

Navigation Interface for Recommending Home Medical Products

Gang Luo

IBM T.J. Watson Research Center, 19 Skyline Drive, Hawthorne, NY 10532, USA

luog@us.ibm.com

Abstract Based on users' health issues, an intelligent personal health record (iPHR) system can automatically recommend home medical products (HMPs) and display them in a sequential order. However, the sequential output interface does not categorize search results and is not easy for users to quickly navigate to their desired HMPs. To address this problem, we developed a navigation interface for retrieved HMPs. Our idea is to use medical knowledge and nursing knowledge to construct a navigation hierarchy based on product categories. This hierarchy is added to the left side of each search result Web page to help users move through retrieved HMPs. We demonstrate the effectiveness of our techniques using USMLE medical exam cases.

Keywords Search engine · Personal health record · Home medical product · Nursing knowledge · Navigation interface

1. Introduction

About half of Americans have chronic conditions and need home medical products (HMPs) to facilitate their daily activities of living [13, 14]. However, most of these people encounter difficulty in finding HMPs that can help them. Their physicians receive little training on HMPs. Due to a lack of medical knowledge, the average person cannot come up with appropriate keywords to search HMP catalogs effectively [7].

To address this issue, we recently developed an intelligent personal health record (iPHR) system that can automatically recommend HMPs based on users' health issues [7, 8]. This personalized healthcare information can facilitate people's daily activities of living. Our key idea is to use medical knowledge and nursing knowledge to automatically generate keyword queries that are submitted to a search engine. Using expert system technology and Web search technology, iPHR recommends HMPs in the following way.

For each health issue (e.g., disease, symptom, surgery), iPHR stores a list of "search guide" phrases pre-compiled using disease/symptom treatment knowledge and nursing knowledge. These phrases serve to bridge the semantic gap between the literal meaning and the underlying medical meaning of the health issue. For example, the phrase *grab*

bar is pre-compiled for *muscular dystrophy* because grab bar can reduce a muscular dystrophy patient's risk of falling. For a user concerned with this health issue, iPHR uses a search engine and each of the pre-compiled search guide phrases as a query to retrieve some relevant HMPs. Then iPHR combines all retrieved HMPs together, diversifies search results, and returns them to the user.

By default, iPHR displays its recommended HMPs in a sequential order traditionally used by search engines. Initially when the user has no idea what HMPs are available on the market, this sequential output interface can let her quickly reach all kinds of HMPs that potentially can help her. Often after little browsing, she will form a rough idea of some preferred categories (e.g., eating aids) that her desired HMPs should fall into. In practice, some users can have such ideas even before they start to use iPHR, e.g., if they have a particular urgent need. In this case, the sequential output interface is no longer optimal. It mixes a large number of product categories and does not focus on HMPs in a particular category. Hence, the user cannot use it to easily find desired HMPs in preferred categories.

- < Health & Personal Care
 - < Health Care
 - < Medical Supplies & Equipment
 - < Daily Living Aids
 - < **Eating & Drinking Aids**
 - * Dinnerware (43)
 - * Cups & Glasses (34)

Fig. 1 An example product navigation hierarchy used on Amazon.com.

This paper presents a HMP navigation interface to address this problem. In the navigation interface, a navigation hierarchy based on product categories is displayed on the left side whereas recommended HMPs are displayed sequentially on the right side. The number of HMPs contained in each product category is also supplied. As shown in Fig. 1, such kind of navigation interface has been widely used in online product catalogs. However, existing product category hierarchy is static and manually constructed. It does not provide a good coverage of HMPs

recommended by iPHR because only a small number of categories are available at each level of the hierarchy. Users would easily miss the HMPs omitted by the navigation hierarchy. In our approach, all recommended HMPs are retrieved using treatment knowledge and nursing knowledge. We can use the same knowledge to dynamically construct a product category hierarchy that has better coverage of recommended HMPs than the existing one. To improve users' search experience, we also sort categories in the navigation hierarchy in such a way that important, information-rich, and diverse categories are ranked higher.

We implemented our techniques in a prototype iPHR system and evaluated their effectiveness using USMLE medical exam cases [6]. Our experiments show that the HMP navigation interface significantly improved users' search experience and helped them reach desired HMPs more quickly.

In related work, traditional recommender systems recommend items based on various factors such as users' prior ratings, previous purchases, and profiles [17]. Our iPHR can be regarded as a recommender system that goes beyond these factors and recommends personalized healthcare information using user profile, medical knowledge, and nursing knowledge.

