

INFSCI 2930 - INDEPENDENT STUDY

**A profile-based exploratory search and
recommendation system**
for finding research advisors

Behnam Rahdari and Jaime Fawcett
(ber58@pitt.edu , jaf157@pitt.edu)
Fall 2019

Abstract

Finding research advisors is a known challenge since students' research interests are frequently vague and not expressed in terms used by their prospective advisors. In this paper, we present an exploratory search and recommendation system, which assists students in establishing their interest profiles and uses it to recommend the most relevant advisors. A user study with a target audience demonstrated that students extensively used exploratory features and revealed interesting patterns in student adoption of exploratory and control features.

Introduction and Background

A number of real-life search tasks require a considerable amount of learning during the process of search to achieve adequate results. These tasks are known as exploratory search tasks. Since simple search systems are usually not efficient in supporting exploratory search, a range of advanced exploratory systems have been developed and evaluated (Roth., 2009). A few projects in this area demonstrated that the effectiveness of exploratory search could be improved by using a personalized system, which builds a profile of user interests and adapts to the individual user (Brusilovsky, 2013), (Amanda Goncalves Dias, 2019) and (Jaakko Peltonen, 2017). The work presented in this paper investigates the ideas of profile-based exploratory search in the context of finding research advisors.

Finding an advisor for a research project or a graduate study program is a good example of an exploratory search task since prospective students rarely express or understand their interests using clear scientific terms that experienced researchers may use. For example, simple search will hardly help someone interested in working on "big data in medicine" to find and recognize an advisor who applies machine learning approaches to health records. The *Grapevine system* presented in this paper combines features of exploratory search systems and a controllable recommender interface (Bart P. Knijnenburg, 2012), (John O'Donovan, 2008) to help students discover and create their interest profiles (Susan Gauch, 2007) which, in turn, is used to recommend relevant advisors with matching interests. This paper presents *Grapevine's interface*, stressing its exploratory search and profile-based recommendation components. The paper also details a controlled user study used to assess the system and reviews the results from this study.

The Grapevine System

Grapevine is an exploratory search system enabling novice students to gain familiarity with research topics and find research advisors with similar research interests. *Grapevine's* intelligent user interface is driven by a knowledge graph, a tightly connected network of research topics and prospective advisors. *Grapevine's* novelty and value added is its ability to enable its users to gradually form their interest profiles by discovering new research topics and keywords that match their not-yet-well-formulated interests. Following the nature of exploratory search, *Grapevine* combines focused search with exploratory browsing activities (Roth., 2009). In particular, it connects multiple effective exploratory search systems, which provide users with different opportunities to recognize (rather than to formulate) research topics and interests shown in different contexts (Claypool, 2008) (Cecilia di Sciascio, 2016) (Hearst., 2006) (Jaakko Peltonen, 2017).

In this section, we review the interface's features and explain the process of building the knowledge graph.




Figure 1: Grapevine system interface. Features of the interface include A: Search box, B: Recommended keywords, C: Interest profile and sliders, D: Results list, E: Heat bar and F: Final list.

The Interface Design

During exploratory search, the *Grapevine* system engages users in building their interest profile, which is visualized as a list of preferred keywords (Figure 1-C). *Grapevine* offers users several ways to discover their interests. First, users are able to find keywords of interest and add them to their profile directly using a search box (Figure 1-A). To better engage recognition (vs. recall) and to expand the user's vocabulary, the system engages the knowledge graph to recommend semantically similar keywords (Figure 1-B). Similarly, the knowledge graph is used to generate a full list of research interests for each advisor recommended and explored (Figure 2}). Relevant keywords found in each of these contexts can be added to the user's profile. Once a user's interest profile is established, the system morphs into a controllable recommender system. To allow greater user control in finding a relevant advisor, each specific interest (represented as a keyword or keyword phrase in the user's profile) can be weighted separately using sliders (Figure 1-C).

Search box

The search-box (Figure 1-A) is the gateway to the system. Using an instant search approach, it allows users to discover relevant keywords without a fully formulated query. When a user starts typing a query, a series of textually similar keywords appears, which helps the user to discover a range of matching concepts (e.g., Data Mining and Educational Data Mining). All keywords in the knowledge graph are considered when generating the list of instant search matches.




