

How Users Search and What They Search for in the Medical Domain

Understanding Laypeople and Experts Through Query Logs

João Palotti · Allan Hanbury · Henning Müller · Charles E. Kahn Jr.

Received: date / Accepted: date

Abstract The internet is an important source of medical knowledge for everyone, from laypeople to medical professionals. We investigate how these two extremes, in terms of user groups, have distinct needs and exhibit significantly different search behaviour. We make use of query logs in order to study various aspects of these two kinds of users. The logs from America Online (AOL), Health on the Net (HON), Turning Research Into Practice (TRIP) and American Roentgen Ray Society (ARRS) GoldMiner were divided into three sets: (1) laypeople, (2) medical professionals (such as physicians or nurses) searching for health content and (3) users not seeking health advice. Several analyses are made focusing on discovering how users search and what they are most interested in. One possible outcome of our analysis is a classifier to infer user expertise, which was built. We show the results and analyse the feature set used to infer expertise. We conclude that medical experts are more persistent, interacting more with the search engine. Also, our study reveals that, conversely to what is stated in much of the literature, the main focus of users, both laypeople and professionals, is on disease rather than symptoms. The results of

this article, especially through the classifier built, could be used to detect specific user groups and then adapt search results to the user group.

1 Introduction

Among all topics available on the internet, medicine is one of the most important in terms of impact on the user and one of the most frequently searched. A recent report states that one in three American adult Internet users have sought out health advice online to diagnose a medical condition [18]. This tendency is the same in Europe, where a recent report from the European Commission estimates that 60% of the population have used the Internet to search for health-related information in 2014 [16], with numbers even higher in several member states. Both reports show that the most common tasks performed are either searching for general information on health-related topics, such as diet, pregnancy and exercise, or searching for information on specific injuries or diseases. They also found that mostly the search starts in a search engine and young users are more likely to search for this kind of information.

Physicians are also very active Internet users [31]. PubMed which indexes the biomedical literature reports more than one hundred million users [25], where two-thirds are experts [32]. Nevertheless, studies on how experts search on the Internet for medical content are relatively rare [57].

We divide the users of medical search engines into *laypeople* and *experts*, where laypeople are considered to be searchers that do not have a deep knowledge about the medical topic being searched, and experts do have a deep knowledge about the medical topic being searched. Our assumption is that laypeople wish to see more in-

J. Palotti · A. Hanbury
Vienna University of Technology
Austria
E-mail: palotti,hanbury@ifs.tuwien.ac.at

H. Müller
University of Applied Sciences and Arts Western Switzerland
(HES-SO)
Switzerland
E-mail: henning.mueller@hevs.ch

C. E. Kahn Jr.
University of Pennsylvania
United States
E-mail: charles.kahn@uphs.upenn.edu

troductory material returned as search results, whereas experts wish to see detailed scientific material returned as search results. At first glance, this could easily be interpreted as a division into patients and medically-trained professionals. Nevertheless, it often occurs that a patient or patient's relatives become experts on a disease or condition affecting themselves or a family member, sometimes becoming more knowledgeable in a narrow domain than medically-trained professionals. There is also the case of a medical professional searching in a medical topic outside of his/her main expertise (e.g. a cardiac surgeon looking for information on a skin disease), where the information need may be initially satisfied by less scientific documents, although likely not very basic documents due to the medical background. For these reasons, we specifically avoid defining medical professional and patient classes.

Distinguishing laypeople and experts can significantly improve their interactions with the search engine [52, 38]. Currently, users may get different results for their queries if they are in different locations, but not if they have different levels of expertise. We make the assumption that it is possible to distinguish the level of expertise of the searcher based on the vocabulary used and the search style. While it would be realistic to represent a continuum of expertise levels, we define two classes (laypeople and experts) in this study, allowing us to investigate the most relevant differences between the classes.

Recently, many studies showed successful cases of exploring the user's expertise, in particular for general search engines [52, 44, 11]. Schwarz et al. [44] show that the popularity of a webpage among experts is a crucial feature to help laypeople identify credible websites. Collins-Thompson et al. [11] discuss that re-ranking general search engine results to match the user's skills of readability can provide significant gains, however estimating user profiles is a non-trivial task and needs to be further explored.

This study investigates how users search for medical content, building profiles for experts and laypeople. Understanding the needs of these two distinct groups is important for designing search engines, whether it is used for boosting easy-to-read documents or for suggesting queries to match the search expertise. Additionally, whenever it is possible, we also compare search for medical information with regular search for other topics.

This work is conducted through the analysis of user interactions logged by search engines. Log analysis is unobtrusive and captures the user behaviour in a natural setting [27]. We used Metamap¹, which is the state-

of-the-art tool to recognise and map biomedical text to its corresponding medical concepts, to provide a richer set of information for each query. Little is known in the literature on how to identify medical concepts in short Web queries, therefore we also evaluated Metamap for this task.

In particular, this work addresses the following questions:

1. How suitable is MetaMap for analysing short queries?
2. Which characteristics allow laypeople and experts to be distinguished based on
 - (a) How they search in medical content?
 - (b) What they search for in medical content?
3. To what extent do these characteristics match or disagree with those identified in other published studies?
4. What are the most useful features to automatically infer user expertise through the query logs?

In our analysis, we use health related queries from the America Online (AOL) query log, as well as the Health on the Net (HON) search engine log to represent the logs generated to a significant extent by laypeople. Medical professionals also use general search engines to seek health content, however their queries are drowned in the laypeople queries. White et al. [52], for example, hypothesise that search leading to PubMed was made by experts. Using this hypothesis, only 0.004% of the whole AOL log was issued by medical professionals (also referred to as experts).

Besides the fact that PubMed is more frequently used in a research environment rather than in a clinical environment [31], it is also frequently visited by laypeople [32]. Therefore, we use the logs from the evidence-based search engine TRIP Database and the radiology image search engine American Roentgen Ray Society (ARRS) GoldMiner to represent queries entered by physicians usually when facing a practical problem.

Several analyses are presented: from general statistics of the logs to complex inference on what is the search focus in each individual search session. We contrast our results with others from the literature and provide our interpretation for each phenomenon found.

The remainder of this paper is organised as follows. Section 2 presents a literature review and positions our work with respect to other articles in the literature. In Section 3, we describe the datasets used and the preprocessing steps applied. In Section 4 we present and evaluate MetaMap, the tool used to enrich the information contained in the query logs. We start our analysis in Section 5, where we examine the general user behaviour and the most popular queries, terms and topics searched. In Section 6, we introduce the concept of

¹ <http://metamap.nlm.nih.gov/>

search session into our analysis. In Section 7, we present a Random Forest classifier to infer user expertise and analyse the feature’s importance. Section 8 presents our findings and limitations. Finally, conclusions and future work are presented in Section 9.

