

SIGIR 2017 Tutorial on Health Search (HS2017)

A Full-day from Consumers to Clinicians

Guido Zuccon
Queensland University of Technology
Brisbane, Australia
g.zuccon@qut.edu.au

Bevan Koopman
Australian e-Health Research Centre, CSIRO
Brisbane, Australia
bevan.koopman@csiro.au

ABSTRACT

The HS2017 tutorial will cover topics from an area of information retrieval (IR) with significant societal impact – health search. Whether it is searching patient records, helping medical professionals find best-practice evidence, or helping the public locate reliable and readable health information online, health search is a challenging area for IR research with an actively growing community and many open problems. This tutorial will provide attendees with a full stack of knowledge on health search, from understanding users and their problems to practical, hands-on sessions on current tools and techniques, current campaigns and evaluation resources, as well as important open questions and future directions.

CCS CONCEPTS

•Information systems → Expert search;

ACM Reference format:

Guido Zuccon and Bevan Koopman. 2017. SIGIR 2017 Tutorial on Health Search (HS2017). In *Proceedings of ACM SIGIR conference, Tokyo, Japan, July 2017 (SIGIR'17)*, 5 pages.
DOI: 10.475/123_4

1 INTENDED AUDIENCE

Researchers of all levels seeking to understand the challenges, tasks and recent developments in information retrieval related to health (health search), be it consumer-oriented search, clinician search, or biomedical search. No prior knowledge in health search is required, making this tutorial ideal for those unfamiliar with this domain. The tutorial is also suitable for those familiar with health search as they will acquire insights from a hands-on session. The tutorial will also provide an analysis of successes and failures of current techniques, and an outline of the opportunities for IR research in the health domain.

2 MOTIVATION AND OVERVIEW

With modern medicine increasingly reliant on information technology, the demand for IR systems that search medical content has grown significantly. The increasing need to retrieve medical advice (by both consumers and clinicians), and the adoption of electronic medical records are two factors driving the demand for

health search. IR research has much to offer here by developing new tools and techniques specific to this domain [17].

The range of health information available (primary research sources, secondary research sources, patient records, web pages and popular publications, etc.), the range of end users (health consumers, different clinicians – general practitioners, specialists, researchers, etc.), and the range of tasks (searching evidence-based-medicine literature [36], searching patient records and cohort selection [43], searching for medical advice on the Web [54], searching the literature for drug-drug interactions and co-morbidities [9], searching for clinical trials [23], etc.) leads to complex requirements that often require novel solutions to these different problems.

The key challenge in health search is how to bridge the *semantic gap*: the mismatch between the raw data and the way a human being interprets it. Although particularly prevalent in health search, the semantic gap problem is found in all domains [1]; advances in health search can thus advance the whole field. Key challenges include: how to leverage semantics and domain-knowledge resources for a better representation of documents and information need [26]; what characterises relevance, in particular how topicality is complemented by other dimensions of relevance [52] (understandability, authoritativeness, etc.), how bias and time pressure affects perception of relevance and decisions [33, 48] and how these influences the search process and evaluation.

Advances in health search will require familiarity with the tasks, users, successes, failures, and domain-specific resources. This tutorial will introduce researchers to the challenges and opportunities in health search, providing insights into current techniques and their results. It will also offer a hands-on overview of tools specific to the health domain made available by the clinical informatics and natural language processing communities.

This tutorial also draws on complementary efforts from other computer science fields in the health domain. Efforts from the clinical informatics and natural language processing communities have produced a wide array of tools that can dovetail with IR techniques. High quality domain knowledge resources (e.g., the UMLS¹ meta-thesaurus and SNOMED CT ontology²) have been developed to encode medical knowledge – these can be used for reasoning and inference within IR techniques [25, 53]. Natural language processing resources that identify medical concepts (from the aforementioned domain-knowledge resources) from free text have been developed (e.g., Metamap [5]). Similar tools have been created to extract other information from medical documents, such as negations, assertions and medications [15, 44]. The tutorial will provide

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).

SIGIR'17, Tokyo, Japan

© 2017 Copyright held by the owner/author(s). 123-4567-24-567/08/06...\$15.00

DOI: 10.475/123_4

¹<https://www.nlm.nih.gov/research/umls/>

²<http://www.snomed.org/snomed-ct>

hands-on demonstrations of how these tools and techniques can be exploited by IR systems.

