

Flexible Entity Search on Surfaces

Tuukka Ruotsalo², Khalil Klouche^{1,3}, Diogo Cabral^{1,4}
Salvatore Andolina², Giulio Jacucci^{1,2}

¹Helsinki Institute for Information Technology HIIT, Department of Computer Science,
University of Helsinki, Finland

²Helsinki Institute for Information Technology HIIT, Department of Computer Science, Aalto University,
Finland

³Aalto University, School of Arts, Design and Architecture, Finland

⁴Madeira-ITI, University of Madeira, Portugal

¹first.last@helsinki.fi, ²first.last@aalto.fi ⁴first.last@staff.uma.pt

ABSTRACT

Surface computing allows flexible search interaction where users can manipulate the representation of entities recommended for them to create new queries or augment existing queries by taking advantage of increased screen estate and almost physical tactile interaction. We demonstrate a search system based on 1) Direct Manipulation of Entity Representation on Surfaces and 2) Entity Recommendation and Document Retrieval. Entities are modeled as a knowledge-graph and the relevances of entities are computed using the graph structure. Users can manipulate the representation of entities via spatial grouping and assigning preferences on entities. Our contribution can help to design effective information exploration systems that take advantage of large surfaces.

Author Keywords

Entity search; Surface computing; Information exploration

ACM Classification Keywords

H.5.2. Information Interfaces and Presentation: User Interfaces; H.3.3. Information Storage and Retrieval: Information Search and Retrieval

INTRODUCTION

Information seeking is no longer limited to conventional desktop settings, but search activities are increasingly carried out using touch surfaces available in mobile devices, but also public collaborative spaces. For example, teachers seeking supporting materials in a class room or visitors exploring information about artworks, artists, or exhibitions at museums. Surface computing technologies (i.e., interactive screens) hold great potential for enhancing information seeking activities as they are not limited to the conventional input or presentation techniques.

Surface computing has been recently utilized for open-ended information exploration [6, 8, 10] and surface interaction have been shown to increase the amount of interaction, encourage reflection with the presented information, and foster effective and more intuitive interaction designs [7, 21]. Morris et. al. introduce surface computing search specifically for co-located collaboration [13, 14] as surface computing few years ago was generally a synonym for tabletop collaboration. However, now large surfaces and multitouch input have become central to individual users on a variety of screen form factors including smart phones, tablets and large format touch screens.

Several visual approaches for desktop set ups attempt to better support search in different ways: supporting sense-making by incrementally and interactively exploring the network of data [3, 19], showing how visualization supports user involvement in the recommendation, providing rationale behind suggested items [5, 20, 22], and visualizing relations of different queries and result sets [1]. The present article extends this line of work on using touch interfaces [11] to utilize large surfaces [8, 14] applying entity search [17, 18, 16]. We report an initial user study demonstrating the usability and user acceptance of a system implementing entity search on large surfaces.

Designing intelligent search user interfaces for surfaces face two major challenges: 1) allowing users with flexible query construction to augment text entry input [23] and 2) taking advantage of the benefits and attractiveness of the direct, almost physical, interaction of a tactile interface [13].

Such user interface and interaction designs are changing users' expectations on how search systems should function; the single search-box and result listing paradigm is being increasingly challenged by more intelligent user interfaces that can help users to supply additional input data beyond the keyword query. There is also increasing evidence that users adopt these practices, if using them is made sufficiently effortless for the users to do so, for example, through specialized interfaces or query assistance services [20].