2 Related Work

As soon as modern search engines appeared, the first studies on query logs started [29,45]. Jansen et al. [29] and Silverstein et al. [45] analysed the logs from Excite and Altavista respectively, popular search engines at that time. Both articles point out some important results such as the fact that the vast majority of users issue only one single query and rarely access any result page beyond the first one.

The most recent general search engine to disclose query logs to researchers was America Online (AOL) in 2006 [41]. The AOL data were afterwards used in various studies, such as Brenes et al. [7], which provides methods to group users and their intents, and Torres et al. [14], who analyse queries targeting children’s content. In this work, we compare the analysis made in the literature for general search engines [29,45] with medical domain search engines, and we adopt a method similar to [14] to divide the AOL logs into queries related or not to health. It is important to mention that the AOL log had known privacy problems in the past, resulting in some users being identified even though the logs were supposedly anonymised. Despite this problem, we opt to use this dataset for several reasons. One reason is that it can be freely downloaded, as well as the code used for all the experiments of this paper, making the experiments reproducible². Another reason is that studies of how medical annotation tools such as MetaMap perform in the wild are not well known. Finally, in the absence of a more recent large search engine query log we consider that the AOL logs are still the best choice for researchers in academia. A complete reference of the previous 20 years of research on log analysis and its applications is well described by Silvestri [46].

There are a number of studies analysing query logs in the medical domain. We highlight here some important work for this research, including work based on general search engines [47,52,8,54], as well as specialised ones [22,25,48,34,59]. We also highlight some important work on user expertise and behaviour. Figure 1 depicts each one of these areas, including relevant work on general search for non-health-related content. For a matter of organisation, we divide the rest of this section into 3 parts, one for each topic. As shown in

Figure 1, some papers may be relevant to more than one topic.

2.1 General Search Engines

We describe here studies on health-related query logs in general search engines, starting with Spink et al. [47], who studied medical queries issued in 2001 in Excite and AlltheWeb.com. They showed that medical web search was decreasing since 1999, suggesting that users were gradually shifting from general-purpose search engines to specialised sites for health-related queries. Also, they found that health-related queries were equivalent in length, complexity and lack of reformulation to general web searching.

More recently, White and Horvitz [53,54] studied how users start looking for a simple symptom and end up searching for a serious disease, a phenomenon they named cyberchondria. They used the logs of the Windows Live Toolbar to obtain their data and list of keywords to annotate symptoms and diseases in queries, while we used the US National Library of Medicine MetaMap to do the same. Similar to our work, they define user sessions as a series of queries followed by a period of user inactivity of more than 30 minutes and they made use of the Open Directory Project (ODP) hierarchy to identify medical sessions.

Another important work is Cartright et al. [8]. The authors presented a log-based study of user behaviour when searching for health information online. The authors classified user queries into three classes: symptoms, causes and remedy. They analysed the change of search focus along a session, and showed that it is possible to build a classifier to predict what is the next focus of a user in a session. We decided to use the same classes in order to make our study comparable, however we used the semantic annotator of MetaMap instead of hand coded rules.

Not studying the query logs, but the ranking lists of major general search engines, Wang et al. [50] compared the results of Google, Yahoo!, Bing, and Ask.com for one single query *breast cancer*. Among their conclusions is the fact that results provided rich information and highly overlapped between the search engines. The overlap between any two search engines was about half or more. Another work that compares a large number of search engines is Jansen and Spink [28], in which 9 search engines with logs varying from 1997 to 2002 were used. Nevertheless, the latter did not focus on medical queries.

² <https://github.com/joaopalotti/logAnalysisJournal>

2.2 Expertise and Search Engines

One of the first studies to report how expertise influences the process of search dates from the 1990's. In this work, Hsieh-Yee [24] reported that experienced library science students could use more thesauri, synonymous terms, combinations of search terms and spend less time monitoring their searches than novices. Later, Bhavnani [5] studied search expertise in the medical and shopping domains. He reported that experts in a topic can easily solve the task given even without using a search engine, because they already knew which website was better adapted to fill their needs. Bhavnani also reported that experts started their search by using websites such as MedlinePlus³, instead of a major search engine, while laypeople started with Google.

White et al. [52] showed a log-based analysis of expertise in four different domains (medicine, finance, law, and computer science), developing an expertise classifier based on their analysis. Apart from showing that it is possible to predict user expertise based on their behaviour, they showed that experts have a higher success rate only in their domain of expertise, with success in a session being defined as a clicked URL as the final event in a session. Therefore, an expert in finance would have a comparable or worse success rate in medicine than a non-expert. An important difference between our work and White's work is the approach used to separate experts from non-experts. They assume that search leading to PubMed was made by medical experts and search leading to ACM Digital library (ACM-DL)⁴ was made by computer science experts. In the medical domain this is a weak premise for two reasons: (1) it is estimated that one-third of PubMed users are laypeople [32], (2) PubMed is more important for medical researchers than practitioners [31]. Tracing a parallel between medicine and computer science, a general practitioner would be like a software developer that does not necessarily need to consult the ACM-DL (the correspondent for PubMed) to perform his/her work. One could manually expand the list of expert sites to include, for example, StackOverflow⁵ or an API website for experts in CS and treatment guidelines or drug information sites for medicine but it would be a laborious task and unstable over time. Hence, to cope with this challenge, we use the logs of different search engines made for distinct audiences.

³ MedlinePlus is a web-based consumer health information system developed by the American National Library of Medicine (NLM): <http://www.medlineplus.gov/>

⁴ <http://dl.acm.org/>

⁵ <http://stackoverflow.com/>

An important user study was conducted by Wildemuth [55]. He evaluated how the search tactics of microbiology students changed over an academic year, while the students' topic knowledge was increasing. The students were asked questions about the topic at 3 different times: before starting the course, when finishing the course, and 6 months after the course. As their expertise increased, the users were able to perform a better term selection for search, being more effective. The most common pattern used across all three occasions was the narrowing of the retrieved result set through the addition of search concepts, while at the beginning users were less effective in the selection of concepts to include in the search and more errors were made in the reformulation of a query. Later, Duggan and Payne [15] explored the domains of music and football to evaluate how the user knowledge of a topic can influence the probability of a user answering factual questions, finding that experts detect unfruitful search paths faster than non-experts.

Recently, there have been a few user studies in user expertise prediction. For example, Zhang et al. [58] and Cole et al. [10] are based on TREC Genomics data. The former employed a regression model to match user self-rated expertise and high level user behaviour features such as the mean time analysing a document and the number of documents viewed. They found that the user's domain knowledge could be indicated by the number of documents saved, the user's average query length, and the average rank position of opened documents. Their model, however, needs to be further investigated because the data was limited, collected in a controlled experiment, and from only one domain. Similarly, but using only eye movement patterns as features, the latter conducted a user study instead of log analysis and employed a linear model and random forests to infer the user expertise level. Their main contribution is demonstrating that models to infer a user's level of domain knowledge without processing the content of queries or documents is possible, however they only performed one single experiment and in one single domain.