Figure 2: Detailed view showing demographic information and a list of the advisor's research interests.

Recommended Keywords

When at least one keyword is added to the user's profile, a series of five semantically similar keywords will appear in the *Similar Keywords* area of the interface (Figure 1-B). Users can add these recommended keywords to their interest profiles by clicking on the + button to the right of each keyword. As the user's profile grows and refines, the set of recommended keywords is updated since the system recommends instances similar to *all keywords* in the user's profile.

Profile and Sliders

The user profile area (Figure 1-C) is both a placeholder for the interest keywords and a mechanism for user-controlled recommendation. Based on these keywords, the system generates and ranks the list of recommended advisors (Figure 1-E). In addition to adding keywords using the methods described above, users can remove less relevant keywords (using the red x) as they discover more relevant keywords or explore different interests. Sliders associated with each keyword enable users to reflect the relative importance of the keyword compared to others in their profile, ranging from 1 (least important) to 10 (most important). The initial value of the sliders is set to five but can be changed at any time. All actions within the profile (adding, removing, or adjusting sliders) immediately affects the list of recommended advisors.

Recommended Advisors

As soon as the user adds the first keyword to the profile of interests, a list of the 20 most relevant advisors will appear as individual cards (Figure 1-D). Each card contains a photo and brief information about the advisor. There is also a relevance bar (Figure 1-E) to reflect the advisor's relevance to the user's profile and two action buttons. The button *Details* opens the Details view (Figure 2) that provides more details about the advisor. The button *Select* adds the advisor to the final results list (Figure 1-F).

Details view

Details view (Figure 2) provides users with additional information about the advisors such as affiliation, research impact, and research interests (shown as the list of keywords). To highlight advisor's relevance to the user profile, the research interests that match the user profile are shown in green. Since faculty usually explore a set of related topics, we expected that the research interests or the prospective advisor not yet added to the user profile (shown in blue) might be also of relevance to the given user. To support the “interest discovery” process, these blue keywords can be added to the user's profile of interests with one click.

The Knowledge Graph

The knowledge graph consists of multiple entries extracted from Google Scholar¹, enriched by Wikipedia and hosted on a native graph database (Neo4j²).

Building the knowledge graph included the following stages.

Data Collection

For the purpose of user study presented in this paper, we retrieved information on the top 1000 (based on citations) faculty members in two specific areas of research from 15 universities in the Atlantic Coast Conference³. We extracted academic information, including name, affiliation, number of citations, h-index, i10-index, self-defined areas of research (represented by keywords), and a list of the 20 most recent publications. We only include individuals with at least one publication within the last two years to ensure all individuals are still active as prospective advisors.

Next, we used the list of recent publications to generate keywords (unigrams, bigrams and trigrams) that represent each advisor's recommended research interests. Finally, we captured the top 20 co-authors for each advisor to represent the social connections between individuals.

Verification and Enrichment

We used the Wikipedia API⁴ to filter all extracted keywords; only keywords with an entry in Wikipedia were kept in the knowledge graph. It assured keyword quality and made it easier to weight each link between a keyword and a faculty using cosine similarity between the text of the Wikipedia page and the faculty publications. Furthermore, we added up to 10 Links to the knowledge graph for the Categories in the Wikipedia page that represent a keyword. The connections

added a semantic layer to the knowledge graph and were used to recommend semantically similar keywords to the users.

¹ https://en.wikipedia.org/wiki/Google_Scholar

² <https://en.wikipedia.org/wiki/Neo4j>

³ https://en.wikipedia.org/wiki/Atlantic_Coast_Conference

⁴ <https://pypi.org/project/wikipedia>




Figure 3: Graph Schema. This figure shows the entities of knowledge graph and the relationship between them

Graph Schema

Figure 3 presents the schematic representation of the knowledge graph. Advisors are interconnected by the relation *Works_With* (based on co-authorship) and connected to Wikipedia-verified keywords (self-defined and extracted from publications) by the relation *Has_Key*. The latter carries a weight that determines the strength of the relationship between each keyword and the advisor's research.