Entity oriented search, as opposed to conventional document search, is becoming a key enabling component for next generation search user interfaces [4, 18]. Query construction can be effectively supported via enriching the search result page with entities to enable direct exploration points

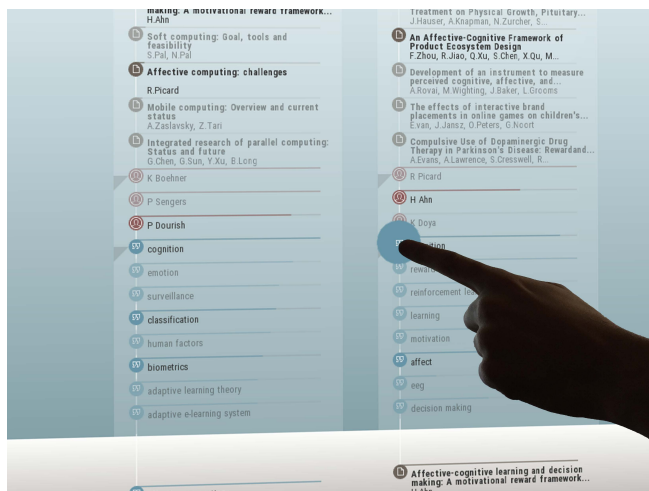
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) 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.

MUM '16, Rovaniemi, Finland, December 12–15, 2016


Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-4860-7/16/1...\$15.00

DOI: <http://dx.doi.org/10.1145/3012709.3012732>



(a) Search streams can be directly manipulated



(b) Multi-touch surfaces allow collaborative search

Figure 1: a) Search results and entity recommendations are visualized for interaction as search streams on a large interactive surface. The interface enables users to directly interact with entities and search results by spatially organizing them on a surface to create new search streams or augment existing streams. b) Multi-touch surfaces allow collaborative search.

to novel and serendipitous results [2]. For example, a user searching for “robotics” (Figure 2) could immediately recognize relevant entities related to her query, such as “mobile robots”, “human-robot interaction” or “social robots” to continue the search and exploration process. In this way entity oriented search is able to bridge the gap between unstructured keyword-search and structured knowledge [12]. We demonstrate the design and implementation of the resulting prototype application (Figure 1). The prototype indexes over 50 million scientific documents and the associated entity information, such as authors, publication venues, and other article metadata. The user interface allows spatial organization of entities on the surface that can be used to compose several parallel search streams. Users can create new queries or augment existing ones by direct manipulation of entity representation by dragging them to different spatial areas.

ENTITY SEARCH ON SURFACES

Our entity search approach is based on 1) Direct Manipulation of Entity Representation and 2) Entity Recommendation and Document Retrieval. These allow users to retrieve relevant information, make sense of the result space, and flexibly continue the search by directly manipulating entity representation and using them as parts of queries and search streams.

Direct Manipulation of Entity Representation

The user interface (Figure 2) is horizontally divided in two areas: the *Results Area* on top, and the lower part which is called the *Dock*, in which the user composes queries. The entity representation can be directly manipulated by using the following interaction operations:

1. **Query Augmentation** - The user is able to compose queries or augment existing queries by moving corresponding entities in the same area on the Dock of the surface

(Figure 2, left). The user can directly manipulate the entity representation by dragging one or several entities from the *Results Area* to the *Dock* (Figure 2, center). A query always consists of a single entity or a group of entities, each of which must be close to each other. The unity of a query set - along with its boundaries - are visualized by a variation of the background color. If an entity is dragged within those boundaries, it will be part of the existing query, otherwise it will start a new search stream.

2. **Entity Inspection** - The user can inspect an entity or relationships between the entities by holding on an entity (Figure 3b). Tapping an entity will display a snippet next to the entity (Figure 3c).
3. **Search Streams** - The user can manage and visualize numerous search streams simultaneously. Each spatial group on the Dock will start a new search stream, but search streams can be also combined or separated by using specific gestures on the surface (Figure 2, right). We use a fixed distance threshold and two entities get associated if they are within a threshold distance from each other. The user can navigate horizontally across the search streams by scrolling the screen with a specific gesture (Figure 3a).