2.3 Medical-Specialised Search Engines

For specialised medical search engines, Herskovic et al. [22] analysed an arbitrary day in PubMed, the largest biomedical database in the world. They concluded that PubMed may have a different usage profile than general web search engines. Their work showed that PubMed queries had a median of three terms, one more than what is reported for Excite and Altavista. Subsequently, Dogan et al. [25] studied an entire month of PubMed log data. Their main finding comparing PubMed and

general search engines was that PubMed users are less likely to select results when the result sets increase in size, users are more likely to reformulate queries and are more persistent in seeking information. Whenever possible, our analysis is compared with the statements made for PubMed.

Meats et al. [34] conducted an analysis on the 2004 and 2005 logs of the TRIP Database, together with a usability study with nine users. Their work concluded that most users used a single term and only 12% of the search sessions utilised a Boolean operator, under-utilising the search engine features. Tsikrika et al. [48] examined query logs from ARRS GoldMiner, a professional search engine for radiology images. They studied the process of query modification during a user session, aiming to guide the creation of realistic search tasks for the ImageCLEFmed benchmark. Meats used 620,000 queries and Tsikrika only 25,000, while we use nearly 3 and 9 times more queries from TRIP and GoldMiner, respectively, allowing us to perform a deeper analysis.

Zhang [59] analysed how 19 students solved 12 tasks using MedlinePlus. The tasks were created based on questions from the health section of Yahoo! Answers. Although the log analysis made is very limited due to the artificial scenario created and the small number of users, Zhang could investigate browsing strategies used by users (amount of time searching and/or browsing MedlinePlus) and the users' experience with MedlinePlus (usability, usefulness of the content, interface design) through questionnaires and recording the users performing the tasks. Our study is limited to only the query logs, however a large analysis is made for different websites and the user behaviour is captured in a very natural setting.

2.4 This Work

As illustrated in Figure 1, this work closes a gap. It studies both general and specialised search engines, as well as taking into consideration different user expertise levels. Throughout the rest of this work, we compare our methodology and results with the studies cited in this section.

3 Datasets and Pre-processing Steps

In this section, we describe the datasets used in this study and the preprocessing steps applied to them.

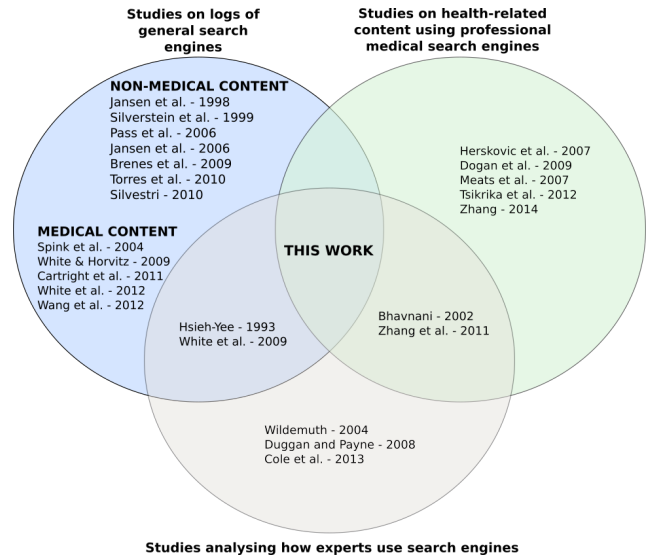


Fig. 1: Our work studies both general and specialised search engines and investigates how users with different expertise levels search for health content

3.1 Sorting the Data by Expertise Level

We make the assumption that experts and laypeople are more likely to use different search engines to satisfy their information needs. Therefore we assume that almost all queries entered into a particular search engine are entered by only one of the two classes of users under consideration. This assumption is justified as we are using search logs from search engines clearly aimed at users of specific expertise. This assumption is also more inclusive than another assumption that has been used to separate medical experts from laypeople: that only searches leading to PubMed were made by medical experts [52]. As discussed in [38], this assumption would only tend to detect medical researchers, as medical practitioners make less use of PubMed [31]. We do not take into account that many users are in between laypeople and experts as levels can vary.

On one extreme, we have AOL laypeople users. There might be a few medical experts using AOL, but their queries are drowned in the laypeople queries. Also focused on patients, HON is a search engine for laypeople searching for reliable health information. The main target audience is laypeople concerned about the reliability of the information they access. On the other extreme, mainly targeting physicians looking for medical evidence, the TRIP database can also be accessed by patients but these few patients might be already considered specialists on their diseases. Finally, the GoldMiner search engine is made by radiologists and for radiologists, patients have practically no use for this kind of information, but a variety of physicians might ac-

cess the system. We position each dataset on an expertise axis in Figure 2, to help understanding how each dataset relates to each other.



Fig. 2: The datasets used here are plotted on an expertise scale. The expertise level increases as a dataset is placed more to the right-hand side of the scale.

3.2 Data

Four query logs from search engines taking free text queries were divided into five datasets in our analysis: two focused on laypeople queries, two made up of queries from medical professionals and one consisted of queries not related to health or medical information.

The query logs that are assumed to consist almost completely of queries submitted by laypeople were obtained from medical-related search in America Online’s search service [41]⁶ and from the Health on the Net Foundation website (HON⁷).

The AOL logs were obtained from March to May 2006. We divided them into two non-overlapping sets: **AOL-Medical** and **AOL-NotMedical**. For this purpose, the click-through information available in the AOL data was used. A common approach to decide what the topic of a URL is, is checking if it is listed in the Open Directory Project (ODP)⁸ [8, 11, 14, 52, 54]. For the clicked URLs that are not present in ODP, some researchers use supervised learning to automatically classify them [11, 52, 54]. However, it is very important to note that this approach cannot be used here, as 47% of the AOL log entries lack the clicked URL information.

Alternative approaches can be designed. One is to keep only queries in which the clicked URL is found in ODP, excluding all the rest. Although valid, this approach results in removing 73% of all queries, as only 27% of the queries had a clicked URL found in ODP. This has a strong impact in the behaviour analysis, such as a strong reduction in the number of queries per session. Another possibility is doing as in Cartright’s work [8], in which a list of symptoms was used to filter sessions on health information. However this approach creates a strong bias when analysing what users are

searching for, as it certainly results in a dataset in which everyone searches for symptoms.

Our solution is based on user sessions – this approach is not as restricted as when analysing single queries and does not suffer from the bias of filtering by keywords. First we divide the query log into user sessions, continuous queries from the same user followed by an inactivity period exceeding 30 minutes. After this, we attribute one of the following labels for each clicked URL, if any: (1) *Medical*, (2) *Not Medical*, or (3) *Not Found*. This depends on whether the URL is (1) found in any Medical category listed in Table 1; (2) found in any other category: News, Arts, Games, Health/Animals, Health/Beauty, etc; or (3) not found in either of these. Last, we assign to the whole session the Medical label only if the proportion of URLs on Medical information is greater than a threshold t . Medical search sessions classified this way are attributed to the set **AOL-Medical**, while the rest goes to the **AOL-NotMedical** set. Figure 3 illustrates the session assignment procedure. For the experiments performed in this work, we use $t = 0.5$. This value is a fair trade-off between two extremes: considering an entire session as being on medical information because one single URL on medical information was clicked (see second part of Figure 3), and considering an entire session as being on medical information only if all the known clicked links are on medical information (see the first part of Figure 3).

