DOI:10.1145/2656334

The system should let users incrementally direct their search toward relevant, though not initially obvious, information.

BY TUUKKA RUOTSALO, GIULIO JACUCCI, PETRI MYLLYMÄKI, AND SAMUEL KASKI

Interactive Intent Modeling: Information Discovery Beyond Search

COMBINING INTENT MODELING and visual user interfaces can help users discover novel information and dramatically improve their information-exploration performance.

Current-generation search engines serve billions of requests each day, returning responses to search queries in fractions of a second. They are great tools for checking facts and looking up information for which users can easily create queries (such as "Find the closest restaurants" or "Find reviews of a book"). What search engines are not good at is supporting complex information-exploration and discovery tasks that go beyond simple keyword queries. In information exploration and discovery, often called "exploratory search," users may have difficulty expressing their information needs, and new search intents may emerge and be discovered only as they learn by reflecting on the acquired information.[8,9,18] This finding roots back to the "vocabulary mismatch problem"[13] that was identified in the 1980s but has remained difficult to tackle in operational information retrieval (IR) systems (see the sidebar "Background"). In essence, the problem refers to human communication behavior in which the humans writing the documents to be retrieved and the humans searching for them are likely to use very different vocabularies to encode and decode their intended meaning.[8,21]

Assisting users in the search process is increasingly important, as everyday search behavior ranges from simple look-ups to a spectrum of search tasks[23] in which search behavior is more exploratory and information needs and search intents uncertain and evolving over time.

We introduce interactive intent modeling, an approach promoting resourceful interaction between hu-

» key insights

- **Current search engines offer limited assistance in complex search tasks; users are distracted by having to focus their cognitive effort on finding navigation cues rather than on learning and selecting relevant information.**
- **Interactive intent modeling enhances human information exploration through computational modeling (visualized for interaction), helping users search and explore via user interfaces that are highly functional but not cluttered or distracting.**
- **Interactive intent modeling can improve task-level information-seeking performance by over 100%.**

IMAGE BY KHALIL KLOUCHE, UNIVERSITY OF HELSINKI

mans and IR systems to enable information discovery that goes beyond search. It addresses the vocabulary mismatch problem by giving users potential intents to explore, visualizing them as directions in the information space around the user's present position, and allowing interaction to improve estimates of the user's search intents.

Interactive intent modeling is based on two scientific principles (see Figure 1):

Visualization. Visualizing the current search intent and directions in the information space; and

Adaptation. Interactive adaptation of the intent model, balancing exploration of the information space and exploitation of user feedback; the intent model must be able to rigorously handle uncertainty due to limited, possibly suboptimal, user feedback.

By visualizing query and data elements (such as keywords), this approach enables the system to show its understanding of user search intent to the user and also provide a view of available search directions around the user's current position in the information space. The initial evidence concerning user search intent is often limited. The intent model is thus manageable for the user only if the system is able to predict a sufficient subset of the potentially relevant intents. Given the visualization of the intent model and its relation to the information space, the user is able to provide feedback for the intent model, allowing the system to improve intent estimates on subsequent iterations, retrieve and rank data, and update the visualization of directions in the information space.

Interactive Intent Modeling Example

The SciNet system for scientific literature search (http://augmentedresearch.hiit.fi/) is an example of the two principles in interactive intent modeling (see Figure 2).[20] The system currently indexes more than 50 million scientific articles and is designed to assist users exploring information related to a particular research topic through rapid feedback loops and in making sense of the available information around the initial query context.[14,20]

In the Figure 2 scenario, a user is trying to learn about "3D gestures" and types in the corresponding query. The user is visualized with an estimate of his or her present search intents, as well as potential intents, and directions in the information space on a radar screen. The user then navigates by directly manipulating the estimated intents on the display.

