User Exploration of Slider Facets in Interactive People Search System

Shuguang Han[1], Danchen Zhang[1], Daqing He[1] and Qikai Cheng[2]
[1]School of Information Sciences, University of Pittsburgh, USA
[2]School of Information Management, Wuhan University, China

Abstract
People search is an important search task where the goal is to find relevant people instead of relevant documents. Providing search facets in a people search system can help users better describe their search intents. In existing studies, a search facet is either represented by checkboxes if the facet can be divided into discrete values, or by a slider in which users can tune for continuous values. Slider facets enable a people search system to handle both discrete-value and continuous-value facets, and thus are shown to be more effective. Yet, the ways of how users interact with slider facets in people search are rarely studied. Based on a user study with 24 participants using an interactive people search interface with three slider facets, we find that users indeed utilize sliders consistently in their search processes to fine-tune search results. We also find that although slider tuning brings performance boost but users are lack of abilities to locate the optimal facet-values, which indicates the necessity of providing automatic facet-value suggestion.

Keywords: People search; faceted search; faceted people search

DOI:10.9776/16324.

Acknowledgements: Thanks for the useful suggestions provided by Zhen Yue (Yahoo! Labs, zhenyue@yahoo-inc.com)

Contact: shh69@pitt.edu, daz45@pitt.edu, dah44@pitt.edu and chengqikai0806@gmail.com

1 Background & Motivation

Although people's major information need is still on searching for web documents, researchers (Weerkamp et al., 2011; Hofmann, Balog, Bogers, & De Rijke, 2010) have observed an increasing trend of people search need. Existing studies of people search usually targeted for the academic purposes, such as finding experts (Balog, Azzopardi, & De Rijke, 2006), assigning paper reviewers (Karimzadehgan, Zhai, & Belford, 2008) or discovering potential collaborators (Han, He, Brusilovsky, & Yue, 2013). With the advancement of social networking services, people search is studied on the platforms such as Facebook and LinkedIn (Huang, Tunkelang, & Karahalios, 2014; Spirin, He, Develin, Karahalios, & Boucher, 2014). Studies (Han, He, Jiang, & Yue, 2013; Hofmann et al., 2010; Yarosh, Matthews, & Zhou, 2012) also find that there are multiple factors when users are searching for people, particularly when the tasks are exploratory. For example, a recent study (Han, He, Jiang, & Yue, 2013) finds that scholars are more aware of social similarity when finding collaborators while care more about the reputation when seeking for mentors/advisors.

Therefore, researchers have started to investigating the ways to leverage multiple factors for people search. Past studies either tried to design automatic supervised learning algorithms (Tang, Tang, & Tan, 2010) or employed interactive search interfaces (Han, He, Jiang, & Yue, 2013) that allow users manually control their preferences over different facets. The latter resembles the study of faceted search (Tunkelang, 2009), which usually involves a list of facets and their associated values for user manipulation. A study finds (Kules, Capra, Banta, & Sierra, 2009) that users indeed look at facets with about 31% of retrieval task time in a faceted web search system. Search facets are often represented as checkboxes (users can select and/or de-select a facet-value checkbox to add/remove certain candidates) or sliders (users can tune values on sliders to emphasize and de-emphasize the corresponding facet). Checkbox-based facet search system usually involves a large number of facet-value checkboxes, which may make users feel overwhelmed (Hearst, 2006). Slider-based facets could provide a concise search interface but is unable to give an overview of the data corpus (Han, He, Jiang, & Yue, 2013; Haveliwala, Jeh, & Kamvar, 2014; Parra, Brusilovsky, & Trattner, 2014). In this paper, we are interested in studying the slider-based facet search systems.

Slider tuning is a popular and intuitive interaction paradigm for discovering and mining applications.

Sample use cases include Google news[1] and thesaurus.com[2]. In these systems, slider tuning only changes the rank of each candidate instead of filtering out certain results. While faceted search has been extensively studied in previous studies (Ben-Yitzhak et al., 2008; Koren, Zhang, & Liu, 2008; Liberman & Lempel, 2012), the effectiveness of utilizing slider facets in supporting exploratory people search tasks has not yet been tested. It is unclear that how people would interact with the slider facets, and whether utilizing these facets can indeed help users find the optimal search results. These topics are the focus of this paper. As a result, we want to study the following two research questions: (1) how do people use the slider facets in people search interfaces? and (2) are users able to reach the optimal facet-values when using the slider facets?

