

# IntentStreams: Smart Parallel Search Streams for Branching Exploratory Search

Salvatore Andolina<sup>1</sup>, Khalil Klouche<sup>1,4</sup>, Jaakko Peltonen<sup>2,3</sup>, Mohammad Hoque<sup>1</sup>, Tuukka Ruotsalo<sup>2</sup>, Diogo Cabral<sup>1</sup>, Arto Klami<sup>1</sup>, Dorota Glowacka<sup>1</sup>, Patrik Floréen<sup>1,2</sup>, and Giulio Jacucci<sup>1,2</sup>

Helsinki Institute for Information Technology HIIT

<sup>1</sup>University of Helsinki, Department of Computer Science, Finland

<sup>2</sup>Aalto University, Finland <sup>3</sup>University of Tampere, School of Information Sciences, Finland

<sup>4</sup>Aalto University, Media Lab, Finland

## ABSTRACT

The user's understanding of information needs and the information available in the data collection can evolve during an exploratory search session. Search systems tailored for well-defined narrow search tasks may be suboptimal for exploratory search where the user can sequentially refine the expressions of her information needs and explore alternative search directions. A major challenge for exploratory search systems design is how to support such behavior and expose the user to relevant yet novel information that can be difficult to discover by using conventional query formulation techniques. We introduce IntentStreams, a system for exploratory search that provides interactive query refinement mechanisms and parallel visualization of search streams. The system models each search stream via an intent model allowing rapid user feedback. The user interface allows swift initiation of alternative and parallel search streams by direct manipulation that does not require typing. A study with 13 participants shows that IntentStreams provides better support for branching behavior compared to a conventional search system.

## Author Keywords

User interface design; information exploration; parallel browsing.

## ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

## INTRODUCTION

Exploratory search activities confront users with problems in formulating queries and identifying directions for information exploration. Studies show that searchers tend to perform more than one task simultaneously: approximately 75% of submitted queries involve a multitasking activity [22, 21]. Users engage in multitask search with and without parallel

browsing, but parallel browsing is a common activity and more prevalent than linear browsing. In parallel browsing, also called branching [13], users visit web pages in multiple concurrent threads, for example, by opening multiple tabs or windows in web browsers [14]. Branching in browsing has been studied extensively [6], but little has been done to support nonlinear and parallel browsing. Recent visual search user interfaces have shown the effectiveness of interacting visually with query elements, however, there are no solutions to support fluid branching and parallel search.

We introduce IntentStreams, a system supporting parallel browsing and branching during search without the need to open new tabs. It presents parallel streams of searches, where each stream shows a list of resulting documents and keywords, and a display of the underlying queries as keywords representing the search intent of the stream. New streams are initiated by the user, where the search intent of a new stream is initialized by typing a traditional query or by dragging keywords available in any of the streams. In each stream, in addition to the user-chosen keywords, the system proposes other relevant keywords and orders them vertically by their predicted relevance. The users can change the relative relevance of keywords in the query intent of each stream and branch new streams by simply dragging keywords. IntentStreams was tested using 25 million news articles crawled from public news sources in a comparative study with 13 subjects.

The experimental results show that IntentStreams better supports parallel search and branching behavior when compared to a conventional search system.

**Exploratory search** is commonly distinguished from lookup search and has been described as combining exploratory browsing and focused searching [23, 5]. These and other information seeking frameworks [23, 27] partially overlap, emphasizing different aspects but similarly trying to characterize iterative sense making and refinement of search intents. As a response, personalization techniques have been developed to support query formulation and relevance feedback. Several techniques exist for supporting query formulation or for processing results to help re-rank, filter [31] or expand the query [8, 26]. For example, relevance feedback [16] and query or term suggestions [17] can be effective in short-term navigational search, but give limited support for exploratory

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [Permissions@acm.org](mailto:Permissions@acm.org).

*IUI 2015*, March 29–April 1, 2015, Atlanta, GA, USA.

Copyright © 2015 ACM 978-1-4503-3306-1/15/03 ...\$15.00.