Figure 2a is the system's response to the initial query on 3D gestures, offering directions to, say, "video games," "user interfaces," "gesture recognition," and "virtual reality." In Figure 2b, the user has selected "gesture recognition" and is offered further options to continue the exploration to more specific topics (such as "nearest neighbor approach" and "hidden Markov models") but also to more general topics (such as "pointing gestures" and "spatial interaction") estimated to be relevant for the user's interaction history. The modeling is based on a fast online regression model that estimates task-level search intents.[20] The model estimates relevance related to potential search intent and uncertainty related to these estimates based on user feedback. Search intents are visualized as keywords, and selection of which intents to visualize is determined through the exploration-exploitation paradigm. The trick is to present to the user not only choices estimated to be most relevant but those for which the upper confidence bound is the greatest. The user decides whether to explore or exploit, as both relevant and uncertain keywords are visualized; for example, if the user first selects "gesture recognition" and then "hidden Markov models," the system would then suggest specific hidden Markov model applications in gesture recognition, which would be exploitative, as they are estimated to be most relevant and also allow the user to continue to explore more uncertain directions (such as other computational techniques in gesture recognition).

While the idea of providing the user uncertain interaction options may be counterintuitive from a conventional

Figure 1. Exploring information with interactive intent modeling is based on two principles: visualizing current search intent and direction; and balancing exploration and exploitation of user feedback. The user's cognitive effort is thus reduced, as it is easier to recognize items instead of having to remember them when reformulating queries.

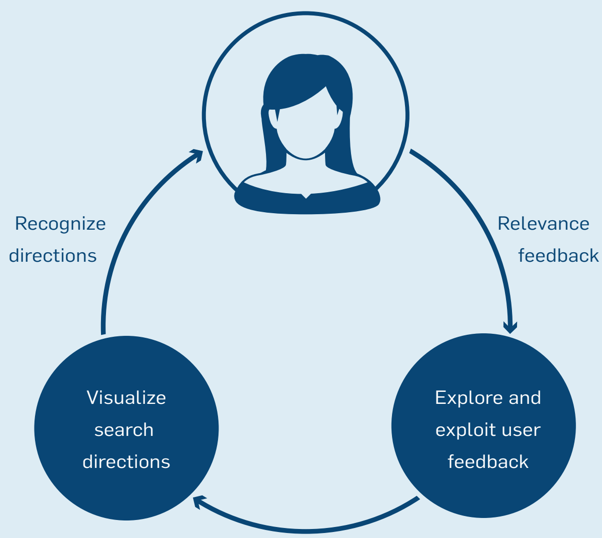

Figure 2. SciNet system search user interface.

(a) The user issues the query "3D gestures," and the system visualizes an intent model on the radar screen consisting of potentially interesting intent as keywords and a ranked document list. The estimated intents, for which the results on the right side have been retrieved, are visualized for the user (inner darker-gray area). The angular distance corresponds to similarity of intent and the radial distance from the center to relevance. Predicted potential future intents, which help users orient themselves on the radar, is visualized in the outer (lighter-gray) area. The user provides positive feedback by dragging keywords closer to the center of the radar and negative feedback by dragging them further away. Multiple keywords can be dragged in each iteration. Online learning methods make it possible for the system to respond in less than one second.

(b) The user has increased the relevance of "gesture recognition" by dragging the corresponding keywords to the center of the radar screen. The system then visualizes new estimated relevant intents as a set of keywords (such as "pattern recognition," "pointing gestures," "recognition rates," "nearest neighbor approaches," and "hidden Markov models").

IR perspective, which is based on the principle of maximizing relevance, our interactive intent modeling helps users overcome the vocabulary-mismatch problem, as the system provides them interaction resources to continuously direct the search and actively explore relevant, though not initially obvious, information. Experiments show user-task-level performance can be improved significantly.