IR has a long history of rigorous empirical evaluation; this is also the case for work related to health search. This tutorial will cover topics specific to health search evaluation: available test collection, evaluation resources, evaluation campaigns (TREC, CLEF, etc.), as well as insights on successes, failures and difficulties encountered. From an IR research perspective, novel aspects will include the consideration of multiple dimensions of relevance and new frameworks to evaluate systems based on task completion rather than relevance.

3 OBJECTIVES

The main aims of HS2017 will be to:

- (1) Summarise the basics of search in the health domain;
- (2) Present the different end user requirements for multiple user groups interested in health search, including tasks;
- (3) Provide an overview of the current use of IR techniques in the health domain;
- (4) Provide a hands-on introduction to domain-specific tools which can be exploited in health search;
- (5) Present resources and campaigns for evaluation in health search, including novel evaluation approaches;
- (6) Present challenges and opportunities for further research in the health domain and discuss how these could be met.

This knowledge will allow IR researchers to identify promising ways of applying their work to the health domain, allowing them to contribute to a domain of rapidly growing importance.

4 FORMAT AND SCHEDULE

4.1 Session 1: Background and theory

Introduction to the health domain and to the tutorial

The tutorial begins with an introduction to IR in health, giving an overview of the topics that will be covered in the tutorial and why they are important.

Duration: 15m

Types of health information

Health information comes in a myriad of forms. This section covers the characteristics of different types of health information sources important for health search. These range from patient-based information (e.g., electronic health records), knowledge-based information (e.g., scientific papers), through to consumer-based information (e.g., patient forums on the web). Also included are sources of domain knowledge such as medical ontologies, terminologies and classification systems, all of which are playing an increasingly important role in state-of-the-art IR systems.

Duration: 30m

End users and tasks

An analysis of the end user (from consumer [10, 49, 54] to clinician [34]) characteristics and tasks in health search is presented. Some groups of end users are addressed, and the information needs and search tasks they undertake are described. For example, the

group of “physicians” can be divided into groups ranging from general practitioners operating from individual practices who require practice-oriented secondary literature, through specialists in a medical domain who wish to access more specialised works in this domain, to research physicians who have an interest in the primary literature.

Duration: 45m

4.2 Session 2: Techniques and methods

Methods in health search

This section covers the state-of-the-art in health search, summarising the most important research methods and results in this area with respect to the different tasks discussed in the previous session and highlighting common trends across tasks. This session will cover methods across different health search tasks, including: query expansion and reformulation [2, 32, 39, 42, 50, 51], use of domain knowledge and inference mechanisms [12, 24, 26, 29, 53], learning to rank and other learning methods [3, 6, 30, 31, 41], task-based information [21, 28], and specifically handling clinical text [7, 8, 16, 20, 22, 27]. We provide examples of health search systems already in use and lessons that can be drawn from their use.

Duration: 90m

4.3 Session 3: Practical

Use of clinical NLP tools for medical IR: Hands-on session

The clinical informatics and clinical natural language processing community have developed a number of tools for extracting clinical information free text. For example, statistical and dictionary-based named entity recognition systems have been developed to identify medical entities [5, 18, 40]; algorithms have been developed to extract associations, relationships and contexts and to enhance textual content with semantic information [11, 15, 47]. This part of the tutorial presents an overview of such techniques, providing a hands-on demonstration of how these tools work. In addition, the tutorial provides an outlook at how these tools have been used in the literature or can be integrated to enhance information representation and the whole information retrieval process.

Duration: 90m

4.4 Session 4: Evaluation and future directions

Evaluation

The tasks and challenges in evaluating health search are covered in this section. Evaluation that considers multiple dimensions of relevance (topicality, reliability, understandability [52], bias [48]) is presented. We touch on new frameworks to evaluate systems based on task completion rather than relevance. Evaluation campaigns and resources in this domain are presented, including TREC Medical Records Track [43, 45, 46], TREC Clinical Decision Support Track [36–38], CLEF eHealth (consumer health search [13, 14, 35, 55] and as of 2017 search systems for the compilation of systematic reviews), i2b2 Shared Task Challenges³, ALTA Shared Task

³<https://www.i2b2.org/NLP/>

(Query-based summarisation for evidence-based medicine) [4], clinical trial retrieval [23], and the use of ICD coded data as an automated relevance judgement mechanism [19].

