WSDM 2019 Tutorial on Health Search (HS2019)

A Full-day from Consumers to Clinicians

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

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

ABSTRACT

The HS2019 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 information on current tools and techniques, evaluation resources, as well as important open questions and future directions. Tutorial material is available at https://ielab.io/health-search-tutorial/.

ACM Reference Format:
Bevan Koopman and Guido Zuccon. 2019. WSDM 2019 Tutorial on Health Search (HS2019): A Full-day from Consumers to Clinicians. In *The Twelfth ACM International Conference on Web Search and Data Mining (WSDM '19), February 11–15, 2019, Melbourne, VIC, Australia.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3289600.3291379

1 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 [9].

The range of health information available (primary research sources, secondary research sources, patient records, web pages and popular publications, etc.), plus the range of end users (health consumers, different clinicians — general practitioners, specialists, researchers, etc.), and the range of tasks (searching evidence-based-medicine literature [22], searching patient records and cohort selection [30], searching for medical advice on the Web [34], searching the literature for drug-drug interactions and co-morbidities, searching for clinical trials [13], searching literature to form systematic reviews [10], etc.) all leads to complex requirements that often require novel solutions to these different problems.

The key challenge in health search is how to bridge the *sematic 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 needs [14]; what characterises relevance, in particular how topicality is complemented by other dimensions of relevance [32] (understandability, authoritativeness, etc.), how bias and time pressure affect perception of relevance and decisions [21, 31] 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 UMLS[1] metathesaurus and SNOMED CT ontology[2]) have been developed to encode medical knowledge — these can be used for reasoning and inference within IR techniques [14, 33]. Natural language processing resources that identify medical concepts (from the aforementioned domain-knowledge resources) from free text have been developed (e.g., Metamap [3]). Similar tools have been created to extract other information from medical documents, such as negations, assertions and medications. The tutorial will provide 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 in health search. This tutorial will cover topics specific to health search evaluation: available test collections, evaluation resources, evaluation campaigns (TREC, CLEF, etc.), as well as insights on successes, failures and difficulties encountered.

Health related topics have become a common theme within IR. A number of venues, including WSDM, have workshops, tutorials or tracks dedicated to health search. Even so, a number of important aspects of health search were never covered and are thus the focus of this tutorial. Specifically, this tutorial will:

WSDM '19, February 11–15, 2019, Melbourne, VIC, Australia

ACM ISBN 978-1-4503-5940-5/19/02.
https://doi.org/10.1145/3289600.3291379

[1] https://www.nlm.nih.gov/research/umls/
[2] http://www.snomed.org/snomed-ct

- 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.
- provide attendees with hands-on experience with health search techniques, tools and problems.
- present an analysis of open questions in the domain.

2 OBJECTIVES

The main aims of HS2019 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 will allow IR researchers to identify promising ways of applying their work to health problems, allowing them to contribute to a domain of rapidly growing importance.

3 TUTORIAL TOPICS AND STRUCTURE

Session 1: Types of health information, end users and tasks. This section covers the characteristics of different types of health information sources important for health search, e.g. patient related (e.g., electronic health records [13, 30]), knowledge related (e.g., scientific papers [10, 22]), consumer related (e.g., patient forums [35]). We also discuss sources of domain knowledge such as medical ontologies, terminologies and classification systems. In addition, an analysis of the end user (from consumers [6, 31, 34] to clinicians [20, 28]) characteristics and tasks in health search is presented.

Session 2 &3: Techniques, methods and tools. 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 tasks in health search, highlighting common trends across tasks. This will cover techniques such as query expansion and reformulation, (e.g. [23, 27, 29], use of domain knowledge and inference mechanisms (e.g. [7, 8, 14, 17, 24, 33]), learning to rank and other learning methods (e.g., [2, 4, 18, 19, 23, 26]), task-based information (e.g., [11, 16]), and specifically handling clinical text (e.g. [5, 12, 15, 25]). This part of the tutorial will also present an overview of tools for extracting clinical and biomedical information, providing a hands-on demonstration of how these tools work and an outlook of how they have been used to enhance information representation and the whole IR process.

Session 4: Evaluation and open challenges. Tasks and challenges in evaluating health search are covered in this section, including evaluation techniques specific to health search (e.g. [11, 32], and datasets for evaluating health search (e.g. [10, 13, 22, 30, 35]).

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 IR in Lorne. In *SIGIR Forum*, Vol. 46. ACM, 2–32.

[2] Mohammad Alsulmi and Ben Carterette. 2016. Learning to predict the performance of clinical queries using an integrated approach. In *BIBM*. IEEE, 930–937.

[3] A.R. Aronson. 2001. Effective mapping of biomedical text to the UMLS Metathesaurus: the MetaMap program.. In *AMIA*. 17.

[4] Saeid Balaneshin-kordan and Alexander Kotov. 2016. Optimization method for weighting explicit and latent concepts in clinical decision support queries. In *ICTIR*. ACM, 241–250.

[5] Florian Boudin, Jian-Yun Nie, and Martin Dawes. 2010. Clinical information retrieval using document and PICO structure. In *HLT-NAACL*.

[6] 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.