<http://dx.doi.org/10.1145/2678025.2701401>

search. However, increasing evidence in information retrieval research suggests relevance feedback in exploratory search often leads to a context trap where after a few iterations of feedback users have specified their context so strictly the system is unable to propose anything new [12, 23].

**Visual search interfaces.** Recently, search systems use visualization of the resulting information for faster relevance judgment and effective feedback [11, 15, 20, 24, 32]. Several approaches to visualize results have been explored, including multiple linked lists, scatter plots, graphs and their combinations [30, 20]. Such visual search systems are distinguished from traditional query composition ones through their emphasis on rapid filtering to reduce result sets, progressive refinement of search parameters, continuous reformulation of goals, and visual scanning to identify results [1]. Visual approaches provide better support for exploration in a variety of ways: supporting sense making by incrementally and interactively exploring the network of data [7], showing how visualization supports user involvement in the recommendation providing rationale behind suggested items [34], visualizing relations of different queries and result sets [2], supporting entity search on touch devices [18]. Recent work shows how users can view and manipulate their search intent models provided by the system [2, 3, 29, 9, 28]. However, there is still a lack of solutions that specifically address branching behavior while supporting query formulation.

## INTENTSTREAMS

A user performing exploratory search typically needs to steer the exploration by query refining or reformulating as she makes better sense of the information space. We identified two challenges impairing a proper steering support:

1. Supporting information comparison from several queries: Traditional search systems show one set of results at a time. To compare information from distinct queries, one must rely on memory or complicated browser windows arrangements.
2. Formulating queries in unknown areas: Exploratory search tasks usually require the user to iteratively formulate queries in a field she is unfamiliar with. To progress with the search, the user needs to go through long lists of results with low information gain to acquire new terms and notions to reformulate the query.

We tackle these challenges through a combination of a user interface supporting parallel search streams, and user intent modeling.

### The User Interface

IntentStreams provides a unique horizontally scrollable workspace divided in two areas: the keywords area at the bottom and the results area on top (Figure 1). By clicking the workspace, the user is prompted to type a first query. The system returns a list of relevant documents in the results area and a set of related keywords in the keywords area. Keywords are positioned vertically by weight and horizontally by topic proximity. The vertical arrangement is called a stream and can be easily manipulated, modified and refreshed. The content of a document can be seen by clicking the title. A click

and hold on a document highlights keywords directly related to it. A click and hold on a keyword highlights related documents. By moving keywords vertically, the user can change their weight; by hitting the refresh button, the stream then updates and presents a new set of documents and keywords. New parallel streams can be created by clicking next to an existing stream and typing a new query, or simply by dragging a keyword outside of its stream. Since the workspace is horizontally scrollable, the amount of parallel streams a user can create is limited only by computer memory. The amount of parallel streams that can be shown simultaneously is determined by the display resolution. Streams can be dragged and rearranged. A button lets the user delete streams.


### Interactive Intent Model

For each search stream, the interactive intent model is similar to the model in a previous non-parallel system [29] and has two parts: a model for retrieval of documents, and a model for estimating the user’s search intent (relevance of keywords to the user’s information need). We describe both below.

**Document retrieval model.** For each stream, we estimate a relevance ranking where documents are ranked by their probability given the intent model for the stream. We use a unigram language model. The intent model yields a vector  $\hat{\mathbf{v}}$  with a weight  $\hat{v}_i$  for each keyword  $k_i$ . The  $\hat{\mathbf{v}}$  is treated as a sample of a desired document. Documents  $d_j$  are ranked by probability to observe  $\hat{\mathbf{v}}$  as a sample from the language model  $M_{d_j}$  of  $d_j$ . Maximum likelihood estimation yields  $\hat{P}(\hat{\mathbf{v}}|M_{d_j}) = \prod_{i=1}^{|\hat{\mathbf{v}}|} (\hat{P}_{mle}(k_i|M_{d_j}))^{\hat{v}_i}$ . We regularize probabilities  $\hat{P}_{mle}(k_i|M_{d_j})$  in  $d_j$  towards overall keyword proportions in the corpus by Bayesian Dirichlet smoothing. In each stream the  $d_j$  are ranked by  $\alpha_j = \hat{P}(\hat{\mathbf{v}}|M_{d_j})$ . To expose the user to more novel documents we sample a document set from the ranking and show them in rank order. As in [9, 29] we use Dirichlet Sampling based on the  $\alpha_j$ , and favor documents whose keywords got positive feedback by increasing their  $\alpha_j$  [10].