Duration: 50m

Open challenges for health search and conclusions (including discussion)

There are many open problems in health search which are fertile ground for information retrieval (IR) research. Examples include: (i) searching for “similar” anonymised patient records or “similar” medical images within a hospital to assist in diagnosis or treatment; (ii) linking treatment guidelines to patient records based on their content; (iii) searching within a patient record to obtain an overview of the medical history. This part of the tutorial briefly presents some of the areas of medicine and consumer-health experience that have a potential to be improved through the use of IR techniques, leading to a discussion with the participants on meeting these challenges.

Duration: 40m

5 RELEVANCE TO THE IR COMMUNITY

Health related topics have become a common theme within IR. A number of venues, including SIGIR, have dedicated workshops, tutorials or tracks dedicated to health search. In recent times, SIGIR has had both tutorials (2012) and workshops (2014 and 2016). Similarly, other venues such as WSDM, WWW, KDD, and ACL have all hosted health related tutorials or workshops. This shows the interest from the community in health search. Even so, a number of important aspects of health search were never covered and are thus the focus of this tutorial:

- This tutorial will clearly provide an understanding of the users, their information needs, tasks and challenges that exist in this domain. This is critical as many of these differ from other domains.
- Practical sessions will provide attendees with hands-on experience with health search techniques, tools and problems.
- The tutorial will conclude with an analysis of open questions in the domain.

It is our intention to provide attendees with a full stack of knowledge in health search, from understanding the problems, practical solutions and fruitful areas of future work. It is our hope that this will equip them to contribution to an area of growing interest in the field with significant societal impact.

A PRESENTERS

Guido Zuccon

Guido is a lecturer within the School of Electrical Engineering and Computing Science at the Queensland University of Technology. His research interests include formal models of search and evaluation methods, in particular applied to health search. Guido has actively contributed to the areas of document ranking, search result diversification, formal models, and evaluation. Guido is the IR Task leader for the CLEF eHealth Evaluation Lab, a shared-task that aims to evaluate systems for consumer health search. He has

already presented tutorials at SIGIR 2015, CIKM 2016 and ICTIR 2016 on formal models of search.

Guido received a Ph.D. in Computing Science from the University of Glasgow in 2012. Before joining the Queensland University of Technology as a lecturer in 2014, he was a postdoctoral research fellow at the CSIRO, Australia, working on health search technology at intersection of academic research and applied research with real customers.

Bevan Koopman

Bevan is a Research Scientist at the Australian e-Health Research Centre, part of the Commonwealth Science and Industrial Research Organisation (CSIRO), based in Brisbane, Australia. He leads projects dedicated to health search, that is the research and development of novel search engine technology to improve access, retrieval and analysis of different health data. He is working on search engines for electronic patient records and search engines for evidence-based medicine. This work is done at intersection of academic research and applied research with real customers.

He received a PhD from the Queensland University of Technology in 2014 with a thesis titled *Semantic search as inference: applications in health informatics* and holds a visiting researcher position at that university. He also spent a number of years in industry and will therefore provide both an academic and commercial expertise to this tutorial.

B TYPE OF SUPPORT MATERIALS TO BE SUPPLIED TO ATTENDEES

A selected list of references is presented at the end of this proposal. In addition to those, the materials will consist of (1) slides and the corresponding hand-outs, (2) an annotated bibliography of works in health search, (3) a virtual machine distribution and/or Git repository containing software and example programs that will be demonstrated in the hands-on session.

Participants will require WiFi access and power for the practical sessions.