To deliver such support, the interface provides a nonintrusive relevance-feedback mechanism in which the user pulls keywords closer to the center of the radar screen to increase their importance and pushes keywords away from the center of the radar screen to decrease their importance. The keywords can be enlarged with a fisheye lens that follows the mouse cursor anywhere on the radar screen. In response, the system updates intent visualization and search results. The radar screen's radial layout represents good balance between the amount of information shown and comprehensibility compared to alternative visualizations with lower or higher degrees of freedom that could make interaction with the visualization more difficult.

Interactive Visualization of Search Intent and Direction

The SciNet example demonstrates how visualization can be used to elicit feedback. Feedback can be targeted directly to the intent model (the inner circle of the radar screen in Figure 2) or to possible future directions (the outer rim of the radar screen in Figure 2). Due to the vocabulary-mismatch problem, users often have trouble expressing their needs as written queries and are likely to start their search with imprecise queries. Hence, interaction and feedback mechanisms that engage users to provide feedback on how to direct their search in the subsequent iterations are crucial. This is grounded in a well-known cognitive-science theory stating users find recognition easier than recall.[3] It is usually easier for humans to recognize something they see than describe it from scratch.

However, increasing evidence from IR research supports the finding that while relevance feedback is useful in enabling systems to better serve user search intent, in most cases users do not in fact use feedback mechanisms.[18] This observation is related to two other cognitive science findings: users find it easier to recognize smaller units than more complex ones, and it is easier for them to make sense of information relative to a reference point than in isolation.[6] Assessing the relevance of a full document may be an even more demanding task than formulating a new query.

Information visualization can turn laborious relevance assessment into a more fluent recognition task; for example, visualizing essential document content can be faster for recognizing important directions toward finding relevant information than forcing users to read this information from an original document.[1,17]

Recent visualizations applied in search tasks support sense making of bibliographic data by, for example, incrementally interactively exploring networks of data.[10] While these systems show the importance of visualization for sense making, they are limited by not allowing users to negotiate the intent model with the system, allowing them only to explore information through direct links already present in

Background

Recent behavioral studies show a large portion of user information-seeking activities are exploratory and characterized by the complex and evolving nature of user information needs.[18] As a result, users face the problem of entering correct terms that describe their search intents so the desired information can be retrieved on subsequent iterations. This is one of the major findings of information-seeking research, as identified by Furnas et al.,[13] Saracevic and Kantor,[21] and Bates.[7,8] All demonstrated human communication patterns are not likely well suited to creating written queries. For example, two study subjects in Furnas et al.[13] favored the same search term with a probability smaller than 0.2, which was shown to lead to 80%–90% failure rates in many common search situations; Saracevic and Kantor[21] and Zhao and Callan[25] later obtained similar results. These findings limit the success of various design methodologies for written-query-driven search interfaces, highlighting the importance of the user interface for modeling and discovering users' potential, yet unspecified, search intent.

The simplistic search interfaces of the current generation of search engines not only force users to focus their cognitive effort on the discovery of information potentially relevant for their search intents, in most cases, they deliver only the search result listing that supports this activity. A system's ability to discover potentially relevant search intents could be better supported by search engines in the first place. As user expressions of information needs are often suboptimal, reflecting only limited evidence of their real search intents, opportunities for IR software developers focus on figuring out ways to get users to negotiate with the search engines to better capture their needs. This goal of utilizing the human information-processing system more effectively in the search process has led researchers to combine work in human-computer interaction, information retrieval, and machine learning to overcome barriers between techniques that rely on only human information-processing capabilities and computational methods employed by search engines.

Other researchers have investigated adaptive and interactive search user interfaces and their effect on retrieval performance. For example, Hearst et al.[16] developed a variety of search user interfaces that use filtering and visualization techniques ranging from hierarchical, faceted metadata[24] and search result clustering[16] to visualization of the similarity between query term and result item.[15] Marchionini[18] proposed user interfaces to support exploratory search, including browsing and retrieving video content. While these techniques are all highly functional and have shown increased search effectiveness, they exploit only the information already found by the search engine in response to a user's query. They did not take into account that the initial expression and future refinement of information needs are often suboptimal,[12] especially when users are unfamiliar with the domain and its vocabulary.