**User intent model.** For each stream, the intent model estimates relevance of keywords from feedback to keywords. For a stream launched by a typed query, we use the query with weight 1 as the initial intent model; for a stream launched by dragging a keyword we use the keyword with weight 1. The user gives feedback as relevance scores  $r_i \in [0, 1]$  for a subset of  $J$  keywords  $k_i, i = 1, \dots, J$  in the stream;  $r_i = 1$  means  $k_i$  is highly relevant and the user wishes to direct the stream in that direction, and  $r_i = 0$  means  $k_i$  is of no interest.

Let  $\mathbf{k}_i$  be binary  $n \times 1$  vectors telling which of the  $n$  documents  $k_i$  appeared in; to boost documents with rare keywords we convert the  $\mathbf{k}_i$  to *tf-idf* representation. We estimate the expected relevance  $r_i$  of a keyword  $k_i$  as  $\mathbb{E}[r_i] = \mathbf{k}_i^\top \mathbf{w}$ . The vector  $\mathbf{w}$  is estimated from user feedback by the LinRel algorithm [4]. In each search iteration, let  $k_1, \dots, k_p$  be the keywords for which the user gave feedback so far, let  $\mathbf{K} = [\mathbf{k}_1, \dots, \mathbf{k}_p]^\top$  be the matrix of their feature vectors, and let  $\mathbf{r}^{feedback} = [r_1, r_2, \dots, r_p]^\top$  be their relevance scores from the user. LinRel estimates  $\hat{\mathbf{w}}$  by solving  $\mathbf{r}^{feedback} = \mathbf{K}\mathbf{w}$ , and estimates relevance score for each  $k_i$  as  $\hat{r}_i = \mathbf{k}_i^\top \hat{\mathbf{w}}$ .



**Figure 1. a.** The first query (in this case mobile phone”) returns a search stream composed of news articles most relevant to the query, as well as a set of most relevant keywords extracted from a larger set of related articles. **b.** The user can modify the weight of the keywords by sliding them vertically, after which the stream will refresh, updating articles and keywords accordingly. If dropped outside their initial stream, keywords can either trigger a new search stream or be passed to an already existing parallel stream.

To expose the user to novel keywords, in each stream we show keywords  $k_i$  not with highest  $\hat{r}_i$ , but with highest upper confidence bound for relevance, which is  $\hat{r}_i + \alpha\sigma_i$ , where  $\sigma_i$  is an upper bound on standard deviation of  $\hat{r}_i$ , and  $\alpha > 0$  is a constant for adjusting the confidence level. In each iteration, we compute  $\mathbf{s}_i = \mathbf{K}(\mathbf{K}^T\mathbf{K} + \lambda\mathbf{I})^{-1}\mathbf{k}_i$  where  $\lambda$  is a regularization parameter, and show the  $k_i$  maximizing  $\mathbf{s}_i^T \mathbf{r}^{feedback} + \frac{\alpha}{2} \|\mathbf{s}_i\|$  representing estimated search intent. We optimize horizontal positions of the shown  $k_i$  by dimensionality reduction [33];  $k_i$  get similar positions if their relevance estimate changes similarly with respect to a set of additional feedback.

## EVALUATION

We evaluated the system to find out if and how IntentStreams supports parallel browsing and branching behavior. IntentStreams was compared against a baseline system with an interface similar to a traditional Google search interface. Our hypothesis was that, compared to the baseline, IntentStreams generates (1.) more parallel streams, (2.) more revisits, and (3.) more branches. We used the following metrics: number of parallel streams, number of revisits, and number of branches. In the baseline, the number of parallel streams denotes the number of tabs opened, a revisit indicates returning to an already open tab, and a branch denotes a query updated after a revisit. In IntentStreams, a revisit occurs when a user performs certain activities (opening an article, weight change) on a previously created stream. A branch occurs when a new stream is created from an existing one. That includes both creating a new query by dragging a keyword or updating the existing stream by modifying the weights of its keywords.

