



## UvA-DARE (Digital Academic Repository)

### Conversational Exploratory Search via Interactive Storytelling

Vakulenko, S.; Markov, I.; de Rijke, M.

**Publication date**

2017

**Document Version**

Author accepted manuscript

[Link to publication](#)

**Citation for published version (APA):**

Vakulenko, S., Markov, I., & de Rijke, M. (2017). *Conversational Exploratory Search via Interactive Storytelling*. Paper presented at SCAI'17 — Search-Oriented Conversational AI, Amsterdam, Netherlands. <https://arxiv.org/abs/1709.05298>

**General rights**

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

**Disclaimer/Complaints regulations**

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

# Conversational Exploratory Search via Interactive Storytelling

## Position Paper

Svitlana Vakulenko  
Vienna University of Economics and  
Business  
Vienna, Austria  
svitlana.vakulenko@wu.ac.at

Ilya Markov  
University of Amsterdam  
Amsterdam, The Netherlands  
i.markov@uva.nl

Maarten de Rijke  
University of Amsterdam  
Amsterdam, The Netherlands  
derijke@uva.nl

### ABSTRACT

Conversational interfaces are likely to become more efficient, intuitive and engaging way for human-computer interaction than today's text or touch-based interfaces. Current research efforts concerning conversational interfaces focus primarily on question answering functionality, thereby neglecting support for search activities beyond targeted information lookup. Users engage in exploratory search when they are unfamiliar with the domain of their goal, unsure about the ways to achieve their goals, or unsure about their goals in the first place. Exploratory search is often supported by approaches from information visualization. However, such approaches cannot be directly translated to the setting of conversational search.

In this paper we investigate the affordances of interactive storytelling as a tool to enable exploratory search within the framework of a conversational interface. Interactive storytelling provides a way to navigate a document collection in the pace and order a user prefers. In our vision, interactive storytelling is to be coupled with a dialogue-based system that provides verbal explanations and responsive design. We discuss challenges and sketch the research agenda required to put this vision into life.

### KEYWORDS

Conversational search; Exploratory search; Chatbot

## 1 INTRODUCTION

Exploratory search systems provide guidance for users who are exploring unfamiliar information landscapes [11, 24]. White and Roth [24] differentiate two main activities within the exploratory search paradigm: *exploratory browsing* and *focused searching*. Exploratory browsing is an initial step that provides necessary domain understanding required for focused searching activities. It is related to Radlinski and Craswell [19]'s *system revealment* property: "The system reveals to the user its capabilities and corpus, building the user's expectations of what it can and cannot do."

Lately, conversational agents and conversational search systems are becoming increasingly popular [21]. So far, however, such systems mainly focus on question answering and simple search tasks,

those that are to a large extent solved by web search engines. We argue that conversational agents and search systems should also support exploratory search. While exploratory search is a challenging task in itself, *conversational exploratory search* raises unique research and practical issues, which we discuss in this position paper.

In particular, we argue that the core of conversational exploratory search is *interactive storytelling*, where the document collection underlying a conversational search system is first converted into a set of stories and then a user interactively navigates within a story and between stories by means of a dialogue with the system.

There have been several recent position statements on conversational agents and search. One, by Radlinski and Craswell [19], focuses on a theoretical model of conversational search systems. Another, by Kiseleva and de Rijke [8], focuses on evaluation. In contrast, we focus on solution strategies for a specific conversational search scenario, viz. exploratory search.

Below, we first provide motivating examples in Section 2. In Section 3 we present our view of a conversational exploratory search system. The research agenda associated with this system is presented in Section 4. The paper is concluded in Section 5.

## 2 MOTIVATING EXAMPLES

The literature is full of arguments motivating computational support for exploratory search [23]. Exploratory search is an important enabler for educational purposes that aim to broaden the knowledge of a user and understanding of the domain by enhancing learning processes. Serendipitous discoveries are very important in less structured and content rich domains such as music, videos, design etc., where users often look for inspiration, surprises and novel ideas [26]. Furthermore, the potential benefits of conversational exploratory search for e-commerce applications should not be underestimated. In particular, it can be combined with personal recommendations and persuasion techniques for marketing purposes [16].