Moreover, users are learning the search vocabulary as they make sense of the information space during the search process; for example, if a user searches for "search engines," techniques that exploit the results already found by the search engine may limit the user's options for exploring beyond the initially found results, as the user is not presented different words for the same highly relevant additional topics (such as "information retrieval" and "information seeking behavior"). Despite these findings, the objective of much current IR system theory is toward systems where the user has limited and often only reactive involvement in the search process.

the network data.[1]

By visualizing potential intents, an IR system can give users a spectrum of choices in a form suitable for the human visual system to process rapidly, even when the data changes dynamically as interaction occurs. Such an interface requires advanced data-driven visualizations that can be computed online. Moreover, a visualization should not contain only information already familiar to the user that would be good for recognition but lead to the intent model getting stuck in a "context bubble." Instead, unseen parts of the information space must be offered to the user, facilitating sense making through the relation of these parts to already familiar information when possible.

Balancing Exploration and Exploitation

Given the evolutionary nature of search, as demonstrated in our example search scenario involving SciNet, it is important to not only exploit the feedback elicited from the user but balance it with exploration. Users must be able to focus on a specific location in the information space (exploit) and be able to broaden their search through more general areas (explore).

This insight is particularly important for users exploring information with which they are not familiar. Users often suffer from what psychologists call "anchoring," or the tendency to make insufficient adjustments to initial values when judging under conditions of uncertainty.[22] Users may thus tend to refrain from abandoning their initial expression of their information needs or from adjusting them very much, causing subsequent expression of information needs to be biased toward their current knowledge. This bias reduces the likelihood they will discover something novel.

This behavioral finding has consequences for machine-learning approaches to modeling search intent. A promising direction for predicting intent while still allowing users to be in control of the search process comes from machine-learning methods that learn online. Online learning methods are able to update models one observation at a time so future predictions can be made immediately when feedback is received. The goal in online learning for search is to predict the relevance of content interactively, meaning that soon after prediction is made, the judgment of its usefulness is received from user feedback. This information can then be used by the IR system to refine the prediction hypothesis used by the method.

Online learning methods are able to update models one observation at a time so future predictions can be made immediately when feedback is received.

Standard machine learning for online prediction does not solve the problem of discovering what interaction options are most useful in allowing the learning method to improve its estimates and therefore create visualizations for the user. Straightforward exploitation by choosing the directions currently estimated to be most relevant could lead to converging to suboptimal goals and locking users in context bubbles predetermined by the user's limited prior knowledge; for example, a user searching for "3D gestures" might never explore "pointing gestures," as the initial query scope could already be too specific to allow such exploration. A promising solution for collecting feedback that also allows exploration is the "exploration-exploitation" paradigm of reinforcement learning.[5] In it, the model and its environment (the user) form an online loop, and learning involves finding a balance between exploration (showing items from uncharted information space for feedback) and exploitation (showing items most likely to be relevant, given the current user intent model).

Users can thus be assisted to direct their searches under conditions of uncertainty by learning intent models online based on feedback they give about the models. Due to the limited and imperfect feedback available, the amount of uncertainty about user intent can be substantial. It is therefore important for an IR system to use models capable of handling uncertainty in a rigorous manner. Probabilistic online learning models can be used in the exploitation part of the exploration-exploitation dilemma; moreover, as the models are probabilistic, they quantify uncertainty related to the estimates, using it to determine the optimal alternative choices to be visualized for users.[5]

From Effectiveness to Task Performance

When studying IR systems designed to negotiate user search intent, it is important to realize the utility these systems improve is not retrieval effectiveness at a query-response level but task-level performance. Within interactive IR systems, users are often required to do more work to complete their tasks, and, while some of it may be wasted, they may be more successful in correcting their initially suboptimal actions. To get a complete picture of performance, two aspects of IR systems should be measured simultaneously: IR system effectiveness, given a complete description of an information need; and human task performance, given the system's interaction modes.