Entity Recommendation and Document Retrieval

The query augmentation that the users can perform are targeted to entities and enabled by the entity recommendation system. As the number of potential entities to be used to create queries are likely to be much larger than what can be presented for the user at a single iteration, ranking of the most important entities is an essential feature of our system. The ranking is based on the centrality of entities in a data retrieved for a specific search stream and the users feedback on entities

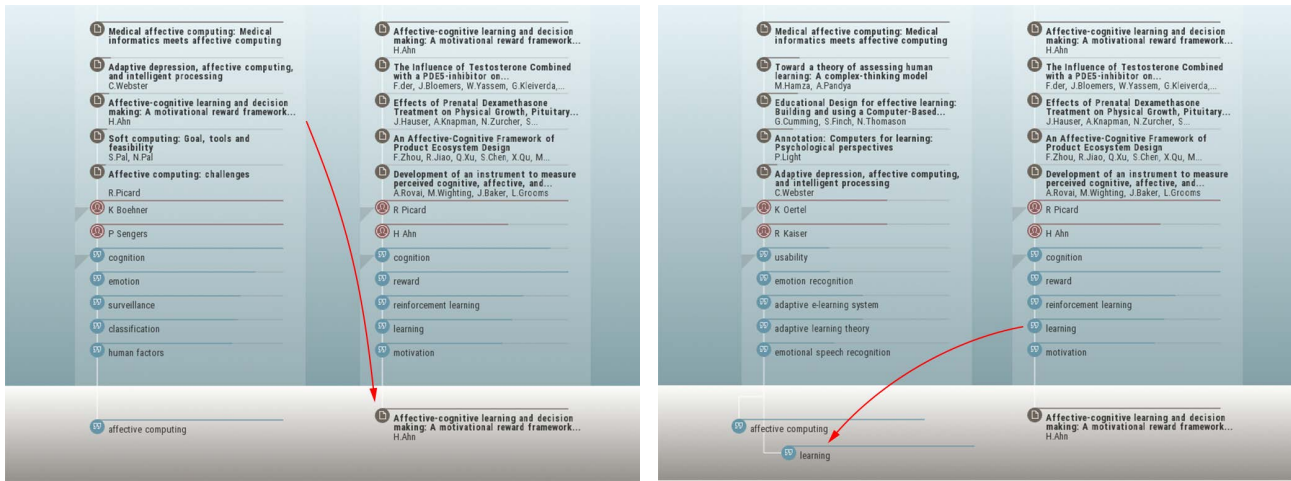


Figure 2: Searching using the prototype system via direct manipulation of entity representation. Left: The user is searching for “Affective Computing”. Above it, a first set of related entities and search results (documents, authors and keywords) appears. The user then chooses to investigate one particular article by directly manipulating the document entity “Affective-Cognitive Learning And Decision Making” and drags it to the dock to create a new search stream. Right: The user then decides to augment the query in the first search stream by dragging a keyword entity “learning” in the spatial area representing the search stream. The resulting documents and entity recommendations are presented in each search stream. The estimated relevance is visualized as the colored bar under each document and entity.

determined by dragging them into the dock. The search session is initiated by either typing a query using a virtual keyboard or tapping an empty area of dock for two seconds to initiate a query via speech input.

We model the data as an undirected and labeled knowledge-graph G consisting of a disjoint union set of entities and documents and the set of edges between the entities. Each entity in the graph is connected to a document (which is here considered also to be an entity) it describes. For example, an article written by a specific author has an author entity and an edge between the author entity and the article entity.

The user’s query consists of one or more entities or documents $e \in G$ dragged into the dock and defined to belong to a single search stream. These entities are called preference entities q in the graph, where $|q| = 1$ and q_j denotes the preference for entity j . We use the Personalized PageRank method [9] to compute the ranking of the entities. Computing the personalized PageRank vector v for a given preference nodes can be then formalized as follows. Let an individual entity be denoted as e , and by $I(e)$ and $O(e)$ denote the set of in-neighbors and out-neighbors of e in G respectively. Let A be the matrix corresponding to the entity graph G , where

$$A_{ij} = \frac{1}{|O_{ij} \cup I_{ij}|}$$

if the entity i links to the entity j or vice versa, and $A_{ij} = 0$ otherwise. For a given q , the Personalized PageRank equation can be written as

$$v = (1 - c)Av + cq,$$

where c is the teleportation rate. Following [9], we set $c = 0.15$. The solution v is a steady-state distribution of

random surfers, where a surfer teleports at each step to an entity e with probability $c \cdot q(e)$, or moves to a random neighbor otherwise. We compute the steady distribution by using the power iteration method with 50 iterations.