For the first part of Figure 3, it is important to note that the first query could be considered to belong to another session, as the user intent might be different from the rest of the session. The second and third queries, drug names that are clearly for medical content, were not used to calculate whether the session was on medical information or not, as their clicked URLs were not found in ODP. After the label estimation is done, all the queries of a session are assigned to the same class, therefore all six queries in S_i are assigned to AOL-Medical.

While only 27% of the queries have their URLs found in ODP, using the session approach described above allows us to have 50% of all sessions with at least one URL found in ODP. Altogether, 68% of all AOL queries were evaluated, as they belong to sessions that had at least one clicked URL in ODP. A more accurate way to define sessions is a field of research by itself [21, 30, 19] and it is not the goal of this work.

The **HON** dataset is composed of anonymous logs ranging from December 2011 to August 2013. This non-governmental organisation is responsible for the HON-code, a certification of quality given to websites fulfilling a pre-defined list of criteria [6]. HON provides a search engine to facilitate the access to the certified

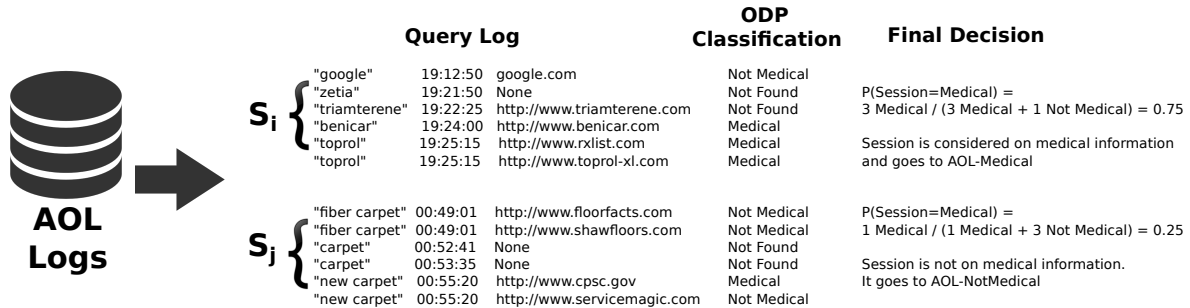
⁶ Obtained from <http://www.gregsadetsky.com/aol-data/>

⁷ <http://www.hon.ch/HONsearch/Patients/index.html>

⁸ <http://www.dmoz.org/>

Table 1: ODP categories used to filter the AOL-Medical. These categories are the most relevant ones related to Medicine in ODP hierarchy (see <http://www.dmoz.org/Health/Medicine/>)

| ODP Category | URL Examples |
|-------------------------------------|--|
| \Top\Health\Medicine | http://www.nlm.nih.gov http://www.webmd.gov |
| \Top\Health\Alternative | http://www.acupuncturetoday.com http://www.homeopathyhome.com |
| \Top\Health\Dentistry | http://www.dental--health.com http://www.animated-teeth.com |
| \Top\Health\Conditions_and_Diseases | http://www.cancer.gov http://www.cancer.org |
| \Top\Health\Organisations\Medicine | http://www.ama-assn.org http://www.aafp.org |
| \Top\Health\Resources | http://health.nih.gov http://www.eyeglassretailerreviews.com |

Fig. 3: Two real user sessions extracted from AOL logs, S_i is classified as a search for medical content, while S_j is not.

sites. Although the majority of the queries are issued in English, the use of French or Spanish is frequent. Aiming to reduce noise, every query in the HON dataset was re-issued in a commercial search engine and the snippets of the top 10 results were used as input for an automatic language detection tool [33], which presented a precision of 94% in filtering English queries.

As expert datasets, we use the logs from the Turning Research Into Practice (**TRIP**) database⁹ and ARRS **GoldMiner**¹⁰. The former is a search engine indexing more than 80,000 documents and covering 150 manually selected health resources such as MEDLINE and the Cochrane Library. Its intent is to allow easy access to online evidence-based material for physicians [34]. The logs contain queries of 279,280 anonymous users from January 2011 to August 2012. GoldMiner consists of logs from an *image* search engine that provides access to more than 300,000 radiology images based on text queries of text associated with the images. Although the usage of an image search engine is slightly different from document search, previous work in the liter-

ature [48,23] showed that the user search behaviour is similar. We had access to more than 200,000 queries with last logged query being issued in January 2012. Due to a confidentiality agreement, we cannot reveal the start date of this collection. The GoldMiner search engine is interesting because its users are so specialised and it therefore represents the particular case of catering to experts in a narrower domain inside medicine. As GoldMiner is so specialised, the number of laypeople using it is likely small. It is therefore a good example of the extreme specialisation end of the expert continuum, allowing the effects of this specialisation on the vocabulary and search behaviour of the users to be found.

3.3 Pre-processing Log Files

The first challenge dealing with different sources of logs is normalising them. Unfortunately, there is clickthrough URL information available only for the AOL and HON datasets, limiting a detailed click analysis. Therefore, we focus on a query content analysis, using only the intersection of all possible fields: (1) timestamp, (2)

⁹ <http://www.tripdatabase.com/>

¹⁰ <http://goldminer.rrs.org>

anonymous user identification, and (3) keywords. Neither stop word removal nor stemming were used.

Sessions were defined as follows. They begin with a query and continue with the subsequent queries from the same user until a period of inactivity of over 30 minutes is found. This approach for sessions, as well as the 30-minutes threshold, is widely used in the literature [8, 54, 30]. We excluded extremely prolific users (over 100 queries in a single session), since they could represent “bots” rather than individuals.

4 Enriching the Query Logs with MetaMap

The US National Library of Medicine MetaMap was intensively used in this work to enrich the information contained in the query logs, adding annotations regarding the concepts searched in the queries. MetaMap is widely used to map biomedical text to the Unified Medical Language System (UMLS) Metathesaurus, a compendium of many controlled vocabularies in the biomedical sciences [1]. This mapping can serve for different tasks, such as query expansion [4, 20], concept identification and indexing [2, 36], question answering [12], knowledge discovering [51], and more related to this work, enrich query logs to understand user goals [22, 25]. To explain how mapping queries to UMLS can give us some insights about the user intent, we first have to explain what UMLS is and how MetaMap maps text to UMLS. We explain how the mapping works in the next section and we evaluate the mapping in Section 4.2.

4.1 MetaMap

A Metathesaurus can be defined as a very large, multi-purpose, and multi-lingual vocabulary resource that contains information about biomedical and health related concepts, their various names, and the relationships among them [37]. In its 2013 version, the UMLS Metathesaurus has more than one hundred different controlled vocabulary sources and a large amount of internal links, such as alternative names and views of the same concept.