2 Methodology

2.1 Research Platform

To answer our research questions, we develop an slider-based faceted people search system named PeopleExplorer[3]. The main interface is shown as Figure 1, which includes a query input box, three sliders and a result area to display returned candidates. Following a recent study (Han, He, Jiang, & Yue, 2013),our three sliders include (1) the content relevance that measures the degree of content match between a query and a candidate; (2) the authority of each candidate; and (3) the social similarity between the searcher and each candidate. Our system uses the title, abstract and author information of 219, 677 conference papers from ACM Digital Library[4] between 1990 and 2012, which involves 253,390 authors and 953,685 coauthor connections.

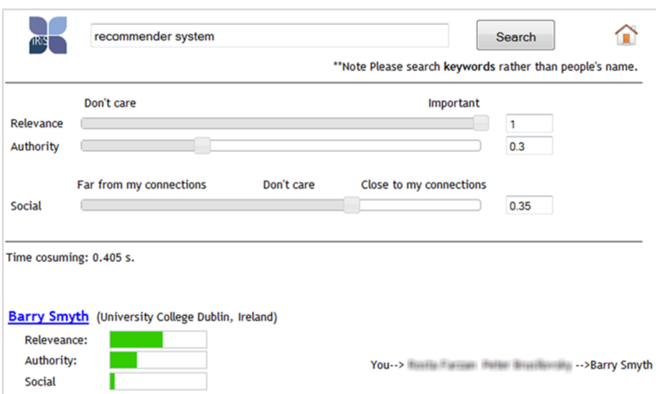

Figure 1: A screenshot of PeopleExplorer system

Candidates for each query q is ranked based on q and user-tuned weights for three sliders (w_c for Relevance, w_p for Authority and w_s for Social similarity). Followed the method proposed in (Han, He, Jiang, & Yue, 2013), Relevance $p(ca|q)$ is measured by the content matching between q and each candidate ca, which is the aggregation of document relevance if the candidate is one of the document authors. Authority is computed based on the Page Rank algorithm running on coauthor networks. Each candidate in our data collection will receive an authority value (i.e., pr_{ca}). Social similarity computes the social closeness between the user and each candidate, which is measured by the Jacarrd similarity between the candidate's coauthors and user's coauthors (i.e., $S_{u,ca}$). The final ranking list is generated based on the weighted score in Formula 1. The weights for three sliders all range from 0 to 1, and are directly obtained from user tuning.

$$S_{ca|u,q} = w_c \log p(ca|q) + w_p \log pr_{ca} + w_s \log S_{u,ca} \tag{1}$$

[1]https://news.google.com/
[2]http://www.thesaurus.com/
[3]http://crystal.exp.sis.pitt.edu:8080/PeopleExplorer/
[4]http://dl.acm.org/proceedings.cfm

2.2 User Study

We obtained users' real interactions on the provided sliders in PeopleExplorer through a user study. The user study was conducted between December 2012 and January 2013, where we recruited 24 participants (all PhD students majoring in information science or computer science). In the beginning, each participant was asked to provide his/her social connections (such as advisors, colleagues) to build his/her social profile so that we could compute his/her social similarity with each candidate in our system. In the post-task questionnaire, the participants were asked to rate the relevance of each marked candidate in a Five-point Likert scale (1 as non-relevant and 5 as the highly relevant). This data is used as the ground-truth information for our later evaluations. The study involved four different people search tasks directly adopted from (Han, He, Jiang, & Yue, 2013). The tasks are: (1) conference mentor finding (Task 1); (2) collaborator finding (Task 2); (3) external thesis committee member finding (Task 3); and (4) reviewer suggestion (Task 4). Each task aimed to find 5 candidates satisfying users' needs. The tasks orders were rotated based on Latin square. These four tasks are directly adopted from (Han, He, Jiang, & Yue, 2013).

3 Result Analysis

3.1 RQ1: How do people tune the sliders?

We firstly analyze how do people interact with sliders with two sub-questions: (1) were people indeed using the sliders; (2) were there any behavioral patterns when people tuned sliders?