The SciNet system was recently studied in task-based experiments in which users were given 30 minutes to solve research tasks using IR systems operating on a database of more than 50 million scholarly articles. We compared a system setup with interactive intent modeling against a conventional IR system with a list-based visualization and interaction with typed queries. We quantified the quality of retrieved information, adoption of the visualization, and feedback mechanisms separately.[14,20] We found interactive search intent modeling significantly improved users task performance. We also found the task outcomes graded more highly by experts, and the search user interface enhanced interaction without compromising task execution time. We attributed the improved task performance to the improved quality of retrieved information and to the improved visualizations and interaction modes offered by the system. In particular, interactive intent modeling increased recall of novel information without losing precision. This performance demonstrates the power of the interactive intent modeling technique in supporting exploration and discovery of information that can be difficult to find with systems that rely on conventional search user interfaces (see the table here).

Making Intent Modeling Ubiquitous

Engaging users to interact with IR systems is crucial for such systems to be able to offer better interaction modes and reduce uncertainty related to user expression of search intent. Despite significant improvement in user task performance on the example SciNet system discussed earlier, we are only scratching the surface of human-centered computing as part of the search activity. Intent-aware IR systems can benefit from ubiquitous computing in at least two ways, as discussed next.

Wearable User Interfaces and Augmented Reality

IR systems can be extended by augmenting a real scene with predictions of what the user might find useful, shown as augmented reality on head-mounted displays (HMDs). Users' implicit and explicit reactions to visualized content can reveal their intent and help improve the user intent model contextualized to the immediate setting. Figure 3a shows how suitable information (such as topics, research group, and publications) the user can recognize and act upon can be visualized on a HMD superimposed on the real scene;[2] for example, augmenting a user's environment when visiting a poster session at a conference with visual cues and information can help the system collect information about the user's intent even when the user is not actively engaged with a search engine.

Implicit feedback from physiological computing. Recent advances in wearable computing have facilitated capturing users' affective and cognitive states (such as wearable electroencephalography, or EEG, systems). Moreover, other physiological sensors (such as galvanic skin response and heart-rate sensors) are being integrated into wrist-wearable products like smartwatches. Such physiological signals give researchers additional sources of feedback information not previously available.

Key benefits of interactive intent modeling.

Improved task performance	Interactive intent modeling improves users' task performance compared to state-of-the-art retrieval methods and alternative search user-interface techniques.[20]
Quality of retrieved information	Interactive intent modeling helps users go beyond their initial query context, allowing them to increase recall significantly while preserving precision, particularly for novel information, leading to session-level improvement of 100%.[14,20]
Enhanced interaction	Interactive visualizations enhance interaction without compromising task execution time. Users in our experiments chose visualization as their main user-interface component for making sense of returned information and for expressing their search intent.[20]

Figure 3. Making intent modeling ubiquitous.

(a) Suitable information the user can recognize and act upon can be visualized on a display as augmented reality.[2] While the user visits a poster session at a conference, the IR system suggests information by augmentations on the data glasses. The system can then, based on implicit and explicit interactions, iterate the intent model and propose new information.

(b) An experiment involving term-relevance prediction from brain signals via EEG to automatically detect the relevance of textual information directly from brain signals. Wearable EEG and other techniques can be used for implicit relevance feedback to improve prediction of the intent model to complement or substitute explicit relevance judgments.[11]

It has been shown that affective state information can be used for relevance judgment prediction,[4] and affective and psychophysiological signals are being employed in multimedia search systems with encouraging results.[19] Figure 3b shows an example of an EEG sensor setup used to demonstrate term-relevance prediction from brain signals. The experiment shows it is possible to automatically detect the relevance of text information visualized for the user directly from brain signals by analyzing neural activity of participants while providing relevance judgments to text terms for a given topic.[11] Employing such physiology-based relevance detection for implicit relevance feedback on visualized information can be used by IR system developers to improve the prediction of the intent model to complement or substitute explicit user relevance ratings.