The white row of Table 2 is the original version of the classical UMLS example from [37]. It illustrates how different atoms can have the same meaning. Atoms are the basic building blocks from which the Metathesaurus is constructed, containing the concept names or strings from each of the source vocabularies. The atoms shown are part of two vocabularies PSY (Psychological Index Terms), and MSH (Medical Subject Headings, MeSH), mapping different strings and terms to the same concept, C0004238, which states that atrial fibrillation is a

pathological function. The other row of this table shows another NCI (National Cancer Institute), which states that atrial fibrillation can be an adverse event associated with the use of a medical treatment or procedure, although we do not know which medical treatment or procedure.

The job of MetaMap is to map a biomedical text to its corresponding concept(s). MetaMap generates a candidate set for a piece of text, based on its internal parser and variant generation algorithm, which takes into account acronyms, synonyms, inflections and spelling variants of the text. Then, based on metrics such as centrality, variation, coverage and cohesiveness, MetaMap ranks each candidate [1]. Occasionally, more than one candidate may have the same score. We collect all the top candidate(s) and its (their) associated semantic type(s), shown in bold below the CUIs in Table 2. In the running example, a text containing only ‘atrial fibrillation’ is mapped to both C0004238 and C1963067 with the same top score, and the types ‘Pathologic Function’ and ‘Finding’ are assigned to the query. To the best of our knowledge, MetaMap is the state of the art for mapping biomedical text to UMLS concepts.

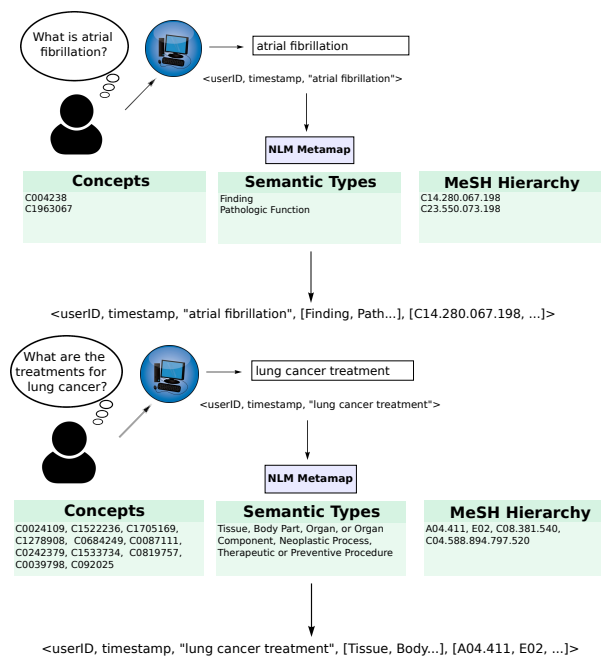


Fig. 4: Two different user queries are enriched with information extracted with MetaMap. In the top part, the same example used in Table 2 is processed by MetaMap. In the bottom part, the query “lung cancer treatment” is more ambiguous and results in different mappings, such as *Lung (Entire lung) / Cancer Treatment (Cancer Therapeutic Procedure)* and *Lung Cancer (Malignant neoplasm of lung) / Treatment (Therapeutic procedure)*

Table 2: A concept is potentially linked to various AUI (atom), SUI (string), and LUI (term). We used MetaMap to map a user query, e.g. “Atrial Fibrillation” to the different existing concepts (C0004238, C1963067). Note that each concept is associated to one single semantic meaning.

| Concept (CUI) | Terms (LUIs) | Strings (SUIs) | Atoms (AUIs) |
|--|---|--|--|
| C0004238 [Pathologic Function] Atrial Fibrillation (preferred) Atrial Fibrillations Auricular Fibrillation Auricular Fibrillations | L0004238 Atrial Fibrillation (preferred) Atrial Fibrillations | S0016668 Atrial Fibrillation (preferred) | A0027665 Atrial Fibrillation (from MSH) |
| | | S0016669 (plural variant) Atrial Fibrillations | A0027667 Atrial Fibrillation (from PSY) |
| | L0004327 (synonym) Auricular Fibrillation Auricular Fibrillations | S0016899 Auricular Fibrillation (preferred) | A0027930 Auricular Fibrillation (from PSY) |
| | | S0016900 (plural variant) Auricular Fibrillations | A0027932 Auricular Fibrillations (from MSH) |
| C1963067 [Finding] Atrial fibrillation (Atrial Fibrillation Adverse Event) | | | Auricular Fibrillations (from NCI) |

An interesting way to capture the user intent is mapping the queries to a well known domain corpus. In this work we use the Medical Subject Headings, MeSH, as it is a rich and well structured hierarchy that has already been studied to examine user query logs [22], allowing us to compare the behaviour of the users studied here with PubMed users. The whole MeSH hierarchy contains more than 25,000 subject headings in the 2013 version, the one used in this work, containing 16 top categories such as ‘Anatomy’ and ‘Diseases’. Figure 5

1. * Anatomy [A]
2. * Organisms [B]
3. - Diseases [C]
 - o Bacterial Infections and Mycoses [C01] +
 - o Virus Diseases [C02] +
 - o Parasitic Diseases [C03] +
 - o Neoplasms [C04] +
 - o Musculoskeletal Diseases [C05] +
 - o Digestive System Diseases [C06] +
 - o Stomatognathic Diseases [C07] +
 - o Respiratory Tract Diseases [C08] +
 - o Otorhinolaryngologic Diseases [C09] +
 - o Nervous System Diseases [C10] +
 - o Eye Diseases [C11] +
 - o Male Urogenital Diseases [C12] +
 - o Female Urogenital Diseases and Pregnancy Complications [C13] +
 - o Cardiovascular Diseases [C14] +
 - o Hemic and Lymphatic Diseases [C15] +
 - o Congenital, Hereditary, and Neonatal Diseases and Abnormalities [C16] +
 - o Skin and Connective Tissue Diseases [C17] +
 - o Nutritional and Metabolic Diseases [C18] +
 - o Endocrine System Diseases [C19] +
 - o Immune System Diseases [C20] +
 - o Disorders of Environmental Origin [C21] +
 - o Animal Diseases [C22] +
 - o Pathological Conditions, Signs and Symptoms [C23] +
 - o Occupational Diseases [C24] +
 - o Chemically-Induced Disorders [C25] +
 - o Wounds and Injuries [C26] +
4. * Chemicals and Drugs [D]
5. * Analytical, Diagnostic and Therapeutic Techniques and Equipment [E]
6. * Psychiatry and Psychology [F]
7. * Phenomena and Processes [G]
8. * Disciplines and Occupations [H]
9. * Anthropology, Education, Sociology and Social Phenomena [I]
10. * Technology, Industry, Agriculture [J]
11. * Humanities [K]
12. * Information Science [L]
13. * Named Groups [M]
14. * Health Care [N]
15. * Publication Characteristics [V]
16. * Geographicals [Z]

Fig. 5: MeSH hierarchy with the Disease branch expanded

shows an example of the MeSH hierarchy with the first level of the disease branch expanded.