The rest of the paper is organized as follows. Section 2 provides an overview of iPHR. Section 3 presents our navigation interface for retrieved HMPs. Section 4 evaluates our navigation interface. Section 5 concludes this paper.

2. Overview of iPHR

In this section, we provide an overview of iPHR. More details of iPHR are available in [7, 8, 9]. iPHR uses a set of standardized nursing languages to link health issues to search guide phrases [8, 9]. Each health issue links to one or more *nursing diagnoses* [1] that are clinical judgments about individual, family, or community responses to actual or potential health problems [11]. Each nursing diagnosis links to a list of *nursing interventions* representing treatments that can be performed to enhance patient/client outcomes [4, 5]. Each nursing intervention includes a list of *home nursing activities* (HNAs) representing the actions that patients and caregivers can perform at home or in the community [4]. For each HNA, a nurse pre-compiles a set of search guide phrases that are stored in iPHR's knowledge base.

Using nursing diagnoses, nursing interventions, and HNAs as intermediate steps, we can link each health issue to multiple search guide phrases compiled using nursing knowledge, as shown in Fig. 2. Moreover, using the concepts of virtual nursing diagnosis, virtual nursing intervention, and virtual HNA, we can link health issues to search guide phrases compiled from other sources (e.g., treatment knowledge) [8]. For instance, for a disease, such

compiled search guide phrases include its name, its treatment methods, the names of its symptoms, and the treatment methods of its symptoms.

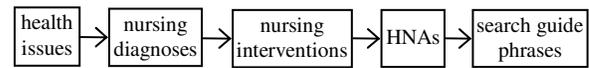


Fig. 2 Linking health issues to search guide phrases.

iPHR contains the user's personal health record as one of its components [7]. iPHR automatically extracts from the personal health record the health issues related to the user, e.g., her current disease. For a given health issue, iPHR uses each search guide phrase linked to it and a search engine to retrieve some HMPs. The combination of all retrieved HMPs is returned to the user. As mentioned in the introduction, the user often cannot use the sequential output interface to quickly reach her desired HMPs. Moreover, this problem cannot be solved by the hierarchical output interface described in [8], where HMPs are listed under their corresponding HNAs. This is because a layman user typically has little nursing knowledge and cannot well understand HNAs. Frequently a health issue links to hundreds of HNAs with long descriptions [9]. At the level of the hierarchical output interface listing HNAs, the user can be overwhelmed by the large number of unfamiliar HNAs as well as frustrated by the fact that no HMP is available for direct viewing.

3. Navigation interface

To help users quickly find desired HMPs, we need a navigation interface displaying both retrieved HMPs and the navigation hierarchy simultaneously. The key of this interface is the navigation hierarchy that is based on product categories. One could use some automatic clustering algorithm to construct the navigation hierarchy, but the quality of the resulting hierarchy is usually unsatisfactory [10, 15]. A better method for constructing the navigation hierarchy is to adjust an existing, manually-built, high-quality product category hierarchy.

3.1 Product category hierarchy

Typically, a product category hierarchy is manually built and maintained for each product shopping Web site. All category nodes in the hierarchy are carefully crafted so that layman users can easily understand them. For the results of a keyword search query, the Web site constructs its product navigation interface according to this category hierarchy by displaying a few appropriate category nodes, each of which contains at least one retrieved product. In this category hierarchy, each category that is not a leaf node has a small number (e.g., 10) of "major" sub-categories as its child nodes. Each sub-category has a descriptive name and usually includes many products. Combined together, these

sub-categories cover a large percentage (e.g., 60%), but not all, of the products in the category. In other words, many products in the category are excluded from these sub-categories and put into a default sub-category entitled “others.” In theory, we could construct a large number of “minor” sub-categories, each containing a small number of products, to cover the products in the “others” sub-category. However, this is both labor-intensive and unnecessary, as the navigation hierarchy can display only a small number of sub-categories without overwhelming the user.

iPHR uses a vertical search engine by crawling Web pages from one or more high-quality HMP shopping Web sites. In the rest of this section, we focus on one Web site and assume that the user selects one topic of concern in the input interface of iPHR’s HMP recommendation function [8]. In the case of multiple Web sites or multiple topics, we can add an additional level of Web site name or topic of concern into the navigation hierarchy.