Conclusion

Recent work demonstrates there is significant room for improving the support provided to users involved in exploratory forms of search. Overall, researchers recognize the need for search and information-exploration systems that combine the information-processing capabilities of humans and computers.

Interactive intent modeling is a theoretically motivated, empirically proven way to support information exploration and discovery. It can increase users' capacity for information processing and discovery through computing technologies that assist users navigating complex information spaces.

Interactive intent modeling provides additional resources for users to better learn about the information space and give increased feedback for the system so it can efficiently adapt its understanding of user-search intent.

Engaging users to adopt interactive feedback mechanisms for information exploration and sense making requires user-interface techniques that go beyond search boxes and lists of links to enable them to better interact with the system and have control over their findings. Modeling user intent online as interaction occurs and even in situations where user feedback is noisy and suboptimal requires machine-learning models that learn online and are able to explore, not just exploit. IR system design must ultimately integrate interactive visualizations, intent prediction, multimodal feedback, and a higher-level context of tasks and goals.

IR systems must be able to help users solve tasks, not just retrieve documents. Users need search engines and user interfaces that adapt to their capabilities and search behavior, rather than require them to adapt to them.

Acknowledgments

This work was partly supported by the Academy of Finland (278090, Multivire 255725 and the COIN Center of Excellence 251170), TEKES (D2I and Re:KnoW), and the European Commission through the FP7 Project MindSee 611570. Certain data included here is derived from the Web of Science prepared by Thomson Reuters, Inc., Philadelphia, PA. 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. *Information Processing & Management 49*, 5 (Sept. 2013), 1139–1164.
2. Ajanki, A., Billinghurst, M., Gamper, H., Järvenpää, T., Kandemir, M., Kaski, S., Koskela, M., Kurimo, M., Laaksonen, J., Puolamäki, K., Ruokolainen, T., and Tossavainen, T. An augmented reality interface to contextual information. *Virtual Reality 15*, 2-3 (June 2011), 161–173.
3. Anderson, J.R. *Learning and Memory: An Integrated Approach, First Edition.* John Wiley & Sons, New York, 2000.
4. Arapakis, I., Athanasakos, K., and Jose, J.M. A comparison of general vs. personalised affective models for the prediction of topical relevance. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Geneva, Switzerland, July 19–23). ACM Press, New York, 2010, 371–378.
5. Auer, P. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research 3* (Mar. 2003), 397–422.
6. Avons, S., Wright, K., and Pammer, K. The word-length effect in probed and serial recall. *The Quarterly Journal of Experimental Psychology 47*, 1 (1994), 207–231.
7. Bates, M.J. Subject access in online catalogs: A design model. *Journal of the American Society for Information Science 37*, 6 (Nov. 1986), 357–376.
8. Bates, M.J. Indexing and access for digital libraries and the Internet: Human, database, and domain factors. *Journal of the American Society for Information Science 49*, 13 (1998), 1185–1205.
9. Belkin, N.J., Oddy, R.N., and Brooks, H.M. Ask for information retrieval: Part I: Background and theory. In *Readings in Information Retrieval*. Morgan Kaufmann Publishers Inc., San Francisco, CA, 1997, 299–304.
10. 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* (Vancouver, BC, Canada, May 7–12). ACM Press, New York, 2011, 167–176.