To answer the first sub-question, we compute the number of slider tuning. Since each user worked on two tasks (the other two are served as baselines), we test whether users kept tuning sliders across two tasks. If a user used the sliders consistently, there would be no difference between the first and second task. We also test whether users consistently tuned sliders within one task. For each task, we divided user interactions into two sessions: the session in the first half and the session in the second half. As the results shown in Table 1, the participants indeed tuned the sliders to facilitate their exploration of candidates. Although it seems that the second half session (and task) has fewer slider tuning than the first half (and task), there are no significant difference. It is also interesting to see that the variance from the second session (task) is smaller than first session (task), which may be due to the learning effect - people learned the way of using facets.

Conditions	#slider tuning
Overall	6.761(5.29)
1st half session in both tasks	3.696(3.64)
2nd half session in both tasks	3.065(2.58)
First task	7.261(5.67)
Second task	6.261(4.96)

Table 1: Mean (S.D.) of #slider tuning in each condition

To answer the second sub-question, we analyze users' search logs. We find that users often perform several rounds of fine-tuning of facet-values before they finally marked a candidate. To analyze the participants' tuning behaviors, we define a facet-value tuning distance (d_i) to measure the efforts for each round of facet tuning. In this paper, we only analyze the distances between two consecutive facet-value configurations. Suppose that user has facet-value configuration (F_i) at the round i, and he/she changes it into (F_{i+1}) in the next round. F_i is a triple $<w_{ci}, w_{pi}, w_{si}>$ that indicates users' settings of three sliders at round i. d_i (see Formula 2) is defined as the Euclidean distance between F_i and F_{i+1}. We also compute the distance for each facet, in which we only use the value from one facet (either w_s, w_p or w_c).

$$d_i = \sqrt{(w_c^i - w_c^{i+1})^2 + (w_p^i - w_p^{i+1})^2 + (w_s^i - w_s^{i+1})^2} \tag{2}$$

After computing all d_i, we plot them against the round i in Figure 2, from which we find that participants change the facet values primarily in the 1st round and conduct relatively smaller tuning in the follow-up rounds. We also observe a fluctuation of facet-values in the later rounds, which may indicate the second attempts for tuning. The same trend is observed for the aggregated over three facets and each separate facet. Note that not every query receives 8 rounds of facet-value tuning. Therefore, we count the percentage of queries completed in each round, where we find that around 40% of queries are completed in the 2nd round, which is because it has a reasonable good nDCG but also in an early stage. For those that have more than 2 times of tuning may be difficult queries and thus may have relatively low nDCG values.

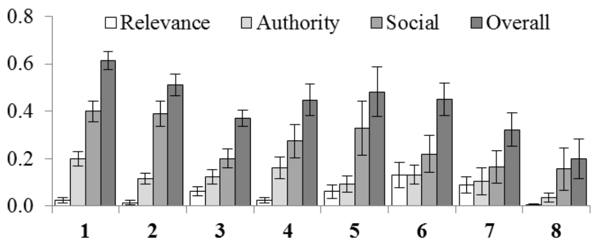

Figure 2: user tuning distance d_i (S.E.) over round i

3.2 RQ2: Can users reach the optimal facet-value settings?

When users marking a candidate, our system logs the query q_t and corresponding facet-value pairs F_i. In this section, we compare the search utility (measured by nDCG) in two conditions - one is without facet tuning (None) and the other is using user-tuned facet-values (Tuned). The purpose is to demonstrate the usefulness of facet tuning. We also consider a third condition with the best achievable nDCG for the given query q_i (Optimal). This is obtained through an exhaustive search over all combinations of three slider facets (values for each facet was divided uniformly into 20 parts) and the largest nDCG is chosen to denote the optimal one. We find that (from Table 2) the tuning for sliders indeed brings significant nDCG promotion, but the increase has strong task effect as there are only significant increase on task 3 and 4, but no significant difference on task 1 and 2. And users are able to reach a reasonable facet-value but are also hard to tune for the optimal one, which indicates the necessity of providing system-mediated suggestions for facet-value configurations.

	Overall	Task 1	Task 2	Task 3	Task 4
None	0.3947	0.5217	0.3156	0.3064	0.4177
Tuned	**0.4327**↑	0.4899	0.3121	**0.4321**↑	0.4950*
Optimal	***0.6917***↑	***0.7578***↑	***0.6458***↑	***0.6760***↑	***0.6784***↑

Table 2: nDCG comparisons on three conditions. Number in bold (italic) means significance over None (Tuned) using Wilcoxon Sign Test. *:p-value=0.054

4 Conclusions and Future Work

Faceted people search systems that provide user control over search facets can be used to facilitate exploratory people search tasks. Most researches of faceted search only consider the facets with discrete values; while the slider facets with continuous values are rarely explored. This paper analyses users' behaviors of using the sliders. Based on a lab-controlled user study, we find: 1) users consistently interact with the slider facets during their search processes, and there are many fine-tuning behaviors for achieving a better results ranking; 2) while users are able to configure reasonable facet-values that are better than having no tuning, they are still far behind the optimal facet-values. It is important to provide system-mediated suggestions for the facet-values tuning. In the future, we would like to explore provide system-mediated suggestions for the facet-values configurations when users cannot reach the optimal facet-values configurations.