We use the approach of Herskovic et al. [22] in this paper, mapping each query onto one or more MeSH terms with MetaMap. As shown in Figure 4, one query can be mapped to multiple MeSH identifiers. For example, the query ‘atrial fibrillation’ is mapped to both MeSH ids *C14.280.067.198* and *C23.550.073.198*, both in the topmost *Disease* category (represented by the starting letter ‘C’ as show in Figure 5). After the mapping to MeSH is done, we can easily have an overview regarding the subjects the users are more interested in. In this case we would conclude that this user is interested in diseases, as her/his only query maps only to category ‘C’, more specific in cardiovascular diseases, C14, and pathological conditions, C23.

After preprocessing, each query is converted into the following format: $\langle timestamp, userID, query, semanticTypes, meshIDs \rangle$, where the *timestamp*, *userID* and *query* are originally query log fields, while *meshIDs* and *semanticTypes* are the set of semantic types and MeSH identifiers generated by MetaMap. These two fields are examined in details in Sections 5.2.2 and 6.2. Figure 4 illustrates how the queries were enriched with the information provided by MetaMap and the final format.

Finally, it is important to mention that the queries were mapped to concepts in the UMLS 2013 AA US-Abase Strict Data and no special behaviour parameter was used. We manually examined the behaviour for two important parameters: allowing acronyms/abbreviations (-a) and using the word sense disambiguation mod-

ule (-y), and decided not to activate them. Our experiments show that activating the former parameter decreases the precision significantly for the sake of a small increase in recall, as MetaMap is already capable of matching some of the most frequently used abbreviations (HIV, HPV, AIDS, COPD). For the latter, we have an inverse scenario, where we had a small gain in precision but a larger loss in recall, as MetaMap always picks only one possibility when more than one concept is possible. It means that MetaMap would be forced to choose between concepts C0004238 and C1963067 of Table 2, even when both are equally likely. The last important reason for not using any other parameter is that we want to compare our results with [22], in which no special option was used either. For the experiments shown in Section 5.2.2 we used the parameter (-R) to restrict MetaMap to use only MeSH as vocabulary source.

4.2 Evaluation of the Mapping

As recently reported by MetaMap’s authors [3], a direct evaluation of MetaMap against a manually constructed gold standard mapping to UMLS concepts has almost never been performed. Usually, indirect evaluations are made, where the effectiveness of a task is measured with and without MetaMap. For example, query expansion using the related concepts of a concept identified by MetaMap versus not using it. Here we are interested in the few articles that evaluate the effectiveness of MetaMap, especially the ones focused on mapping user queries.

In 2003, Pratt and Yetisgen-Yildiz [42] compared MetaMap mappings to UMLS with mappings made by 6 physicians and nurses. For the 151 concepts in their ground truth, MetaMap could match 81 concepts exactly, 60 partially and could not match only 10 concepts, of which 6 were not available in UMLS. In a scenario considering partial matches (e.g., mapping to ‘*angiomas*’ instead of ‘*leptomeningeal angiomas*’), MetaMap had an F1 of 76%. In another experiment in the same year, Denny et al. [13] built a bigger gold standard dataset of 4,281 concepts to evaluate MetaMap, reaching a precision of 78%, recall of 85% and F1 of 81%.

More recently, Névéal et al. [36] reported results on using MetaMap to detect disease concepts on both literature and query corpus. The results showed that MetaMap had a better effectiveness for long sentences (F1 of 76%) than for short queries (F1 of 70%), but they also pointed out that the average inter-annotator agreement of the 3 assessors for the query corpus was 73%, showing that MetaMap results are not far from humans performing the same task. Using 1,000 queries

from partly the same datasets that are used here: AOL, HON and TRIP, Palotti et al. [39] also showed an F1 of 70% for query mappings.

Névéal et al. [35] created an annotated set of 10,000 queries that were mapped to 16 categories, in a similar way to what is done in Section 6.2, where the semantic types produced by MetaMap are used to define our own categories. We used Névéal’s dataset to calibrate our mappings for our ‘Cause’ and ‘Remedy’ categories (see Section 6.2), as well as to take decisions regarding MetaMap’s parameters. We used the ‘Disorder’ category of Névéal as an equivalent of our ‘Cause’ category, and we combined ‘*Chemical and Drugs*’ (antibiotic, drug or any chemical substance), ‘*Gene, Proteins and Molecular Sequences*’ (name of a molecular sequence) and ‘*Medical Procedures*’ (activity involving diagnosis, or treatment) as the closed possible class of our ‘Remedy’ class. We could reach an F1 of 78% for the ‘Cause’ category (P=75%, R=81%) and 72% for ‘Remedy’ (P=70%, R=73%). These figures are in line with what is known for MetaMap when mapping medical abstracts to concepts, encouraging us to use it for mapping short queries to concepts as well.

4.3 Using the Mappings

In the following sections, we show how we exploit the mappings made by MetaMap to enrich query logs. In Section 5, we use the mappings to analyse individual queries, following a very similar approach carried out by Herskovic et al. [22], being able to compare our results for individual queries. Later, in Section 6, the focus is on the session level. An interesting work which we took as a basis for comparison is Cartright et al. [8], which defines 3 classes: symptoms, diseases and remedies. Note that we could group the MeSH hierarchy into these three classes, but we prefer to use the semantic types provided by MetaMap, as it is more intuitive and it was already done in the literature (for example [26, 35, 38, 39]).

5 Individual Query Analysis

One goal of this section is to study how users search, based first on simple statistics to model their behaviour. Also, we start exploring the content of their queries, but considering all the queries without dividing into user sessions.

5.1 How Users Search

We start by showing a few simple but important statistics about the logs. The aim of this section is to un-

derstand the user behaviour through general statistics, as well as to show how each log is composed. In Table 3 we depict several metrics that are used to characterise user interactions, and compare their values to those in related studies. Torres et al. [14] use AOL logs to study queries performed by kids. White et al. [52] use a keyword-based method to filter domain specific queries and divide them into those issued by laypeople and those issued by experts. Their work also considers other types of queries, such as queries on computer science or financial information. We show only the data for the medical domain. Herskovic et al. [22] and Dogan et al. [25] analyse different periods of PubMed logs. For all datasets, “N/A” is used when the information is not available.

The query logs from the related work shown in Table 3 belong to the same time period as the AOL logs. Query logs from HON, TRIP and GM are considerably newer than the others. Nevertheless, Table 3 shows that AOL-Medical and HON are very similar in many aspects, such as the average number of terms per query and the average time per session. The biggest difference between these two query logs was found for the average number of queries per session, however the difference is small if compared to any other datasets.