Keywords are connected to Links and Categories (*Has_Link* and *Has_Cat*). Links are other Wikipedia keywords not connected directly to advisors. The Link is established between keywords if it has been mentioned in the keywords' respective Wikipedia pages.

Recommendation Method

We used Cypher Querying Language⁵ to generate both advisor and keyword recommendations. For advisor recommendations, at each instance of user interaction with the system (e.g., adding/removing keywords), the system matches all advisors with at least one of the user's designated interests. Then, a score is assigned to each candidate based on the number of matches and the multiplication of the value of the sliders with the similarity measure that is introduced in "*Verification and Enrichment*". Finally, the system ranks the candidates by their scores and presents the top 20 candidates as results. When there are less than 20 candidates to recommend, similar keywords are used to find more advisors.

To generate recommended keywords, for each set of keywords in the user's profile, the system generates three sets of candidate keywords. These sets are created using the co-occurrence of seed keywords and advisors' research interests, links, and categories (using collaborative filtering). Then, the system combines the number of co-occurred keywords in all three sets and uses it as a ranking mechanism. The system presents the top five results to the user.

Evaluation

Participants

To make our evaluating as close to the target case as possible, as our evaluation group, we selected a cohort of students attending a summer institute in an information science school, which prepares

⁵ [https://en.wikipedia.org/wiki/Cypher_\(Query_Language\)](https://en.wikipedia.org/wiki/Cypher_(Query_Language))

students from underrepresented populations for graduate study. These students have vested interest in finding prospective research advisers and offer a good example of our target user population, i.e., students less familiar with the way their prospective advisors describe their research interests. Atlantic conference colleges selected for our study were among colleges that these students considered for their graduate study. In total, 19 students from the summer institute participated in the study. To increase the reliability of data, we recruited 8 additional participants from the undergraduate program of our home department who were the closest match to the summer institute cohort.

In total, we enrolled 27 participants (16 male, 11 female). The average age was 20-years-old ($SD = 1.39$). The majority of participants were completing an undergraduate degree ($N = 24$).

Eighteen participants were pursuing degrees in computer science ($N = 17$) and information science ($N = 7$). All participants completed an electronic consent form prior to participating in the study. Participants were not compensated for participation.

Procedure

The study design involved participants viewing a tutorial of the system, completing a pre-study survey, completing a task using the system, and completing a survey following the task (in this order). The study lasted approximately 50 minutes.

The tutorial highlighted the key functions of the system.

The pre-study questionnaire contained 8 items that asked about basic demographic information and familiarity with recommendation-based systems.

Next, participants completed an exploratory search task using the system. The task was: “Using the interface, identify five faculty members that match your research interests.” We did not provide any other parameters for how to complete the task as we wanted the participants to base their interaction on their actual interests and motivations for pursuing graduate study or identifying advisors with similar interests. Log data and interaction data were collected while participants completed the task.

Finally, participants completed a post-task questionnaire. The questionnaire contained 38 items. The first five items asked participants to describe why they chose each of the five advisors in their final list. The next 25 items asked participants to rate their agreement with a series of statements associated with five constructs: (1) overall satisfaction with the system; (2) satisfaction with the results of the system; (3) confidence in the results of the system; (4) cognitive load of using the system; and (5) their perceived control of the system. These items were based on previously developed items used to evaluate user interactions with recommender systems.

Results and Discussion

Since one of the key goals of the system was to offer several ways to discover of users' research interests and to control the process of recommendation, our first target was to examine how the keyword selection and weighting features were used. As the first data row of the Table 2 shows, all three ways of adding keywords were used with a comparable frequency and the sliders were extensively used. However, a deeper analysis revealed that the use of exploratory features varied considerably between students. Some students treated the system mostly as a regular search system adding keywords exclusively by search, while others extensively used recognition-focused exploration using search only at the start and expanding their interest profiles using recommended

keywords and faculty details. To investigate the effect of different exploration approaches on student success and satisfaction with the system, we split participants into three groups: (1) *Explorers* - 9 participants who primarily added keywords to their profiles using recommended keywords or advisor keywords; (2) *Neutrals* - 10 participants who used a combination of methods to identify keywords; and (3) *Searchers* - 8 participants who added keywords to their profiles using search.