For us, one of the main motivations behind conversational exploratory search comes from the results of analyzing the conversation log of a chatbot demo that some of the authors were involved with [17].<sup>1</sup> This chatbot demo exposes search functionality over an aggregated open data repository [18] via a conversational interface. Manual inspection of the conversation log of the digital assistant revealed that the majority of users experience difficulties formulating adequate queries to the system, i.e., queries that return any matches. This effect is, to a large extent, due to a misconception of

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SCAI'17, October 1, 2017, Amsterdam, Netherlands

© 2017 Copyright held by the owner/author(s).

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

<sup>1</sup><https://m.me/OpenDataAssistant>

the underlying collection of documents, which can potentially be retrieved using a search system.

The observation of a user's mistaken internal representation of a document collection is not new. The information seeking literature is full of examples to this effect and the information retrieval community has proposed a range of technological solutions to help address such mismatches, ranging from algorithms that help recover from possibly empty search engine result pages using query suggestions and rewrites [10] to information visualization techniques to help steer users in possibly useful regions of a document collection [5].

However, visualization approaches may vary significantly and be hard to understand without an animated explanation in natural language or even specialized training. While such methods may be effective in traditional keyboard or touch-based exploratory search scenarios, by and large they are inappropriate to support exploratory search in a conversational setting on mobile devices. Instead, we argue that an approach based on interactive storytelling is called for to support conversational exploratory search.

### 3 CONVERSATIONAL EXPLORATORY SEARCH

Our view of a conversational exploratory search system is represented in Figure 1. It has a number of key components: Document Collection, Knowledge Model, Story Space, Dialog System and User. These components are connected through the Reader, Composer, and Guide modules. The interplay of the system components and modules happens at different stages.

*Knowledge Representation.* Knowledge representation consists of the *Reader* module that extracts concepts and relations from the Document Collection and embeds them into a single Knowledge Model. The Knowledge Model integrates different elements (words, concepts or entities) and describes relations between them. The knowledge can be explicitly modeled by means of a taxonomy or ontology (knowledge graph) but it can also be embedded into a latent (hidden) structure.

*Story Generation.* Story generation consists of the *Composer* module that is able to generate stories by combining elements of the Knowledge Model. To create a story, the Composer has to select elements (characters, words, facts, concepts, relations), choose their ordering, arrange selected elements in time and/or space. The set of all possible stories constitutes the Story Space.

*Interactive Storytelling.* Interactive storytelling consists of the *Guide* module that helps the User to navigate through the Document Collection via the Story Space. The Guide can change the current position within a single story or traverse the space across different stories. Interactive storytelling integrates the Dialogue System to communicate a story to the User and to receive an input from the User. Supporting such a conversation with the User requires natural language (utterance) generation and understanding. Note that the input/output modalities do not have to be restricted to text and speech only and may include images, videos, interactive visualization, virtual reality interactions, etc.

We also argue that a conversational exploratory search system should support the following types of the user-system interactions:

- *Navigation Control* – a user chooses a direction (branch) for exploration and is also able to influence and change the current direction of the narrative at any point in time;
- *Feedback* – a user may provide feedback to the system (positive, neutral, negative) that may help to correct and steer the direction of the story that shall maximize the user satisfaction with the system;
- *Question* – a user may pose questions to the system, e.g., a request for a definition, look up query, etc.

A sample dialogue with interactions of all these types is provided in Figure 2. In this example the dialogue agent provides concise natural language descriptions of the information space structure, suggestions for possible exploration directions, and further support and guidance along the chosen direction for exploration.