The solution v is the Personalized PageRank vector for the preference entities q . The weights of the v can now be directly used for ranking the entities. As our graph size would be hundreds of millions of nodes, the computation is not possible on-line for the complete graph. Therefore, we approximate the set of nodes to be included in the initial graph by using a language model approach of information retrieval [24] and select 3000 documents and the corresponding entities to be included in the graph at each iteration. The same method is used to rank the documents. The ranked entities and documents are then grouped by type, ranked within the groups, and visualized for the user.

The entity inspection can be conducted using the same computation. Any nodes that the user taps on the user interface are set as preference entities q . The top-k entities given the weighted shortest path between the entities in q and the entities already presented for the user are highlighted.

USER EXPERIMENT

We conducted an open-ended user study to gather feedback on the perceived usability of the system, interaction adequacy, perceived learnability, perceived control, and users’ general attitude toward the system.

Database

The system was indexed with over 50 million scientific documents. In this domain, authors, documents and keywords are obvious candidates for entities that the users may want to use to direct their search. These entity types were selected as a minimal representative set.

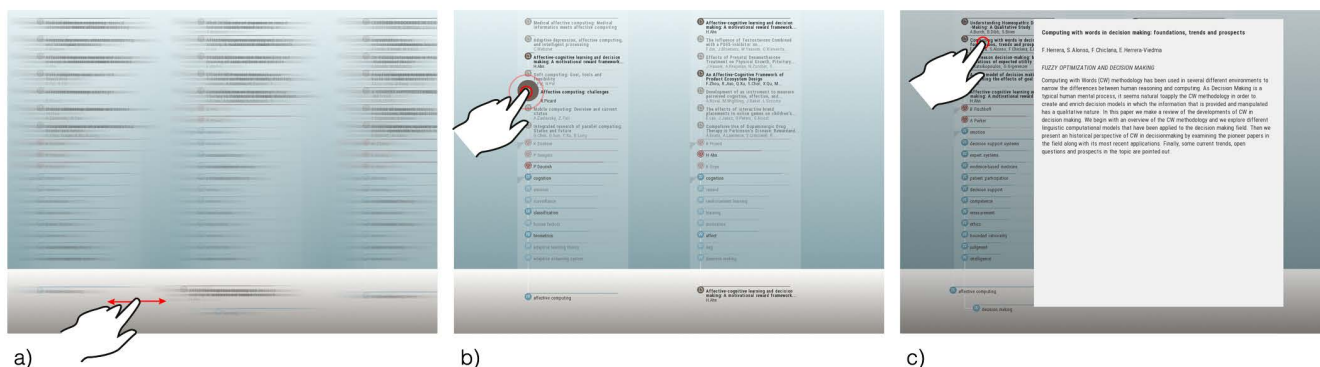


Figure 3: Additional interaction operations. a) Horizontal scrolling through unlimited amount of queries; b) Hold on one or several entities to highlight related entities and connections to documents; c) Single tap on a document title to reveal a content snippet. The main purpose of the multi-touch ability is to allow the user to perform such gestures simultaneously, like scrolling through the user space while moving one or several items.

Participants, Procedure, and Apparatus

We recruited six graduate students (5 male, with average of ages of $M = 30.67, \sigma = 6.68$). The participants were asked to conduct two exploratory search tasks with the system on their own topic of interest: to explore a range of relevant papers, authors, and keywords related to the selected topic. We limited the study to a solo search activity and the participants performed the tasks alone. The participants were allowed to use five minutes for each task and they were familiarized with the system before the actual task. The interactions with the system were logged. After the experiment, the participants answered selected questions from the ResQue questionnaire [15], designed for evaluating recommender systems. Then the participants were interviewed. The apparatus setup was equal to the one in Figure 1.