## Method

We evaluated the system with 13 volunteers (4 female). The participants’ age ranged from 19 to 36 with mean of 28.4 ( $SD = 4.05$ ). Their levels of education were: 8% PhD, 46% Master, 38% Bachelor, 8% High School. Each participant received two movie tickets for their participation. We used a within-subject design, where participants were asked to perform two tasks, one with IntentStreams and one with the baseline. We counterbalanced by changing the order in

which the two tasks were performed and the order in which the two systems were used.

The task was set in an essay writing scenario and formulated as follows: *You have to write an essay on recent developments of X where you have to cover as many subtopics as possible. You have 20 minutes to collect the material that will provide inspiration for your essay. You have additional 5 minutes to write your essay.* The two tasks performed by the participants covered two topics: (1.) NASA, and (2.) China Mobile.

Experiments were run in a laboratory on a laptop with OS X operating system. Each participant signed a consent form. To determine the eligibility, we asked candidates how familiar they were with each chosen topic on a 1-5 scale, where 1 means “no knowledge” and 5 means “expert knowledge”. Only those with a score lower than 3 were considered eligible. Before the experiment, participants received detailed instructions and performed a 5-minute training session.


To evaluate the system, we connected it to a news repository of English language editorial news articles crawled from publicly available news sources from September 2013 to March 2014. The database contains more than 25 million documents. The documents were originally collected for monitoring media presence of numerous interested parties, and hence the collection has wide topical coverage. All the documents were preprocessed by the Boilerpipe tool [19] and the keyphrases were extracted with the Maui toolkit [25].

The baseline system was connected to the same news repository. In the baseline system, users could type queries and receive a list of relevant news articles. To start a new parallel query, a new tab had to be opened.

## FINDINGS

Table 1 shows the results of the log analysis. In the 20-minute long sessions, IntentStreams on average generated 7.84 more queries ( $SD = 7.27$ ), 6.38 more parallel streams ( $SD = 4.03$ ), 4.54 more revisits ( $SD = 4.52$ ), and 3.62 more branches ( $SD = 4.01$ ). A paired t-test indicates that all those differences are statistically significant ( $p < 0.01$ ).

Participant 11 – Baseline (BL) – Topic : NASA



Vertically : rows = parallel streams  
 Horizontally : Chronological sequence of user actions  
 Legend :  
 T Text entry  
 R Revisit  
 B Branch  
 A Search triggered in new query  
 B Search triggered in modified query  
 n Article opened (n = amount if > 1)  
 ♦ Query modified (keywords weight)  
 ⋮ Keyword dragged

Participant 11 – IntentStreams (IS) – Topic : China Mobile




Figure 2. Example of branching behavior from the case study: top - Baseline; bottom - IntentStreams.

**Parallel search supported in IntentStreams.** Results show that users created more parallel streams than opened new tabs. While the system allows the creation of parallel streams, the users revisit earlier ones consistently, which denotes parallel search behavior. In fact, revisits are higher in the IS condition.

**Branching supported in IntentStreams.** In IntentStreams, more queries and parallel streams were created through branching. Figure 2 presents a visual representation of a participant’s search behavior, showing the difference between the linear search behavior in the baseline and the more articulated search behavior in IntentStreams.

Further, IntentStreams supports more exploration. In IS, more exploration of the information space was done as can be seen from the higher number of queries.

**CONCLUSIONS**

We introduced the IntentStreams system for exploratory search of news based on parallel visualization of smart search streams. It models each search stream by an intent model, allows rapid tuning by feedback to keywords, and allows rapid initiation of new streams by keyword interaction without typing. Initial experiments show that users take advantage of the rich parallel search opportunities and engage in much stronger parallel browsing and branching behavior than in a traditional system.