References

Balog, K., Azzopardi, L., & De Rijke, M. (2006). Formal models for expert finding in enterprise corpora. In *Proceedings of the 29th annual international acm sigir conference on research and development in information retrieval* (pp. 43–50).

Ben-Yitzhak, O., Golbandi, N., Har'El, N., Lempel, R., Neumann, A., Ofek-Koifman, S., ... Yogev, S. (2008). Beyond basic faceted search. In *Proceedings of the 2008 international conference on web search and data mining* (pp. 33–44).

Han, S., He, D., Brusilovsky, P., & Yue, Z. (2013). Coauthor prediction for junior researchers. In *Social computing, behavioral-cultural modeling and prediction* (pp. 274–283). Springer.

Han, S., He, D., Jiang, J., & Yue, Z. (2013). Supporting exploratory people search: a study of factor transparency and user control. In *Proceedings of the 22nd acm international conference on conference on information & knowledge management* (pp. 449–458).

Haveliwala, T. H., Jeh, G. M., & Kamvar, S. D. (2014, October 28). *Variable personalization of search results in a search engine.* Google Patents. (US Patent 8,874,567)

Hearst, M. (2006). Design recommendations for hierarchical faceted search interfaces. In *Acm sigir workshop on faceted search* (pp. 1–5).

Hofmann, K., Balog, K., Bogers, T., & De Rijke, M. (2010). Contextual factors for finding similar experts. *Journal of the American Society for Information Science and Technology*, *61*(5), 994–1014.

Huang, S.-W., Tunkelang, D., & Karahalios, K. (2014). The role of network distance in linkedin people search. In *Proceedings of the 37th international acm sigir conference on research & development in information retrieval* (pp. 867–870).

Karimzadehgan, M., Zhai, C., & Belford, G. (2008). Multi-aspect expertise matching for review assignment. In *Proceedings of the 17th acm conference on information and knowledge management* (pp. 1113–1122).

Koren, J., Zhang, Y., & Liu, X. (2008). Personalized interactive faceted search. In *Proceedings of the 17th international conference on world wide web* (pp. 477–486).

Kules, B., Capra, R., Banta, M., & Sierra, T. (2009). What do exploratory searchers look at in a faceted search interface? In *Proceedings of the 9th acm/ieee-cs joint conference on digital libraries* (pp. 313–322).

Liberman, S., & Lempel, R. (2012). Approximately optimal facet selection. In *Proceedings of the 27th annual acm symposium on applied computing* (pp. 702–708).

Parra, D., Brusilovsky, P., & Trattner, C. (2014). See what you want to see: visual user-driven approach for hybrid recommendation. In *Proceedings of the 19th international conference on intelligent user interfaces* (pp. 235–240).

Spirin, N. V., He, J., Develin, M., Karahalios, K. G., & Boucher, M. (2014). People search within an online social network: Large scale analysis of facebook graph search query logs. In *Proceedings of the 23rd acm international conference on conference on information and knowledge management* (pp. 1009–1018).

Tang, W., Tang, J., & Tan, C. (2010). Expertise matching via constraint-based optimization. In *Web intelligence and intelligent agent technology (wi-iat), 2010 ieee/wic/acm international conference on* (Vol. 1, pp. 34–41).

Tunkelang, D. (2009). Faceted search. *Synthesis lectures on information concepts, retrieval, and services*, *1*(1), 1–80.

Weerkamp, W., Berendsen, R., Kovachev, B., Meij, E., Balog, K., & De Rijke, M. (2011). People searching for people: Analysis of a people search engine log. In *Proceedings of the 34th international acm sigir conference on research and development in information retrieval* (pp. 45–54).

Yarosh, S., Matthews, T., & Zhou, M. (2012). Asking the right person: supporting expertise selection in the enterprise. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 2247–2256).