Results

The results from the ResQue questionnaires show that, on average, users had a general positive attitude towards the system ($M = 4.5/6, \sigma = 1.64$), which was considered clear ($M = 5.33/6, \sigma = 0.82$), and easy to learn ($M = 5.17/6, \sigma = 0.41$) and with an adequate graphical design ($M = 4.83/6, \sigma = 1.17$). Furthermore users felt in control when using the system ($M = 4.5/6, \sigma = 1.22$).

Interaction logs reveal that users actively used the entity recommendations provided by the system and constructed on average 6 queries ($\sigma = 3.54$), most of them to create new queries and search streams ($M = 4.18, \sigma = 1.99$). Since the number of queries on average is larger than than new queries to initiate search streams, we can conclude that users also modified previous search streams after creating new streams on an average of more 0.69 entities per created search stream ($\sigma = 0.56$).

Log analysis also showed a unanimous understanding of the usefulness of managing multiple parallel queries by creating new parallel queries ($M = 4.18, \sigma = 1.99$) and going back on average 2 queries ($\sigma = 1.26$) to pick up entities for new queries. We observed that participants were almost immediately able to take advantage of the system and knowingly steer their exploration.

In addition, interestingly, all our test users intuitively developed diverse and personal ways of browsing information, such as branching to many search streams from one interesting item, or physically pinning down interesting items while scrolling in search of related data.

From the analysis of exit interviews we found out that although users were satisfied about the overall experience, the perceived quality of recommended items was highly related with the familiarity of users with such items. This suggests that explanations and improved visualizations of the connections between the entities and the resulting documents could help in making sense of the search space and to direct the search.

DISCUSSION AND CONCLUSIONS

We demonstrated a search user interface designed for interactive search on surfaces where users can flexibly manipulate the representation of entities. Interaction via query augmentation, entity inspection, and creating search streams allow users to turn the tedious typed-query keyword search to a fluid, almost physical, interaction.

We reported a small open-ended user study targeted to demonstrate the overall acceptance and user experience of the system. While users took advantage of the system, they also reported unmet needs, such as the improved explanation functionality for making sense of the connections between the documents and entities. Formal studies on the exploration capabilities and sense making, along with automatic methods to extract the knowledge graph from text, remain future challenges and require further experimentation.

ACKNOWLEDGEMENTS

This research was partially funded by TEKES (Re:Know), the Academy of Finland (278090, 305739), the European Commission through the FP7 Project MindSee 611570 and by FCT/MCTES LARSyS (UID/EEA/50009/2013 (2015-2017)). Certain data included herein are derived from the Web of Science prepared by THOMSON REUTERS, Inc., Philadelphia, Pennsylvania, USA: Copyright THOMSON REUTERS, 2011. All rights reserved. Data is also included from the Digital Libraries of the ACM, IEEE, and Springer.