REFERENCES

- [1] J. Allan, B. Croft, A. Moffat, and M. Sanderson. 2012. Frontiers, challenges, and opportunities for information retrieval: Report from SWIRL 2012 the second strategic workshop on information retrieval in Lorne. In *ACM SIGIR Forum*, Vol. 46. ACM, 2–32.
- [2] Mohammad Alsulmi and Ben Carterette. 2016. Improving clinical case search using semantic based query reformulations. In *Bioinformatics and Biomedicine (BIBM), 2016 IEEE International Conference on*. IEEE, 694–698.
- [3] Mohammad Alsulmi and Ben Carterette. 2016. Learning to predict the performance of clinical queries using an integrated approach. In *Bioinformatics and Biomedicine (BIBM), 2016 IEEE International Conference on*. IEEE, 930–937.
- [4] I. Amini, D. Martinez, and D. Molla. 2012. Overview of the ALTA 2012 Shared Task. In *Australasian Language Technology Workshop*, Vol. 7. 7–9.
- [5] A.R. Aronson. 2001. Effective mapping of biomedical text to the UMLS Metathesaurus: the MetaMap program. In *Proceedings of the AMIA Symposium*. American Medical Informatics Association, 17.
- [6] Saeid Balaneshin-kordan and Alexander Kotov. 2016. Optimization method for weighting explicit and latent concepts in clinical decision support queries. In *Proceedings of the 2016 ACM on International Conference on the Theory of Information Retrieval*. ACM, 241–250.
- [7] Florian Boudin, Jian-Yun Nie, and Martin Dawes. 2010. Clinical information retrieval using document and PICO structure. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 822–830.
- [8] Florian Boudin, Jian-Yun Nie, and Martin Dawes. 2010. Positional language models for clinical information retrieval. In *Proceedings of the 2010 Conference*