## 4 RESEARCH AGENDA

We identify the following research questions with respect to the components and interaction types described in the previous section:

- RQ1.** *Reader:* How to model the information space structure, represent documents and relations between them for the purpose of story generation?
- RQ2.** *Composer:* How to generate a coherent narrative (story) that efficiently describes a knowledge model?
- RQ3.** *Guide:* How to efficiently traverse/navigate a story space?
- RQ4.** *Dialogue system:* How to provide support for the following three types of the system-user interactions:
- RQ4.1.** *Storytelling:* How to communicate a story to a user?
- RQ4.2.** *Question generation:* How to verify user understanding, satisfaction and preferences?
- RQ4.3.** *Response analysis:* How to interpret and correctly react to natural language utterances (or other signals), such as the ones expressing user satisfaction (feedback), communicating the desired directions for traversing the information space (navigation control), checking the terminology and asking other types of questions?

In the following, we organize these research questions into two subtasks, namely, story generation and interactive storytelling.

### 4.1 Story Generation

In the context of conversational AI we are primarily interested in developing an operational knowledge model (**RQ1**), i.e., the structure that the system can act upon, e.g., to answer questions or generate stories. Story generation (**RQ2**) requires accomplishing the following three tasks: (1) select elements of the knowledge model; (2) choose an order in which to present these elements; and (3) communicate the story to the user using the modalities available to the system, e.g., natural language and/or visualization (**RQ4.1**).

Computational narrative intelligence, the ability to craft, tell, understand and respond appropriately to narratives, is a core component of a strong AI system [9, 20]. So far, it has mostly been developed with applications to fiction, in the context of computational creativity. We propose to put it to work for conversational exploratory search. To this end, we first recall some core concepts




Figure 1: Communicating knowledge via an interactive storytelling process.

**System:** I am here to help you explore open data! This collection includes more than 18,000 open datasets from countries such as Italy, Ireland, Mexico, Austria, etc.

**User:** What data do you have about Austria?

**System:** Jobs, healthcare, population statistics, education, environment... Are you interested in anything from the above?

**User:** What else?

**System:** Traffic data, sports, agriculture, tourism.

**User:** Traffic sounds interesting! What kind of data is there?

**System:** There is data for different cities in Austria: Vienna, Linz, Salzburg, Innsbruck.

**User:** ok

**System:** For example, pedestrian zones in Vienna, bridges locations, Electronic timetable information for Linz.

**User:** Great! Show me something else.

**System:** There are 96 datasets about Education and Research in Austria. One of them is a List of all driving schools in Burgenland.

**User:** What is Burgenland?

**System:** Burgenland is the easternmost and least populous state of Austria [Wikipedia].

Figure 2: Sample dialog for exploratory search based on the Open Data Assistant chatbot use case and the faceted-search interface of the Austrian Government Open Data portal <https://www.data.gv.at>.

from the area and then sketch our ideas for putting it to work for conversational exploratory search.

McIntyre and Lapata [13] use genetic algorithms (GAs) to generate children stories from a corpus of fairy tales. They extract schemas from natural language texts using dependency parsing and co-reference resolution tools, then generate a single plot graph by merging these schemas. The plot graph constitutes the story space, where each path is a different story. The algorithm then searches the story space for the best story candidates using a coherence function learned from training data [2]. The produced stories are readable but short and uninformative, and can be considered as a proof-of-concept for the story generation approach.

Martin et al. [12] generate stories in natural language using two sequence-to-sequence recurrent neural networks (RNNs): (1) event representations are extracted from text using dependency parsing, stemming and topic modeling; (2) event2event RNN chains the extracted events together into stories; (3) event2sentence RNN translates the generated story representation into natural language sentences. This approach is applied to a corpus of movie plot summaries extracted from Wikipedia [1]. It is reported to achieve plausible and human-readable sentences.

Huang et al. [6] establish a new task of visual storytelling, in which the system is to generate a story in natural language given a sequence of images as an input. The baseline model for story generation is trained using sequence-to-sequence RNNs.

In our view, the work on algorithms for story generation is sufficiently mature so that it can be successfully used in the context of conversational exploratory search, especially to support dialogue management in conversational exploratory search, thereby offering the potential to address RQ1, RQ2 and part of RQ4, namely, RQ4.1.