This is an important finding as current browsing and searching behavior is already characterized by multitask search (in the same query field users alternate tasks [22, 21]), parallel browsing (users browse on parallel tabs or windows [14]), and branching (a new tab or window is created from a link or result of a previous window or tab [13]). Branching has

Table 1. Comparison between IntentStreams (IS) and the baseline (BL). The number of queries, parallel streams, revisits, and branches, for each participant P1,...,P13.

	queries		par. streams		revisits		branches	
	BL	IS	BL	IS	BL	IS	BL	IS
P1	5	5	2	5	0	7	0	1
P2	5	4	1	2	1	0	1	0
P3	7	17	1	12	1	6	1	4
P4	9	11	1	6	0	2	0	2
P5	14	18	1	12	0	9	0	7
P6	1	8	1	5	0	0	0	0
P7	12	18	7	14	5	7	2	3
P8	22	26	3	12	6	15	1	8
P9	6	18	6	12	4	4	0	0
P10	8	11	7	11	6	5	0	0
P11	8	21	7	11	7	12	0	8
P12	16	35	11	14	3	14	1	9
P13	3	26	3	18	0	11	0	11

been shown to be more important in informational browsing than navigational search [14]. The approach proposed in IntentStreams can be incorporated into other search interfaces to provide an effective way to branch search.

**ACKNOWLEDGMENTS**

This work has been partially supported by TEKES (Re:Know and DIGILE programme D2I) and the Academy of Finland (decisions 251170, 252845, 266969, 278090, and 283311).

**REFERENCES**

1. Ahlberg, C., and Shneiderman, B. Visual information seeking: Tight coupling of dynamic query filters with

- starfield displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (New York, NY, USA, 1994), 313–317.
2. Ahn, J., and Brusilovsky, P. Adaptive visualization for exploratory information retrieval. *Information Processing & Management* 49, 5 (2013), 1139–1164.
  3. Ahn, J.-w., and Brusilovsky, P. Adaptive visualization of search results: Bringing user models to visual analytics. *Information Visualization* 8, 3 (2009), 167–179.
  4. Auer, P. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research* 3 (2002), 397–422.
  5. Bates, M. J. What is browsing—really? a model drawing from behavioural science research. *Inf. Res.* 12, 4 (2007).
  6. Bilenko, M., and White, R. W. Mining the search trails of surfing crowds: identifying relevant websites from user activity. In *Proceedings of the 17th international conference on World Wide Web*, ACM (2008), 51–60.
  7. Chau, D. H., Kittur, A., Hong, J. I., and Faloutsos, C. Apolo: Making sense of large network data by combining rich user interaction and machine learning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, ACM (New York, NY, USA, 2011), 167–176.
  8. Chirita, P.-A., Firan, C. S., and Nejdl, W. Personalized query expansion for the web. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, ACM (2007), 7–14.
  9. Głowacka, D., Ruotsalo, T., Konuyshkova, K., Kaski, S., Jacucci, G., et al. Directing exploratory search: Reinforcement learning from user interactions with keywords. In *Proceedings of the 2013 international conference on Intelligent user interfaces*, ACM (2013), 117–128.
  10. Głowacka, D., and Shawe-Taylor, J. Content-based image retrieval with multinomial relevance feedback. In *Proc. of ACML* (2010), 111–125.
  11. Havre, S., Hetzler, E., Perrine, K., Jurrus, E., and Miller, N. Interactive visualization of multiple query results. In *Proceedings of the IEEE Symposium on Information Visualization 2001*, IEEE Computer Society (Washington, DC, USA, 2001), 105.
  12. Hearst, M. A. *Search User Interfaces*, 1st ed. Cambridge University Press, New York, NY, USA, 2009.
  13. Huang, J., Lin, T., and White, R. W. No search result left behind: Branching behavior with browser tabs. In *Proceedings of the fifth ACM international conference on Web search and data mining*, ACM (2012), 203–212.
  14. Huang, J., and White, R. W. Parallel browsing behavior on the web. In *Proceedings of the 21st ACM conference on Hypertext and hypermedia*, ACM (2010), 13–18.
  15. Käki, M. Findex: Search result categories help users when document ranking fails. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (New York, NY, USA, 2005), 131–140.
  16. Kelly, D., and Fu, X. Elicitation of term relevance feedback: An investigation of term source and context. In *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, ACM (New York, NY, USA, 2006), 453–460.
  17. Kelly, D., Gyllstrom, K., and Bailey, E. W. A comparison of query and term suggestion features for interactive searching. In *Proceedings of the 32nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, ACM (New York, NY, USA, 2009), 371–378.
  18. Klouche, K., Ruotsalo, T., Cabral, D., Andolina, S., Bellucci, A., and Jacucci, G. Designing for exploratory search on touch devices. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI 2015, ACM (2015).
  19. Kohlschütter, C., Frankhauser, P., and Nejdl, W. Boilerplate detection using shallow text features. In *Proceedings of the Third ACM International Conference on Web Search and Data Mining*, ACM (2010).
  20. Kules, W., Wilson, M. L., Schraefel, M. C., and Shneiderman, B. From keyword search to exploration: How result visualization aids discovery on the web. Tech. rep., University of Southampton, 2008.
  21. Lucchese, C., Orlando, S., Perego, R., Silvestri, F., and Tolomei, G. Identifying task-based sessions in search engine query logs. In *Proceedings of the fourth ACM international conference on Web search and data mining*, ACM (2011), 277–286.
  22. Lucchese, C., Orlando, S., Perego, R., Silvestri, F., and Tolomei, G. Discovering tasks from search engine query logs. *ACM Transactions on Information Systems (TOIS)* 31, 3 (2013), 14.
  23. Marchionini, G. Exploratory search: From finding to understanding. *Communications of the ACM* 49, 4 (Apr. 2006), 41–46.
  24. Matejka, J., Grossman, T., and Fitzmaurice, G. Citeology: Visualizing paper genealogy. In *CHI '12 Extended Abstracts on Human Factors in Computing Systems*, ACM (New York, NY, USA, 2012), 181–190.
  25. Medelyan, O., Frank, E., and Witten, I. H. Human-competitive tagging using automatic keyphrase extraction. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing* (2009), 1318–1327.
  26. Micarelli, A., Gasparetti, F., Sciarrone, F., and Gauch, S. Personalized search on the world wide web. In *The Adaptive Web*. Springer, Berlin, Heidelberg, 2007, 195–230.

