, A. V.; Hansberry, D. R.; Agarwal, N.; Heron, D. E.; and Beriwal, S. 2019. Digital era of mobile communications and smartphones: A novel analysis of patient comprehension of cancer-related information available through mobile applications. *Cancer investigation* 37(3):127–133.

Lee, J. S., and Lee, J. C. 2007. Context awareness by case-based reasoning in a music recommendation system. In *International symposium on ubiquitious computing systems*, 45–58.

Li, L.; Chu, W.; Langford, J.; and Schapire, R. E. 2010. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, 661–670.

Middleton, R. G.; Thompson, I. M.; Austenfeld, M. S.; Cooner, W. H.; Correa, R. J.; Gibbons, R. P.; Miller, H. C.; Oesterling, J. E.; Resnick, M. I.; Smalley, S. R.; et al. 1995. Prostate cancer clinical guidelines panel summary report on the management of clinically localized prostate cancer. *The Journal of urology* 154(6):2144–2148.

Ricci, F.; Rokach, L.; and Shapira, B. 2011. Introduction to recommender systems handbook. In *Recommender systems handbook*. 1–35.

Rosa, R. L.; Schwartz, G. M.; Ruggiero, W. V.; and Rodríguez, D. Z. 2018. A knowledge-based recommendation system that includes sentiment analysis and deep learning. *IEEE Transactions on Industrial Informatics* 15(4):2124–2135.

Rutten, L. J. F.; Arora, N. K.; Bakos, A. D.; Aziz, N.; and Rowland, J. 2005. Information needs and sources of information among cancer patients: a systematic review of research (1980–2003). *Patient education and counseling* 57(3):250–261.

Sutton, R. S., and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.

Towle, B., and Quinn, C. 2000. Knowledge based recommender systems using explicit user models. In *Proceedings of the AAAI Workshop on Knowledge-Based Electronic Markets*, 74–77.

Tsang, E. 1993. *Foundations of constraint satisfaction*. Academic Press, London and San Diego.

Zheng, G.; Zhang, F.; Zheng, Z.; Xiang, Y.; Yuan, N. J.; Xie, X.; and Li, Z. 2018. Drn: A deep reinforcement learning framework for news recommendation. In *Proceedings of the 2018 World Wide Web Conference*, 167–176.