## 4.2 Interactive Storytelling

Conversational exploratory search is not a one way traffic. Hence, our perspective on using story generation for the purposes of conversational exploratory search needs to be complemented with conversational aspects. Interactive storytelling is a conversation, in which a storyteller aims to convey a fraction of a knowledge model to a listener (RQ4.1), and the listener can actively influence the direction, flow and manner of the story being told (responsive by design, RQ4.3). Ability of the storyteller to ask questions and expose possible directions for exploration (RQ4.2) aims at encouraging listener’s active engagement with the story and avoiding lengthy monologues in favor of a more balanced dialogue-based interaction with the content.

Approaches developed within the goal-oriented dialogue framework (Dialog State Tracking Challenge [25]) are likely to be useful for dialogue management in the interactive storytelling settings as well. Within this framework the dialogue system is supported by a task-specific domain ontology. The ontology enumerates all concepts and attributes (slots) that a user can specify or request information for [15]. The dialogue management model is trained to correctly classify user intents by matching user utterances to

the elements in the domain ontology. It can also learn to use the distribution over intents to decide whether to execute an action or request a clarification from the user [15].

The results of the Dialog State Tracking Challenge show advantages of end-to-end dialog systems that employ discriminative models and embed a dialog directly as a sequence [25]. Bordes and Weston [3] show how to train such an end-to-end dialog system using the Memory Network architecture. Dhingra et al. [4] use RNNs and reinforcement learning to train a dialogue system that can interactively retrieve items from a single table.

Mrksic et al. [15] avoid the limitations of the exact word matching by loading pre-trained word vectors and composing them into intermediate representations to be able to scale to larger and more complex domains. They carry out an evaluation for a single domain (restaurants), which is described by an ontology with three attributes specifying the goal (information need) and eight attributes available for retrieval. While very promising for the task of conversational exploratory search, the question remains whether the proposed interactive storytelling approaches can scale up from the toy examples considered so far to support meaningful conversations using the full-sized knowledge graphs.

With the fraction of the knowledge model involved in communication getting bigger the major design challenges arise with respect to the balanced composition of the story space (**RQ2**) that allows efficient traversal and communication taking into account cognitive limitations of the human brain (**RQ3**). In addition, the ability to adopt useful shortcuts across the story space will reduce the traversal time and, thereby, improve the experience by avoiding linear search in favor of random access, when it is applicable (**RQ3**).

In addition to scale, another important challenge arises from the fact that interactive storytelling is different from a common conversational search task, where an agent tries to pin-point an item or an information subspace relevant to the user's query [19]. In this respect, interactive storytelling is hard to optimize, since there is no single correct answer. We propose to measure the results of the interactive storytelling process with respect to: (1) the learning outcomes, which constitute the fraction of the knowledge model gained on the listeners' side; and (2) user satisfaction. The datasets available for learning dialogue representations are currently limited to two types of tasks: general chit-chat and goal-oriented dialogues, such as restaurant reservation [7, 14].

There are a few new datasets of conversation transcripts covering more general search scenarios [21, 22], which focus primarily on analyzing different task complexity levels and user experience during the dialogue interactions. To the best of our knowledge, there is currently no publicly available dataset of conversation logs recorded for learning conversational exploratory browsing behavior and evaluation of successful knowledge transfer interactions.

## 5 CONCLUSIONS

In this paper we introduced the idea of enabling conversational exploratory search by means of interactive storytelling. We presented our vision of such a system, its components and modules. We also outlined directions for future research towards development of the computational narrative intelligence, as an enabler of conversational AI, and its application in the exploratory search scenarios,

which go beyond the discrete look-up requests towards continuous interaction sessions with the goal of knowledge transfer, that we refer to as interactive storytelling.

The insights gained in the fields of story generation and dialogue systems suggest that it is feasible to develop a computational model able to learn natural language generation and communication from crowd-sourced examples. We see our task in developing this idea further by adopting it in the context of exploratory search. To begin in this direction, the research community requires a collection of new datasets of dialogue interactions that can be used for evaluation of successful knowledge transfer. Next, evaluation of existing approaches to story generation and learning dialogue policies in this new settings will help to form the baselines for developing novel approaches.