- on *Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 108–115.
- [9] Peter Bruza and Marc Weeber. 2008. *Literature-based discovery*. Springer Science and Business Media.
 - [10] M.A. Cartright, R.W. White, and E. Horvitz. 2011. Intentions and attention in exploratory health search. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*. ACM, 65–74.
 - [11] W.W. Chapman, W. Bridewell, P. Hanbury, G.F. Cooper, and B.G. Buchanan. 2001. A simple algorithm for identifying negated findings and diseases in discharge summaries. *Journal of biomedical informatics* 34, 5 (2001), 301–310.
 - [12] Dina Demner-Fushman and Jimmy Lin. 2007. Answering Clinical Questions with Knowledge-Based and Statistical Techniques. *Computational Linguistics* 33, 1 (March 2007), 63–103.
 - [13] Lorraine Goeuriot, Gareth JF Jones, Liadh Kelly, Johannes Leveling, Allan Hanbury, Henning Müller, Sanna Salanterä, Hanna Suominen, and Guido Zuccon. 2013. ShARe/CLEF eHealth Evaluation Lab 2013, Task 3: Information retrieval to address patients' questions when reading clinical reports. *CLEF 2013 Online Working Notes* 8138 (2013).
 - [14] Lorraine Goeuriot, Liadh Kelly, Wei Li, Joao Palotti, Pavel Pecina, Guido Zuccon, Allan Hanbury, Gareth JF Jones, and Henning Müller. 2014. Share/clef ehealth evaluation lab 2014, task 3: User-centred health information retrieval. In *Proceedings of CLEF 2014*.
 - [15] Henk Harkema, John N Dowling, Tyler Thornblade, and Wendy W Chapman. 2009. ConText: An algorithm for determining negation, experiencer, and temporal status from clinical reports. *Journal of biomedical informatics* 42, 5 (2009), 839–851.
 - [16] Yun He, Qinmin Hu, Yang Song, and Liang He. 2016. Estimating Probability Density of Content Types for Promoting Medical Records Search. In *European Conference on Information Retrieval*. Springer, 252–263.
 - [17] W. Hersh. 2008. *Information retrieval: a health and biomedical perspective*. Springer.
 - [18] Yang Jin, Ryan T McDonald, Kevin Lerman, Mark A Mandel, Steven Carroll, Mark Y Liberman, Fernando C Pereira, Raymond S Winters, and Peter S White. 2006. Automated recognition of malignancy mentions in biomedical literature. *BMC bioinformatics* 7, 1 (2006), 492.
 - [19] Bevan Koopman, Peter Bruza, Laurianne Sitbon, and Michael Lawley. 2011. Evaluating medical information retrieval. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval (SIGIR '11)*. ACM, New York, NY, USA, 1139–1140. DOI:<http://dx.doi.org/10.1145/2009916.2010088>
 - [20] Bevan Koopman, Peter D Bruza, Laurianne Sitbon, and Michael J Lawley. 2010. Analysis of the effect of negation on information retrieval of medical data. In *Proceedings of 15th Australasian Document Computing Symposium (ADCS)*. University of Melbourne.
 - [21] Bevan Koopman, Jack Russell, and Guido Zuccon. 2017. Task-oriented search for evidence-based medicine. *International Journal on Digital Libraries* (2017), 1–13.
 - [22] Bevan Koopman and Guido Zuccon. 2014. Understanding negation and family history to improve clinical information retrieval. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*. ACM, 971–974.
 - [23] Bevan Koopman and Guido Zuccon. 2016. A Test Collection for Matching Patients to Clinical Trials. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. ACM, 669–672.
 - [24] Bevan Koopman, Guido Zuccon, Peter Bruza, Laurianne Sitbon, and Michael Lawley. 2012. Graph-based concept weighting for medical information retrieval. In *Proceedings of the Seventeenth Australasian Document Computing Symposium*. ACM, 80–87.
 - [25] Bevan Koopman, Guido Zuccon, Peter Bruza, Laurianne Sitbon, and Michael Lawley. 2015. Information Retrieval as Semantic Inference: A Graph Inference Model applied to Medical Search. *Information Retrieval* 19, 1 (2015), 6–37.
 - [26] Bevan Koopman, Guido Zuccon, Peter Bruza, Laurianne Sitbon, and Michael Lawley. 2016. Information retrieval as semantic inference: a Graph Inference model applied to medical search. *Information Retrieval Journal* 19, 1-2 (2016), 6–37.
 - [27] Nut Limsopatham, Craig Macdonald, Richard McCreadie, and Iadh Ounis. 2012. Exploiting term dependence while handling negation in medical search. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*. ACM, 1065–1066.
 - [28] Nut Limsopatham, Craig Macdonald, and Iadh Ounis. 2013. A Task-Specific Query and Document Representation for Medical Records Search. In *Proceedings of the 35th European Conference on Information Retrieval (ECIR)*. Moscow, Russia.
 - [29] Nut Limsopatham, Craig Macdonald, and Iadh Ounis. 2013. Inferring conceptual relationships to improve medical records search. In *Proceedings of the 10th Conference on Open Research Areas in Information Retrieval*. 1–8.
 - [30] Nut Limsopatham, Craig Macdonald, and Iadh Ounis. 2013. Learning to combine representations for medical records search. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*. ACM, 833–836.