27. Pirolli, P., and Card, S. Information foraging in information access environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '95, ACM Press/Addison-Wesley Publishing Co. (New York, NY, USA, 1995), 51–58.
28. Ruotsalo, T., Jacucci, G., Myllymäki, P., and Kaski, S. Interactive intent modeling: Information discovery beyond search. *Communications of the ACM* 58, 1 (Dec. 2014), 86–92.
29. Ruotsalo, T., Peltonen, J., Eugster, M., Głowacka, D., Konyushkova, K., Athukorala, K., Kosunen, I., Reijonen, A., Myllymäki, P., Jacucci, G., et al. Directing exploratory search with interactive intent modeling. In *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*, ACM (2013), 1759–1764.
30. Stasko, J., Görg, C., and Liu, Z. Jigsaw: supporting investigative analysis through interactive visualization. *Information visualization* 7, 2 (2008), 118–132.
31. Teevan, J., Dumais, S. T., and Liebling, D. J. To personalize or not to personalize: modeling queries with variation in user intent. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, ACM (2008), 163–170.
32. Terveen, L., Hill, W., and Amento, B. Constructing, organizing, and visualizing collections of topically related web resources. *ACM Transactions on Computer-Human Interaction* 6, 1 (1999), 67–94.
33. Venna, J., Peltonen, J., Nybo, K., Aidos, H., and Kaski, S. Information retrieval perspective to nonlinear dimensionality reduction for data visualization. *Journal of Machine Learning Research* 11 (2010), 451–490.
34. Verbert, K., Parra, D., Brusilovsky, P., and Duval, E. Visualizing recommendations to support exploration, transparency and controllability. In *Proceedings of the 2013 International Conference on Intelligent User Interfaces*, ACM (2013), 351–362.