[7] Dina Demner-Fushman and Jimmy Lin. 2007. Answering Clinical Questions with Knowledge-Based and Statistical Techniques. *Computational Linguistics* 33, 1 (March 2007), 63–103.

[8] Travis R Goodwin and Sanda M Harabagiu. 2017. Knowledge Representations and Inference Techniques for Medical Question Answering. *TIST* 9, 2 (2017), 14.

[9] W. Hersh. 2008. *Information retrieval: a health and biomedical perspective*.

[10] Evangelos Kanoulas, Dan Li, Leif Azzopardi, and Rene Spijker. 2017. CLEF 2017 technologically assisted reviews in empirical medicine overview. In *CLEF*.

[11] Bevan Koopman, Jack Russell, and Guido Zuccon. 2017. Task-oriented search for evidence-based medicine. *IJDL* (2017), 1–13.

[12] Bevan Koopman and Guido Zuccon. 2014. Understanding negation and family history to improve clinical information retrieval. In *SIGIR*. ACM, 971–974.

[13] Bevan Koopman and Guido Zuccon. 2016. A Test Collection for Matching Patients to Clinical Trials. In *SIGIR*. ACM, 669–672.

[14] 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. *IRJ* (2016).

[15] Nut Limsopatham, Craig Macdonald, Richard McCreadie, and Iadh Ounis. 2012. Exploiting term dependence while handling negation in medical search. In *SIGIR*.

[16] Nut Limsopatham, Craig Macdonald, and Iadh Ounis. 2013. A Task-Specific Query and Document Representation for Medical Records Search. In *ECIR*.

[17] Nut Limsopatham, Craig Macdonald, and Iadh Ounis. 2013. Inferring conceptual relationships to improve medical records search. In *OAIR*. 1–8.

[18] Nut Limsopatham, Craig Macdonald, and Iadh Ounis. 2013. Learning to combine representations for medical records search. In *SIGIR*. ACM, 833–836.

[19] Gia-Hung Nguyen, Lynda Tamine, Laure Soulier, and Nathalie Souf. 2017. Learning Concept-Driven Document Embeddings for Medical Information Search. In *AIM*. Springer, 160–170.

[20] João Palotti, Allan Hanbury, Henning Müller, and Charles E Kahn Jr. 2016. How users search and what they search for in the medical domain. *IRJ* (2016).

[21] Frances A. Pogacar, Amira Ghenai, Mark D. Smucker, and Charles L.A. Clarke. 2017. The Positive and Negative Influence of Search Results on People's Decisions About the Efficacy of Medical Treatments. In *ICTIR*. 209–216.

[22] Kirk Roberts, Matthew S Simpson, Ellen Voorhees, and William R Hersh. 2015. Overview of the TREC 2015 Clinical Decision Support Track. In *TREC*.

[23] Harrisen Scells and Guido Zuccon. 2018. Generating Better Queries for Systematic Reviews. In *SIGIR*. ACM.

[24] Harrisen Scells, Guido Zuccon, Bevan Koopman, Anthony Deacon, Leif Azzopardi, and Shlomo Geva. 2017. Integrating the Framing of Clinical Questions via PICO into the Retrieval of Medical Literature for Systematic Reviews. In *CIKM*. ACM, 2291–2294.

[25] Harrisen Scells, Guido Zuccon, Bevan Koopman, Anthony Deacon, Leif Azzopardi, and Shlomo Geva. 2017. A Test Collection for Evaluating Retrieval of Studies for Inclusion in Systematic Reviews. In *SIGIR*. 1237–1240.

[26] Luca Soldaini and Nazli Goharian. 2017. Learning to Rank for Consumer Health Search: a Semantic Approach. In *ECIR*.

[27] Luca Soldaini, Andrew Yates, Elad Yom-Tov, Ophir Frieder, and Nazli Goharian. 2016. Enhancing web search in the medical domain via query clarification. *IR* 19, 1-2 (2016), 149–173.

[28] Lynda Tamine and Cecile Chouquet. 2017. On the impact of domain expertise on query formulation, relevance assessment and retrieval performance in clinical settings. *IP&M* 53, 2 (2017), 332–350.

[29] Carla Teixeira Lopes, Dagmara Paiva, and Cristina Ribeiro. 2017. Effects of language and terminology of query suggestions on medical accuracy considering different user characteristics. *JASIST* 68, 9 (2017), 2063–2075.

[30] E. Voorhees and R. Tong. 2011. Overview of the TREC 2011 medical records track. In *TREC*.

[31] Ryen W White and Eric Horvitz. 2009. Cyberchondria: studies of the escalation of medical concerns in web search. *TOIS* 27, 4 (2009), 23.

[32] Guido Zuccon. 2016. Understandability biased evaluation for information retrieval. In *ECIR*. Springer, 280–292.

[33] Guido Zuccon, Bevan Koopman, Anthony Nguyen, Deanne Vickers, and Luke Butt. 2012. Exploiting medical hierarchies for concept-based information retrieval. In *ADCS*. 111–114.

[34] Guido Zuccon, Bevan Koopman, and Joao Palotti. 2015. Diagnose this if you can. In *ECIR*. Springer, 562–567.

[35] 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*, Vol. 1609. 15–27.