3.2 Navigation hierarchy construction

To construct the navigation hierarchy in iPHR’s HMP navigation interface, our first thought is to use the existing product category hierarchy of the HMP shopping Web site. That is, iPHR and the Web site use the same navigation interface. When the user clicks a category in the navigation hierarchy, iPHR uses the same method as that in [8] to rank retrieved HMPs in that category and displays them sequentially on the right side of the navigation interface. When the user’s topic of concern (e.g., fishing activity) is not a health issue, iPHR uses a single query (the topic’s name) to retrieve HMPs [8] and this navigation hierarchy may be the best that can be done. However, a better navigation hierarchy can be constructed when the user’s topic of concern is a health issue.

Recall that in the existing product category hierarchy of the HMP shopping Web site, many HMPs in a category C are omitted in C ’s “major” sub-categories and have to be put into C ’s “others” sub-category C_o . The user can easily move through the large number of retrieved HMPs in a “major” sub-category C_m by navigating to C_m ’s child categories. However, C_o has no child category and hence the user cannot easily move through the large number of retrieved HMPs in C_o .

To address this problem, we use treatment knowledge and nursing knowledge to add child categories to the “others” sub-category C_o . We notice that unlike traditional keyword search on the HMP shopping Web site, iPHR’s HMP recommendation function has a special property: HMPs are retrieved for a health issue via multiple search guide phrases. These phrases can serve as natural child categories of C_o . Each child category contains all HMPs that are in C_o and retrieved by the corresponding phrase.

Besides these search guide phrases, multiple nursing diagnoses, nursing interventions, and HNAs also link to the health issue. However, those nursing concepts are

unsuitable to serve as child categories of the “others” sub-category C_o . Their descriptions are too long to fit into the limited display space that is on the left side of the navigation interface and allocated to the navigation hierarchy. Layman users with little nursing knowledge also have difficulty in understanding them.

A health issue often links to hundreds of search guide phrases. Consequently, the “others” sub-category C_o can have hundreds of child categories. To avoid overwhelming the user, only a subset of these child categories should be displayed in the navigation hierarchy. That is, we need to sort these child categories in a proper way and display only the top T_c child categories, where T_c is a predetermined constant. In iPHR, the default value of T_c is 20.

To facilitate the user’s navigation process, the child categories need to be ranked according to the following three heuristics. First, each child category c corresponds to a search guide phrase f . f has an overall weight o_{w_f} reflecting the global priority of both c and f . A child category with a larger o_{w_f} should rank higher. Second, a child category including many HMPs tends to provide more HMP information and should rank higher than a child category including only a few HMPs. Third, different child categories can have similarities because their corresponding search guide phrases can link to the same HNA, nursing intervention, or nursing diagnosis. The user would prefer to see dissimilar child categories, especially at the top of the child category list, to quickly gain as much new information as possible [8, 16].

These three heuristics respectively represent three factors: priority, HMP information richness, and diversity. To properly rank child categories, we fold all three factors into a single ranking formula. Each child category c has a score, $score_c$, computed by this formula. All child categories are sorted in descending order of this score.

Factor 1: Priority

Let f be the search guide phrase corresponding to the child category c . f is one of n_A search guide phrases compiled for the HNA A . As described in [8, 9], each nursing diagnosis, nursing intervention, and HNA has a normalized weight reflecting its “normalized” priority. We can similarly define a normalized weight $n_{w_f}=1/n_A$ for f to reflect its “normalized” priority, where all n_A search guide phrases compiled for A are treated equally important. The global priority of f depends on f ’s own priority, A ’s priority, the priority of the nursing intervention I linked to A , and the priority of the nursing diagnosis D linked to I . Consequently, the overall weight o_{w_f} of f is defined as the product of f ’s normalized weight n_{w_f} , A ’s normalized weight n_{w_A} , I ’s normalized weight n_{w_I} , and D ’s normalized weight n_{w_D} . That is, $o_{w_f}=n_{w_f} \cdot n_{w_A} \cdot n_{w_I} \cdot n_{w_D}$. In the case of multiple

linked HNAs, nursing interventions, or nursing diagnoses, all those overall weights are summed into a single number.

Factor 2: HMP information richness

Let N_c represent the number of HMPs contained in the child category c . Since the navigation interface only displays child categories each with at least one retrieved HMP, we always have $N_c \geq 1$. We use $r_c = \ln(N_c + 1)$ to measure the HMP information richness of c , where the addition of 1 is used to prevent this measure from becoming zero.