REFERENCES

1. Ahn, J.-W., and Brusilovsky, P. Adaptive visualization for exploratory information retrieval. *Inf. Proc. Man.* 49, 5 (Sept. 2013), 1139–1164.
2. Balog, K., Bron, M., and De Rijke, M. Query modeling for entity search based on terms, categories, and examples. *ACM Trans. Inf. Syst.* 29, 4 (Dec. 2011), 22:1–22:31.
3. 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 *Proc. CHI '11*, ACM (2011), 167–176.
4. Cheng, T., Yan, X., and Chang, K. C.-C. Supporting entity search: A large-scale prototype search engine. In *Proc. SIGMOD*, ACM (2007), 1144–1146.
5. Dork, M., Henry Riche, N., Ramos, G., and Dumais, S. Pivotpaths: Strolling through faceted information spaces. *IEEE Trans. Vis. and Comp. Graph.* 18, 12 (Dec. 2012), 2709–2718.
6. Hinrichs, U., Butscher, S., Müller, J., and Reiterer, H. Diving in at the deep end: The value of alternative in-situ approaches for systematic library search. In *Proc. CHI*, ACM (2016), 4634–4646.
7. Hinrichs, U., Carpendale, S., Valkanova, N., Kuikkaniemi, K., Jacucci, G., and Vande Moere, A. Interactive public displays. *IEEE Comput. Graph. Appl.* 33, 2 (Mar. 2013), 25–27.
8. Jacucci, G., Morrison, A., Richard, G. T., Kleimola, J., Peltonen, P., Parisi, L., and Laitinen, T. Worlds of information: Designing for engagement at a public multi-touch display. In *Proc. CHI '10*, ACM (2010), 2267–2276.
9. Jeh, G., and Widom, J. Scaling personalized web search. In *Proc. WWW*, WWW '03, ACM (2003), 271–279.
10. Jetter, H.-C., Gerken, J., Zöllner, M., Reiterer, H., and Milic-Frayling, N. Materializing the query with facet-streams: A hybrid surface for collaborative search on tabletops. In *Proc. CHI*, ACM (2011), 3013–3022.
11. Klouche, K., Ruotsalo, T., Cabral, D., Andolina, S., Bellucci, A., and Jacucci, G. Designing for exploratory search on touch devices. In *Proc. CHI*, ACM (2015), 4189–4198.
12. Meij, E., Balog, K., and Odijk, D. Entity linking and retrieval for semantic search. In *Proc. WSDM*, ACM (2014), 683–684.
13. Morris, M. R., Fisher, D., and Wigdor, D. Search on surfaces: Exploring the potential of interactive tabletops for collaborative search tasks. *Inf. Proc. & Man.* 46, 6 (2010), 703 – 717. Collaborative Information Seeking.
14. Morris, M. R., Lombardo, J., and Wigdor, D. Wesearch: supporting collaborative search and sensemaking on a tabletop display. In *Proc. CSCW*, ACM (2010), 401–410.
15. Pu, P., Chen, L., and Hu, R. A user-centric evaluation framework for recommender systems. In *Proc. RecSys '11*, ACM (2011), 157–164.
16. Ruotsalo, T. *Methods and Applications for Ontology-Based Recommender Systems*. PhD thesis, Aalto University School of Science and Technology, Finland, 2010.
17. Ruotsalo, T. Domain specific data retrieval on the semantic web. In *Proc. ESWC*, Springer-Verlag (2012), 422–436.
18. Ruotsalo, T., and Hyvönen, E. Exploiting semantic annotations for domain-specific entity search. In *Proc. ECIR 2015*, Springer International Publishing (2015), 358–369.
19. Ruotsalo, T., Jacucci, G., Myllymäki, P., and Kaski, S. Interactive intent modeling: Information discovery beyond search. *Commun. ACM* 58, 1 (Dec. 2014), 86–92.
20. Ruotsalo, T., Peltonen, J., Eugster, M., Glowacka, D., Konyushkova, K., Athukorala, K., Kosunen, I., Reijonen, A., Myllymäki, P., Jacucci, G., and Kaski, S. Directing exploratory search with interactive intent modeling. In *Proc. CIKM*, ACM (2013), 1759–1764.
21. Shaer, O., Strait, M., Valdes, C., Feng, T., Lintz, M., and Wang, H. Enhancing genomic learning through tabletop interaction. In *Proc. CHI '11*, ACM (2011), 2817–2826.
22. Thudt, A., Hinrichs, U., and Carpendale, S. The bohemian bookshelf: Supporting serendipitous book discoveries through information visualization. In *Proc. CHI '12*, ACM (2012), 1461–1470.
23. Varcholik, P. D., Jr., J. J. L., and Hughes, C. E. Establishing a baseline for text entry for a multi-touch virtual keyboard. *Int. J. of Human-Computer Studies* 70, 10 (2012), 657 – 672.
24. Zhai, C., and Lafferty, J. A study of smoothing methods for language models applied to ad hoc information retrieval. In *c. SIGIR '01*, ACM (2001), 334–342.