## ACKNOWLEDGMENTS

The work of Svitlana Vakulenko has received funding from the EU H2020 programme under the MSCA-RISE agreement 645751 (RISE\_BPM) and the Austrian Research Promotion Agency (FFG) under the project CommuniData (grant no. 855407). Ilya Markov and Maarten de Rijke were supported by Ahold Delhaize, Amsterdam Data Science, the Bloomberg Research Grant program, the Criteo Faculty Research Award program, Elsevier, the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement nr 312827 (VOX-Pol), the Microsoft Research Ph.D. program, the Netherlands Institute for Sound and Vision, the Netherlands Organisation for Scientific Research (NWO) under project nrs 612.001.116, HOR-11-10, CI-14-25, 652.002.001, 612.001-551, 652.001.003, and Yandex. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

## REFERENCES

- [1] David Bamman, Brendan O'Connor, and Noah A. Smith. 2013. Learning Latent Personas of Film Characters. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics, ACL 2013, 4-9 August 2013, Sofia, Bulgaria, Volume 1: Long Papers*. 352–361. <http://aclweb.org/anthology/P/P13/P13-1035.pdf>
- [2] Regina Barzilay and Mirella Lapata. 2008. Modeling Local Coherence: An Entity-Based Approach. *Computational Linguistics* 34, 1 (2008), 1–34. <https://doi.org/10.1162/coli.2008.34.1.1>
- [3] Antoine Bordes and Jason Weston. 2016. Learning End-to-End Goal-Oriented Dialog. *CoRR* abs/1605.07683 (2016). <http://arxiv.org/abs/1605.07683>
- [4] Bhuwan Dhingra, Lihong Li, Xiujuan Li, Jianfeng Gao, Yun-Nung Chen, Faisal Ahmed, and Li Deng. 2017. Towards End-to-End Reinforcement Learning of Dialogue Agents for Information Access. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*. 484–495. <https://doi.org/10.18653/v1/P17-1045>
- [5] Marti A. Hearst. 2009. *Search User Interfaces*. Cambridge University Press.
- [6] Ting-Hao (Kenneth) Huang, Francis Ferraro, Nasrin Mostafazadeh, Ishan Misra, Aishwarya Agrawal, Jacob Devlin, Ross B. Girshick, Xiaodong He, Pushmeet Kohli, Dhruv Batra, C. Lawrence Zitnick, Devi Parikh, Lucy Vanderwende, Michel Galley, and Margaret Mitchell. 2016. Visual Storytelling. In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*. 1233–1239. <http://aclweb.org/anthology/N/N16/N16-1147.pdf>
- [7] Tom Kenter, Alexey Borisov, Christophe Van Gysel, Mostafa Dehghani, Maarten de Rijke, and Bhaskar Mitra. 2017. Neural Networks for Information Retrieval. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, Shinjuku, Tokyo, Japan, August 7-11, 2017*. 1403–1406. <https://doi.org/10.1145/3077136.3082062>
- [8] Julia Kiseleva and Maarten de Rijke. 2017. Evaluating Personal Assistants on Mobile devices. *CoRR* abs/1706.04524 (2017). <http://arxiv.org/abs/1706.04524>
- [9] Boyang Li. 2015. *Learning knowledge to support domain-independent narrative intelligence*. Ph.D. Dissertation. Georgia Institute of Technology, Atlanta, GA.