To reflect the two heuristics on priority and HMP information richness, the score of the child category c , $score_c$, is defined as the product of the overall weight o_{w_f} and the HMP information richness measure r_c : $score_c = o_{w_f} \cdot r_c$. Since r_c is always positive, we can use $score_c$ to differentiate those child categories with the same r_c but different o_{w_f} 's. To reflect the heuristics on diversity, we resort to multi-pass weight discounting and score re-computation.

Factor 3: Diversity

Let N_o represent the number of child categories of the "others" sub-category C_o . Recall T_c is the maximum number of child categories of C_o that can be displayed in the navigation hierarchy. To provide diverse child categories, we re-rank all N_o child categories in $H_o = \min(T_c, N_o)$ passes to sequentially select the H_o child categories for display in the navigation hierarchy. In each pass, among all the child categories that have not been selected before, we pick the one, c_{max} , with the largest score. Then we give appropriate discounts to the normalized weights related to c_{max} . Specifically, suppose the search guide phrase corresponding to c_{max} is compiled for the HNA A , A links to the nursing intervention I , and I links to the nursing diagnosis D . A 's normalized weight n_{w_A} is discounted by d_A . I 's normalized weight n_{w_I} is discounted by d_I . D 's normalized weight n_{w_D} is discounted by d_D . d_A , d_I , and d_D are three constant factors whose default values are all 0.85. The scores of some child categories depend on n_{w_A} , n_{w_I} , and n_{w_D} , and thus need to be re-computed. As a result, the more child categories related to a HNA that have been selected, the less likely the next selected child category will be related to this HNA. A similar property exists for nursing interventions and nursing diagnoses.

4. Experimental results

We implemented the HMP navigation interface in a prototype iPHR system. We conducted experiments under a wide range of medical scenarios to demonstrate the effectiveness of our techniques.

4.1 Experimental setup

We crawled 150,000 HMP Web pages from Amazon [2], the largest online shopping Web site. We compared the navigation interface with the sequential output interface. We used United States Medical Licensing Examination (USMLE) Step 2 CS (Clinical Skills) medical exam cases [6]. Physicians have to pass this USMLE exam to obtain their licenses for practicing medicine. Each exam case has both a sample medical record and a summary that includes a several-page-long, detailed description of the patient's situation.

We randomly selected 30 USMLE medical exam cases as our test cases. Since USMLE covers both the typical cases and almost every aspect of daily medical practice, our random samples have a broad coverage of medical topics. Ten people, six females and four males, served as users. Their median age is 38. All of them are regular, ordinary Internet users without formal medical training, and hence represent iPHR's targeted users. They all have received college education or above.

Each user searched for all 30 medical cases, by inputting the health issues mentioned in the medical cases in the input interface of iPHR's HMP recommendation function. For every medical case, each user randomly selected either the sequential output interface or the navigation interface with equal probability, and had up to 40 minutes to search. The user was asked to focus on quickly finding the HMPs that she would prefer to buy if she were the patient rather than just browsing HMPs for general interest. The search session was terminated when either the user felt she had found enough desired HMPs or time ran out, whichever came first. We allowed users to search for a relatively long time because users care about their health and often spend significant time searching medical information.

In practice, when using iPHR's HMP recommendation function, users usually have no access to nurses, but are aware of their health issues and have some understanding of them. The scenario with our experiment is different. Users read the medical case description on the spot to get to know the situation faced by the patient involved. Since the medical case description often contains medical jargon that they are unfamiliar with, they may encountered difficulty in understanding it. When such difficulty occurred, a nurse was available to explain the medical jargon and provide any needed background information. The nurse's role was limited and she did not assist users with iPHR.

Similar to the TREC interactive track [12] that provides a standard approach for comparing the performance of various information retrieval systems, we use two sets of measures as the performance metrics for the output interface: one set is objective and the other set is subjective. The objective performance measures include the number of search result Web pages viewed and the time spent on the search process. The subjective performance measures from the users' perceptions include the number of desired HMPs

found, ease of using the system, usefulness of the search results, and overall satisfaction with the system. For the navigation interface, the average overall satisfaction with the top 10 child categories of the “others” sub-category is also included. Except for the number of desired HMPs found, all of these subjective performance measures are on a 7-point scale, with 1=low and 7=high [12]. They were obtained from a brief questionnaire that users completed after using the system. For each objective or subjective performance measure, an average is computed for all 30 medical cases and all ten users, and both its mean and its standard deviation are reported. We used ANOVA [3] as the significance test. Our experiments were performed on a computer with two 3GHz processors, 2GB memory, and one 111GB disk.