 - [31] Nut Limsopatham, Craig Macdonald, and Iadh Ounis. 2013. Learning to selectively rank patients' medical history. In *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*. ACM, 1833–1836.
 - [32] Xiaojie Liu, Jian-Yun Nie, and Alessandro Sordani. 2016. Constraining Word Embeddings by Prior Knowledge—Application to Medical Information Retrieval. In *Information Retrieval Technology*. Springer, 155–167.
 - [33] Nina Mishra, Ryen W White, Samuel Jeong, and Eric Horvitz. 2014. Time-critical search. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*. ACM, 747–756.
 - [34] João Palotti, Allan Hanbury, Henning Müller, and Charles E Kahn Jr. 2016. How users search and what they search for in the medical domain. *Information Retrieval Journal* 19, 1-2 (2016), 189–224.
 - [35] João RM Palotti, Guido Zuccon, Lorraine Goeuriot, Liadh Kelly, Allan Hanbury, Gareth JF Jones, Mihai Lupu, and Pavel Pecina. 2015. CLEF eHealth Evaluation Lab 2015, Task 2: Retrieving Information About Medical Symptoms.. In *CLEF (Working Notes)*.
 - [36] Kirk Roberts, Matthew Simpson, Dina Demner-Fushman, Ellen Voorhees, and William Hersh. 2016. State-of-the-art in biomedical literature retrieval for clinical cases: a survey of the TREC 2014 CDS track. *Information Retrieval Journal* 19, 1-2 (2016), 113–148.
 - [37] Kirk Roberts, Matthew S Simpson, Ellen Voorhees, and William R Hersh. 2015. Overview of the TREC 2015 Clinical Decision Support Track. In *Text REtrieval Conference (TREC)*.
 - [38] Matthew S Simpson, Ellen M Voorhees, and William Hersh. 2014. Overview of the TREC clinical decision support track. In *Text REtrieval Conference (TREC)*.
 - [39] Luca Soldaini, Arman Cohan, Andrew Yates, Nazli Goharian, and Ophir Frieder. 2015. Retrieving Medical Literature for Clinical Decision Support. In *ECIR*.
 - [40] Luca Soldaini and Nazli Goharian. 2016. Quickkums: a fast, unsupervised approach for medical concept extraction. In *MedIR Workshop, SIGIR*.
 - [41] Luca Soldaini and Nazli Goharian. 2017. Learning to Rank for Consumer Health Search: a Semantic Approach. In *In Proceedings of ECIR 2017*.
 - [42] Luca Soldaini, Andrew Yates, Elad Yom-Tov, Ophir Frieder, and Nazli Goharian. 2016. Enhancing web search in the medical domain via query clarification. *Information Retrieval Journal* 19, 1-2 (2016), 149–173.
 - [43] N Tracy Edinger, Aaron M Cohen, Steven Bedrick, Kyle Ambert, and William Hersh. 2012. Barriers to retrieving patient information from electronic health record data: failure analysis from the TREC medical records track. (2012).
 - [44] Özlem Uzuner, Imre Solti, and Eithon Cadag. 2010. Extracting medication information from clinical text. *Journal of the American Medical Informatics Association* 17, 5 (2010), 514–518.
 - [45] E. Voorhees and R. Tong. 2011. Overview of the TREC 2011 medical records track. In *The Twentieth Text REtrieval Conference Proceedings TREC*.
 - [46] Ellen M Voorhees and William R Hersh. 2012. Overview of the TREC 2012 Medical Records Track. In *TREC*.
 - [47] Byron C Wallace, Joël Kuiper, Aakash Sharma, Mingxi Brian Zhu, and Iain J Marshall. 2016. Extracting PICO sentences from clinical trial reports using supervised distant supervision. *Journal of Machine Learning Research* 17, 132 (2016), 1–25.
 - [48] Ryen W White and Eric Horvitz. 2009. Cyberchondria: studies of the escalation of medical concerns in web search. *ACM Transactions on Information Systems (TOIS)* 27, 4 (2009), 23.
 - [49] Ryen W White and Eric Horvitz. 2009. Experiences with web search on medical concerns and self diagnosis. In *AMIA*.
 - [50] Dongqing Zhu and Ben Carterette. 2012. Improving health records search using multiple query expansion collections. In *Bioinformatics and Biomedicine (BIBM), 2012 IEEE International Conference On*. IEEE, 1–7.
 - [51] Dongqing Zhu, Stephen Wu, Ben Carterette, and Hongfang Liu. 2014. Using large clinical corpora for query expansion in text-based cohort identification. *Journal of biomedical informatics* 49 (2014), 275–281.
 - [52] Guido Zuccon. 2016. Understandability biased evaluation for information retrieval. In *European Conference on Information Retrieval*. Springer, 280–292.
 - [53] Guido Zuccon, Bevan Koopman, Anthony Nguyen, Deanne Vickers, and Luke Butt. 2012. Exploiting medical hierarchies for concept-based information retrieval. In *Proceedings of the Seventeenth Australasian Document Computing Symposium (ADCS '12)*. ACM, New York, NY, USA, 111–114. DOI:<http://dx.doi.org/10.1145/2407085.2407100>
 - [54] Guido Zuccon, Bevan Koopman, and Joao Palotti. 2015. Diagnose this if you can. In *European Conference on Information Retrieval*. Springer, 562–567.
 - [55] Guido Zuccon, Joao Palotti, Lorraine Goeuriot, Liadh Kelly, Mihai Lupu, Pavel Pecina, Henning Mueller, Julie Budaher, and Anthony Deacon. 2016. The IR Task at the CLEF eHealth evaluation lab 2016: user-centred health information retrieval. In *CLEF 2016-Conference and Labs of the Evaluation Forum*, Vol. 1609. 15–27.