Interaction	Average	Explorers	Neutrals	Searchers
Keywords added	12.59	14.00	13.40	10.00
Added by search	5.56	3.23	6.30	7.25
Added by recom	3.70	5.33	3.80	1.75
Added by details	3.30	5.44	3.30	1.00
Total exploration	7.04	10.77	7.10	2.75
Keyword “churn”	18.62	22.88	19.50	13.50
Deleted keywords	12.59	8.88	6.10	3.50
Active keywords	3.51	4.11	3.81	2.63
Slider use	13.91	14.33	13.90	13.50
Details viewed	15.74	24.22	15.50	7.50

Table 2: System Usage Statistics. This table shows the average usage of different interaction of users with the system for all and three groups of users.

As Table 2 shows, there was a considerable difference between groups beyond the balance of *Added by search* and *Total exploration*, which was used to define the groups. Explorers were more heavily engaged into the work with keywords, as shown by a remarkable difference in *Keyword “churn”*, the sum of all actions with keywords. They also used keywords more extensively for faculty recommendation, as shown by a considerably higher number of *Active keywords*. To calculate the active keywords, we recorded the number of keywords in the users' interest profiles at each action point (opening advisor's details or adding an advisor to their final list) and calculated the average for each student. Finally, the explorers used sliders slightly more frequently. Altogether, the Explorers better embraced the exploratory features of the system than Searchers. The log data provided some evidence that it helped the Explorers to be more productive in their work - they were able to discover 3 times more interesting faculty to examine.

	Average	Neutrals	Searchers	Explorers
Satisfaction	3.14	3.85	3.16	2.33
Results	3.46	3.92	3.43	3.00
Confidence	3.11	3.77	2.83	2.70
Cognitive Load	3.40	4.22	3.37	2.54
Control	3.37	3.88	3.50	2.62

Table 3: User feedback analysis of the post-task survey (five-point Likert scale items) showing the average and per-group results for the five domains.

Table 3 shows the average scores for the post-task survey constructs. The results show that the explorers were more satisfied with results and the system. (Despite exploring three times more advisors, the cognitive load had a very small increase).

	Average	Explorers	Neutrals	Searchers
Computer Use	4.52	4.62	4.41	4.55
Internet Use	4.17	4.42	3.93	4.18
Search Use	3.27	3.23	3.42	3.16
Feeling about Recommenders	3.12	3.68	3.12	2.84

Table 4: User feedback analysis of the pre-test survey.

Our study demonstrated that the provision or novel exploratory features might not be sufficient, the students should embrace and use these features to achieve better success and satisfaction. We were, however, not able to determine what causes students to become explorers rather than searchers.

Table 4 shows that it is unlikely defined by their past experience since there was no difference in past experience across groups. Explorers, however, were slightly more positive about recommendation pointing to the possible attitude differences.

References

- Amanda Goncalves Dias, E. E. (2019). TRIVIR: A Visualization System to Support Document Retrieval with High Recall. *ACM Symposium on Document Engineering*.
- Bart P. Knijnenburg, S. B. (2012). Inspectability and Control in Social Recommenders. *6th ACM Conference on Recommender System*.
- Brusilovsky, J. A. (2013). Adaptive visualization for exploratory information retrieval. *Information Processing and Management*.
- Cecilia di Sciascio, V. S. (2016). Rank As You Go: User-Driven Exploration of Search Results. *21st International Conference on Intelligent User Interfaces (IUI '16)*.
- Claypool, M. a. (2008). The relation browser tool for faceted exploratory search. *Proceedings of the 8th ACM/IEEE-CS joint conference on Digital libraries*.
- Jaakko Peltonen, J. S. (2017). Negative Relevance Feedback for Exploratory Search with Visual Interactive Intent Modeling. *22nd International Conference on Intelligent User Interfaces*.
- John O'Donovan, B. S. (2008). PeerChooser: visual interactive recommendation. *twenty-sixth annual SIGCHI conference on Human factors in computing systems*.
- Roth., R. W. (2009). Exploratory Search: Beyond the Query-Response Paradigm. Morgan and Claypool.
- Susan Gauch, M. S. (2007). User Profiles for Personalized Information Access. *Adaptive Web: Methods and Strategies of Web Personalization*.