4.2 Overall results

iPHR can quickly construct the sequential output interface and the navigation interface, both on average in less than two seconds.

Table 1. Objective performance measures (* means significant at <0.05 level).

mean (standard deviation)	sequential output interface	navigation interface
number of search result Web pages viewed	32* (7)	25 (6)
time (minutes)	23* (5)	17 (5)

The navigation interface organizes retrieved HMPs in a better way than the sequential output interface. The navigation interface clearly marks the HMP categories and provides the user the capability of viewing retrieved HMPs in different categories. In contrast, the sequential output interface mixes together HMPs from all categories and hence creates difficulty for the user to find desired HMPs in preferred categories. Due to the better organization of the navigation interface, a user of the navigation interface views fewer search results and spends less time on the search process than a user of the sequential output interface does (see Table 1). Both differences are statistically significant.

Table 2 shows the subjective performance measures. For the same reason stated above, users find that the navigation interface can help find a larger number of desired HMPs, is easier to use, and is more satisfactory than the sequential output interface. All three differences are statistically significant. Irrespective of the output interface used, iPHR retrieves the same set of HMPs for a given health issue. Hence, the usefulness of the search results is the same for both interfaces. In the navigation interface, if child categories are not provided for the “others” sub-category, the average overall satisfaction with the system will drop from 5.7 to 5.2, as users can have difficulty navigating

through the large number of HMPs in the “others” sub-category.

Table 2. Subjective performance measures (* means significant at <0.05 level).

mean (standard deviation)	sequential output interface	navigation interface
number of desired HMPs found	12* (4)	19 (5)
ease of use	4.8* (0.9)	5.5 (1.0)
usefulness	5.1 (0.9)	5.1 (0.9)
satisfaction with the system	4.8* (1.1)	5.7 (0.8)

iPHR uses three parameters to provide diverse child categories of the “others” sub-category (see Section 3.2): d_D (nursing diagnosis normalized weight discount factor), d_I (nursing intervention normalized weight discount factor), and d_A (HNA normalized weight discount factor). We did three experiments. In each experiment, we fixed one parameter at its default value and varied the values of the other two parameters. The results of these three experiments are similar. We only show the results of one of these three experiments here. In this experiment, we fixed d_A at its default value and varied the values of d_D and d_I .

Fig. 3 shows the impacts of the parameters d_D and d_I on the average overall satisfaction with the top 10 child categories of the “others” sub-category, when d_A is fixed at its default value. When d_D is either too small or too large, the average overall satisfaction decreases. A similar property exists for d_I .

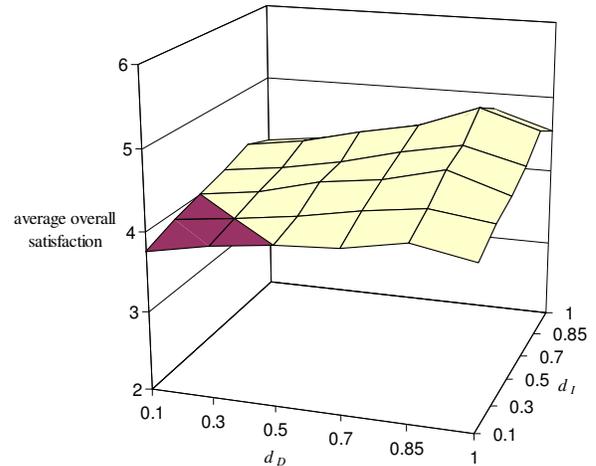


Fig. 3 Average overall satisfaction with the top 10 child categories of the “others” sub-category vs. d_D and d_I .

Using default values of the parameters, the average overall satisfaction with the top 10 child categories of the “others” sub-category is 5.7. If the diversity factor is ignored (or equivalently, $d_D=d_I=d_A=1$) when ranking child categories of the “others” sub-category, the average overall satisfaction will drop to 4.55. In summary, the average

overall satisfaction is maximized around the default values of the three parameters. Each of the three parameters has a not-very-small safe range, within which the average overall satisfaction is insensitive to parameter changes. However, if the value of a parameter is outside its safe range, the average overall satisfaction may drop.