The average terms/characters per query can be an indicator of the complexity and difficulty of the users to express their information needs. We note that AOL-Medical and HON queries are shorter than TRIP queries, and that TRIP logs are similar to PubMed logs in terms of query length. White’s work also found that expert queries are more complex than layperson queries.

The average number of queries per session and time per session, although considerably smaller than what was found by White’s work, follow the same pattern, with TRIP data having longer sessions than HON and AOL-Medical. We could not find an explanation for so long queries in White et al. dataset. We show only the sessions made by experts and laypeople in the medical domain from White’s work, but in their original paper they report that sessions are considerably smaller when the same set of people query in other domains: having a mean session length of less than 5 queries, and the mean time per session is never longer than 800 seconds.

We aggregate the log into 2 groups in Table 4: *laypeople* and *experts*, making the comparison of our datasets with the literature possible. As done by White et al. [52], we use Cohen’s d to determine the effect size of each variable between each pair of groups. We randomly sampled 45,000 users from TRIP and merged them with the 45,090 users from GoldMiner, making all datasets have a comparable number of users. Cohen’s d is a useful metric for meta-analysis [9] that uses the means and

standard deviations of each measurement to calculate how significant a difference is. Although there are controversies about what is a “small”, “medium” or “large” effect size, a recommended procedure is to define a Cohen’s d effect size of 0.2 or 0.3 as a “small” effect, around 0.5 as “medium” effect and greater than 0.8 as a “large” effect [9]. White et al. built a classifier to detect user expertise based on a superset of the features shown in Table 4. They argued that these are valuable features based on Cohen’s d value, as well as feature importance calculated by their regression classifier. Although considered to have a “small” effect, this was big enough to help separate experts from laypeople. We reached very similar Cohen’s d values to White’s paper, hypothesising that the behaviour could be used to predict expertise in other logs as well. In particular, we found the same ranking that White et al. found, among the four features presented in Table 4.

5.2 What Users Search for

In order to understand what users are looking for, we investigate popular terms and queries issued. Also, we use MetaMap to map queries to the MeSH hierarchy, finding the high level topics associated with the user queries.

5.2.1 Terms and Queries

We depict the most popular queries, terms (here excluding the stop words), and abbreviations used in all logs, as well as their frequency among the queries in Table 5. As expected, AOL-NotMedical contains navigational queries and diverse terms related to entertainment. Similarly, some of the most popular queries in AOL-Medical are navigational, with the website ‘*webmd.com*’ appearing twice in the top 10 queries, and the Mayo Clinic also a common query. Both of these navigational queries also appear in the HON search log. The analysis of AOL-Medical terms shows common medicine-related concepts, with people searching for information about different cancer types in more than 3% of the cases.

Most of the top queries in the TRIP log are related to disease. In TRIP logs, we found ‘*area:*’ in 3% of the queries, ‘*title:*’ in 2.2%, ‘*to:*’ in 1.5% and ‘*from:*’ in 1.8%, in total these keywords were used in 6.7% of the queries, however, we do not show these terms in Table 5, as they do not reveal what the users search, but how they search. These patterns were not found in the other datasets. The use of more advanced terms is also found in PubMed logs [22], we hypothesise that some users might just copy and paste their queries from PubMed into the TRIP search engine, resulting in queries such

Table 3: General Statistics describing the query logs. ‘*’ means that the median used as the mean was not informed and ‘N/A’ means that the data was not available.

| Dataset | This Work | | | | | Literature | | | | | |
|--------------------------|----------------------|-----------------------|----------------------|-----------------------|----------------------|--------------------|-------------------|----------------------|------------------|-----------|------------|
| | Laypeople | | Experts | | Non-Medical | AOL-Kids | AOL-NKids | Laypeople | Experts | Pubmed | |
| | AOL-M | HON | TRIP | GM | AOL-NM | Torres et al. [14] | White et al. [52] | Herskovic et al.[22] | Dogan et al.[25] | | |
| Logs Initial Date | Mar 2006 | Dec 2011 | Jan 2011 | N/A | Mar 2006 | Mar 2006 | May 2007 | Jan 2006 | Mar 2008 | | |
| Logs Final Date | May 2006 | Aug 2013 | Aug 2012 | Jan 2012 | May 2006 | May 2006 | Jul 2007 | Jan 2006 | Mar 2008 | | |
| # Users | 47,532 | 47,280 | 279,280 | 45,090 | 655,292 | N/A | N/A | 37,243 | 7,971 | 624,514 | NA |
| # Queries | 215,691 | 140,109 | 1,788,968 | 219,407 | 34,427,132 | 485,561 | N/A | 673,882 | 362,283 | 2,689,166 | 58,026,098 |
| # Unique Queries | 69,407 | 85,824 | 486,431 | 90,766 | 9,695,882 | 10,252 | N/A | N/A | N/A | N/A | N/A |
| # Sessions | 79,711 | 77,977 | 344,038 | 100,848 | 10,555,562 | 21,009 | N/A | 68,036 | 26,000 | 740,215 | 23,017,461 |
| Avg Terms Per Query | 2.61 (± 1.71) | 2.72 (± 2.05) | 3.40 (± 2.33) | 2.28 (± 2.54) | 2.46 (± 1.87) | 3.23 | 2.5 | 2.92 | 3.30 | 3* | 3.54 |
| Avg Char Per Query | 16.22 (± 9.11) | 18.11 (± 11.48) | 21.22 (± 9.69) | 16.64 (± 10.20) | 15.98 (± 9.67) | N/A | N/A | 20.76 | 24.05 | N/A | N/A |
| Avg Queries Per Session | 2.71 (± 2.50) | 1.80 (± 2.48) | 5.20 (± 5.95) | 2.18 (± 2.57) | 3.26 (± 4.65) | 8.76 | 2.8 | 9.90 | 13.93 | N/A | 4.05 |
| Avg Time Per Session (s) | 258 (± 531) | 208 (± 592) | 471 (± 758) | 163 (± 520) | 384 (± 809) | 1238 | N/A | 1549.74 | 1776.45 | N/A | N/A |

Table 4: General Statistics – Stratified by expertise. L for laypeople and E for experts

| Dataset | Laypeople | Experts | Cohen’s d | |
|--------------------------------|-----------------------|-----------------------|-------------|--------------------|
| Total Number of Users | 94,812 | 90,090 | E - L | E - L from [52] |
| Total Number Of Queries | 355,800 | 504,745 | | |
| Total Number Of Unique Queries | 149,648 | 181,051 | | |
| Total Number Of Sessions | 157,688 | 155,965 | | |
| Mean Terms Per Query | 2.65 (± 1.85) | 2.91 (± 2.09) | 0.13 | 0.20 |
| Mean Chars Per Query | 16.97 (± 10.16) | 19.18 (± 10.16) | 0.22 | 0.30 |
| Mean Queries Per Session | 2.26 (± 2.53) | 3.24 (± 4.29) | 0.28 | 0.38 |
| Mean Time Per Session (sec) | 233 (± 562) | 271 (± 629) | 0.06 | 0.11 |

as ‘*palliative care (area:oncology)*’, indicating that the user wants material about palliative care specifically for the area of oncology. ‘Title’ is used in PubMed for performing a search only in the title of the indexed articles, while ‘from:’ and ‘to:’ specify periods of time in which a document was published.