11. Eugster, M.J.A., Ruotsalo, T., Spapé, M.M., Kosunen, I., Barral, O., Ravaja, N., Jacucci, G., and Kaski, S. Predicting term-relevance from brain signals. In *Proceedings of the 37th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Gold Coast, Queensland, Australia, July 6–11). ACM Press, New York, 2014, 425–434.
12. Fu, W.-T. and Gray, W.D. Resolving the paradox of the active user: Stable suboptimal performance in interactive tasks. *Cognitive Science 28*, 6 (Nov. 2004), 901–935.
13. Furnas, G.W., Landauer, T.K., Gomez, L.M., and Dumais, S.T. The vocabulary problem in human-system communication. *Commun. ACM 30*, 11 (Nov. 1987), 964–971.
14. Glowacka, D., Ruotsalo, T., Konuyshkova, K., Athukorala, K., Kaski, S., and Jacucci, G. Directing exploratory search: Reinforcement learning from user interactions with keywords. In *Proceedings of the 2013 International Conference on Intelligent User Interfaces* (Santa Monica, CA, Mar. 19–22). ACM, Press, New York, 2013, 117–128.
15. Hearst, M.A. Tilebars: Visualization of term distribution information in full text information access. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, CO, May 7–11). ACM Press/Addison-Wesley Publishing Co., New York, 1995, 59–66.
16. Hearst, M.A. and Pedersen, J.O. Reexamining the cluster hypothesis: Scatter/gather on retrieval results. In *Proceedings of the 19th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Zurich, Switzerland, Aug. 18–22). ACM Press, New York, 1996, 76–84.
17. Lehmann, S., Schwanecke, U., and Dörner, R. Interactive visualization for opportunistic exploration of large document collections. *Information Systems 35*, 2 (Apr. 2010), 260–269.
18. Marchionini, G. Exploratory search: From finding to understanding. *Commun. ACM 49*, 4 (Apr. 2006), 41–46.
19. Moshfeghi, Y. and Jose, J.M. An effective implicit relevance feedback technique using affective, physiological and behavioural features. In *Proceedings of the 36th International ACM Conference on Research and Development in Information Retrieval* (Dublin, Ireland, July 28–Aug. 1). ACM Press, New York, 2013, 133–142.
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 *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management* (San Francisco, CA, Oct. 27–Nov. 1). ACM Press, New York, 2013, 1759–1764.
21. Saracevic, T., Kantor, P., Chamis, A. Y., and Trivison, D. A study of information seeking and retrieving. I. Background and methodology. *Journal of the American Society for Information Science 39*, 3 (May 1988), 161–176.
22. Strack, F. and Mussweiler, T. Explaining the enigmatic anchoring effect: Mechanisms of selective accessibility. *Journal of Personality and Social Psychology 73*, 3 (Sept. 1997), 437–446.
23. Wilson, M.L., Kules, B., Schraefel, M.C., and Shneiderman, B. From keyword search to exploration: Designing future search interfaces for the Web. *Foundations and Trends in Web Science 2*, 1 (Jan. 2010), 1–97.
24. Yee, K.-P., Swearingen, K., Li, K., and Hearst, M. Faceted metadata for image search and browsing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, FL, Apr. 5–10). ACM Press, New York, 2003, 401–408.
25. Zhao, L. and Callan, J. Term necessity prediction. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management* (Toronto, Canada, Oct. 26–30). ACM Press, New York, 2010, 259–268.

Tuukka Ruotsalo (tuukka.ruotsalo@aalto.fi) is a postdoctoral researcher in the Helsinki Institute for Information Technology HIIT at Aalto University, Helsinki, Finland.

Giulio Jacucci (giulio.jacucci@cs.helsinki.fi) is a professor in the Department of Computer Science, University of Helsinki, Finland, and a co-founder of MultiTaction Ltd.

Petri Myllymäki (petri.myllymaki@cs.helsinki.fi) is a professor in the Department of Computer Science at the University of Helsinki, Finland.

Samuel Kaski (samuel.kaski@hiit.fi) is director of the Helsinki Institute for Information Technology HIIT and a professor in the Department of Information and Computer Science at Aalto University, Helsinki, Finland.