4.3 A Detailed Example

To give the reader a feeling of iPHR's HMP navigation interface, we present detailed results of the navigation hierarchy constructed for a typical health issue *muscular dystrophy*. Fig. 4 shows the constructed navigation hierarchy when the user clicks the top-level "medical supplies & equipment" category followed by its "daily living aids" sub-category.

- < Medical Supplies & Equipment
 - < **Daily Living Aids**
 - * Bath & Body Aids (7159)
 - * Eating & Drinking Aids (1438)
 - * Low Vision Aids (1235)
 - * Dressing Aids (1086)
 - * Low Strength Aids (671)
 - * Medication Aids (610)
 - * Hearing Aid Accessories (303)
 - * Ramps (286)
 - * Hearing Aids (226)
 - * Telephones (87)
 - * Others (856)

Fig. 4 A sample navigation hierarchy constructed for the health issue *muscular dystrophy*.

- * Cane & walker (36)
- * Relaxation (49)
- * Positioning pillow & wedge (16)
- * Pain relief & management (27)
- * Fall prevention (17)
- * Nutritional supplement (3)
- * Foam pad (20)
- * Velcro closure (12)
- * Strength training (4)
- * Weakness (20)
- * Massager (35)
- * Crutch (7)
- * Foam wedge (47)
- * Walking aid (9)
- * Fatigue (31)

Fig. 5 The top 15 child categories of the "others" sub-category shown in Fig. 4.

Fig. 5 shows the top 15 child categories of the "others" sub-category of the "daily living aids" sub-category. In general, for a health issue, iPHR's HMP navigation

interface can classify retrieved HMPs into a reasonable number of appropriately named and sorted (sub-)categories to facilitate users to quickly reach their desired HMPs.

5. Conclusion

This paper presents a HMP navigation interface that can facilitate users to quickly reach their desired HMPs, particularly when they already have a rough idea of some preferred categories that these HMPs should fall into. Our experiments show that due to its better organization, this navigation interface outperforms the sequential output interface and significantly improves user satisfaction.

Acknowledgment We thank Selena B. Thomas, Chunqiang Tang, and Jing Wang for helpful discussions.

References

1. Ackley, B. J., and Ladwig, G. B., *Nursing diagnosis handbook: an evidence-based guide to planning care*, 8th ed. Mosby, 2007.
2. Amazon homepage. <http://www.amazon.com>, 2010.
3. Bickel, P. J., and Doksum, K. A., *Mathematical statistics: basic ideas and selected topics*, Vol. 1. Prentice Hall, 2001.
4. Bulechek, G. M., Butcher, H. K., and Dochterman, J. M., *Nursing interventions classification (NIC)*, 5th ed. Mosby, 2007.
5. Johnson, M., Bulechek, G. M., and Dochterman, J. M., et al., *NANDA, NOC, and NIC linkages: nursing diagnoses, outcomes, and interventions*, 2nd ed. Mosby, 2005.
6. Le, T., Bhushan, V., *First aid for the USMLE step 2 CS (clinical skills exam)*. McGraw-Hill, 2006.
7. Luo, G., Thomas, S. B., and Tang, C., Intelligent consumer-centric electronic medical record. *Proceedings of MIE'09*, pp. 120-124, 2009.
8. Luo, G., Thomas, S. B., and Tang, C., Automatic home medical product recommendation. *JMS*, to appear, available at <http://pages.cs.wisc.edu/~gangluo/device.pdf>.
9. Luo, G., and Tang, C., Automatic home nursing activity recommendation. *Proceedings of AMIA 2009*: 401-405.
10. Manning, C. D., Raghavan, P., and Schutze, H., *Introduction to information retrieval*. Cambridge University Press, 2008.
11. NANDA International. *Nursing diagnoses, 2009-11 edition: definitions and classification*, 2nd ed. Wiley-Blackwell, 2008.
12. TREC interactive track homepage. <http://trec.nist.gov/data/interactive.html>, 2010.
13. *U.S. market for home care products*, 5th ed. Kalorama Information, 2007.
14. Hidola homepage. <http://www.hidola.com>, 2010.

15. Lawrie, D. J., and Croft, W. B., Generating hierarchical summaries for web searches. *Proceedings of SIGIR'03*, pp. 457-458, 2003.
16. Carbonell, J. G., and Goldstein, J., The use of MMR, diversity-based reranking for reordering documents and producing summaries. *Proceedings of SIGIR'98*, pp. 335-336, 1998.
17. Adomavicius, G., and Tuzhilin, A., Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *TKDE* 17(6): 734-749, 2005.