- USA. <http://hdl.handle.net/1853/53376>
- [10] Xinyi Li, Bob Schijvenaars, and Maarten de Rijke. 2017. Investigating queries and search failures in academic search. *Information Processing & Management* 53, 3 (May 2017), 666–683.
- [11] Gary Marchionini. 2006. Exploratory Search: From Finding to Understanding. *Commun. ACM* 49, 4 (April 2006), 41–46. <https://doi.org/10.1145/1121949.1121979>
- [12] Lara J. Martin, Prithviraj Ammanabrolu, William Hancock, Shruti Singh, Brent Harrison, and Mark O. Riedl. 2017. Event Representations for Automated Story Generation with Deep Neural Nets. *CoRR* abs/1706.01331 (2017). <http://arxiv.org/abs/1706.01331>
- [13] Neil Duncan McIntyre and Mirella Lapata. 2010. Plot Induction and Evolutionary Search for Story Generation. In *ACL 2010, Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, July 11-16, 2010, Uppsala, Sweden*. 1562–1572. <http://www.aclweb.org/anthology/P10-1158>
- [14] Alexander H. Miller, Will Feng, Adam Fisch, Jiasen Lu, Dhruv Batra, Antoine Bordes, Devi Parikh, and Jason Weston. 2017. ParlAI: A Dialog Research Software Platform. *CoRR* abs/1705.06476 (2017). <http://arxiv.org/abs/1705.06476>
- [15] Nikola Mrksic, Diarmuid Ó Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve J. Young. 2017. Neural Belief Tracker: Data-Driven Dialogue State Tracking. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*. 1777–1788. <https://doi.org/10.18653/v1/P17-1163>
- [16] Vitobha Munigala, Srikanth Tamilselvam, and Anush Sankaran. 2017. "Let me convince you to buy my product ...". In *Proceedings of Workshop on Machine Learning for Creativity, SIGKDD, Nova Scotia, Canada, August 2017 (ML4Creativity 17)*.
- [17] Sebastian Neumaier, Vadim Savenkov, and Svitlana Vakulenko. 2017. Talking Open Data. *CoRR* abs/1705.00894 (2017). <http://arxiv.org/abs/1705.00894>
- [18] Sebastian Neumaier, Jürgen Umbrich, and Axel Polleres. 2017. Lifting Data Portals to the Web of Data. In *Workshop on Linked Data on the Web co-located with 26th International World Wide Web Conference (WWW 2017)*. <http://ceur-ws.org/Vol-1809/article-03.pdf>
- [19] Filip Radlinski and Nick Craswell. 2017. A Theoretical Framework for Conversational Search. In *Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval, CHIIR 2017, Oslo, Norway, March 7-11, 2017*. 117–126. <https://doi.org/10.1145/3020165.3020183>
- [20] Mark O. Riedl. 2016. Computational Narrative Intelligence: A Human-Centered Goal for Artificial Intelligence. *CoRR* abs/1602.06484 (2016). <http://arxiv.org/abs/1602.06484>
- [21] Paul Thomas, Daniel McDuff, Mary Czerwinski, and Nick Craswell. 2017. MISC: A data set of information-seeking conversations. (2017).
- [22] Johanne R Trippas, Damiano Spina, Lawrence Cavedon, and Mark Sanderson. 2017. How Do People Interact in Conversational Speech-Only Search Tasks: A Preliminary Analysis. In *Proceedings of the 2017 ACM on Conference on Human Information Interaction and Retrieval (CHIIR)*. ACM.
- [23] Ryen W. White. 2016. *Interactions with Search Systems*. Cambridge University Press.
- [24] Ryen W. White and Resa A. Roth. 2009. *Exploratory Search: Beyond the Query-Response Paradigm*. Morgan & Claypool Publishers. <https://doi.org/10.2200/S00174ED1V01Y200901ICR003>
- [25] Jason D. Williams, Antoine Raux, and Matthew Henderson. 2016. The Dialog State Tracking Challenge Series: A Review. *D&D* 7, 3 (2016), 4–33. <http://dad.uni-bielefeld.de/index.php/dad/article/view/3685>
- [26] Yuan Cao Zhang, Diarmuid Ó Séaghdha, Daniele Quercia, and Tamas Jambor. 2012. Auralist: introducing serendipity into music recommendation. In *Proceedings of the Fifth International Conference on Web Search and Web Data Mining, WSDM 2012, Seattle, WA, USA, February 8-12, 2012*. 13–22. <https://doi.org/10.1145/2124295.2124300>