The topmost query in the HON log and its top 3 terms are ‘*trustworthy health sites*’. It shows that many of the queries are from users that do not know which are the medical websites that they can trust, and also demonstrates a misunderstanding by the end users of the nature of the content indexed by the HON search engine (only HONcode-certified websites are indexed).

For the GoldMiner queries and terms, we clearly see the increase in the terminological specificity of the most popular keywords used.

5.2.2 Mapping to MeSH

MeSH is a hierarchical vocabulary used by US National Library of Medicine for indexing journal articles in the life sciences field. A query log analysis using MeSH was also carried out by Herskovic et al.[22] for the PubMed logs in order to understand what are the most popular topics searched by the users. We use the same weighting schema used in Herskovic’s work: if n categories are detected in one query, we give the weight of $1/n$ to these categories.

General statistics calculated for the mapping of user queries to MeSH terms are shown in Table 6. Here, we

are testing MetaMap for the annotation of non-medical queries as well, which to the best of our knowledge was never studied.

An interesting result is the fact that around 50% of AOL-NotMedical queries were successfully mapped to a MeSH concept. To investigate this, we collected a large random sample of mapped queries and analysed them. We found that MetaMap is able to find many concepts not directly linked to medicine, such as geographic locations, animals and plants, food and objects. For example, ‘www’ (*L01.224.230.110.500*), used in 10% of all AOL queries, is recognised and annotated as *Manufactured Object*. Also, locations are usually very commonly found and help to explain the high mean MeSH depth found for this dataset, second row in column AOL-NM (California is mapped to both *Z01.107.567.875.760.200* and *Z01.107.567.875.580.200*). It is important to have this in mind when building systems like in Yan et al. [56], in which the MeSH depth is used to model document scope and cohesion. When looking at false positive mappings, especially the ones mapping to diseases and symptoms, we detected that MetaMap’s errors fall into two main categories: (1) English common words: tattoo (tattoo disorder), pokemon (ZBTB7A gene), and (2) abbreviations: park (Parkinson disease), dvd (Dissociated Vertical Deviation). For both types of errors, MetaMap or a system using it, would have to use the context (words around the mapping) to detect that Pokemon is used as a cartoon or

Table 5: Top queries and terms and their relative frequency (%) among all queries

| Rank. | Laypeople | | | | Experts | | | | | |
|----------------|--------------|------|--------------------------|------|---------------|------|---------------------|------|----------------|------|
| | AOL-Medical | | HON | | TRIP | | GoldMiner | | AOL-NotMedical | |
| | String | Freq | String | Freq | String | Freq | String | Freq | String | Freq |
| QUERIES | | | | | | | | | | |
| 1 | webmd | 0.98 | trustworthy health sites | 4.24 | skin | 0.29 | mega cisterna magna | 0.44 | google | 0.95 |
| 2 | web md | 0.41 | cancer | 0.51 | diabetes | 0.22 | baastrup disease | 0.40 | ebay | 0.40 |
| 3 | shingles | 0.27 | webmd | 0.47 | asthma | 0.17 | toxic | 0.23 | yahoo | 0.37 |
| 4 | mayo clinic | 0.26 | sleep apnea syndromes | 0.27 | hypertension | 0.14 | limbus vertebra | 0.22 | yahoo.com | 0.28 |
| 5 | lupus | 0.25 | lymphoma | 0.22 | stroke | 0.13 | cystitis cystica | 0.20 | mapquest | 0.25 |
| 6 | herpes | 0.20 | breast cancer | 0.21 | osteoporosis | 0.11 | thornwaldt cyst | 0.14 | google.com | 0.23 |
| 7 | diabetes | 0.19 | hypertension | 0.18 | low back pain | 0.10 | buford complex | 0.13 | myspace.com | 0.22 |
| 8 | fibromyalgia | 0.18 | mayoclinic.com | 0.16 | copd | 0.10 | splenic hemangioma | 0.13 | myspace | 0.21 |
| 9 | pregnancy | 0.16 | obesity | 0.16 | breast cancer | 0.09 | throckmorton sign | 0.12 | www.yahoo.com | 0.12 |
| 10 | hernia | 0.16 | drweil.com | 0.14 | pneumonia | 0.09 | double duct sign | 0.12 | www.google.com | 0.12 |
| TERMS | | | | | | | | | | |
| 1 | cancer | 3.40 | health | 6.39 | treatment | 3.03 | cyst | 3.17 | free | 1.24 |
| 2 | hospital | 3.00 | sites | 4.37 | cancer | 2.56 | mri | 1.89 | google | 1.04 |
| 3 | pain | 2.25 | trustworthy | 4.28 | pain | 2.13 | disease | 1.80 | county | 0.65 |
| 4 | symptoms | 2.14 | cancer | 2.74 | care | 2.10 | ct | 1.75 | yahoo | 0.62 |
| 5 | disease | 2.03 | disease | 1.53 | children | 1.98 | fracture | 1.68 | pictures | 0.60 |
| 6 | blood | 1.87 | diabetes | 1.17 | therapy | 1.81 | tumor | 1.65 | lyrics | 0.52 |
| 7 | medical | 1.62 | treatment | 0.96 | diabetes | 1.80 | syndrome | 1.47 | school | 0.51 |
| 8 | webmd | 1.21 | syndrome | 0.87 | disease | 1.78 | liver | 1.26 | myspace | 0.49 |
| 9 | surgery | 1.14 | heart | 0.83 | pregnancy | 1.70 | pulmonary | 1.22 | ebay | 0.46 |
| 10 | syndrome | 1.13 | pain | 0.80 | acute | 1.41 | bone | 1.16 | sex | 0.44 |
| 11 | breast | 1.11 | care | 0.77 | syndrome | 1.39 | renal | 1.13 | florida | 0.45 |
| 12 | center | 1.09 | effects | 0.75 | management | 1.14 | sign | 1.12 | sale | 0.41 |
| 13 | health | 1.04 | medical | 0.67 | stroke | 1.07 | lung | 1.11 | city | 0.40 |
| 14 | heart | 0.90 | blood | 0.65 | surgery | 1.06 | brain | 1.08 | home | 0.39 |
| 15 | diabetes | 0.86 | pregnancy | 0.61 | prevention | 1.05 | cell | 1.00 | state | 0.39 |

a game, and not as a gene name. Specifically for the second case, it would be desirable if MetaMap could allow the use of a pre-defined list of acronyms to increase its precision. In the current implementation, MetaMap has a parameter for user defined acronyms (-UDA), but it is just used to expand more acronyms instead of overwriting its pre-defined ones. Also for AOL-NL, the third and fourth rows indicate the suitability of using mappings to MeSH for distinguishing between medical and non-medical queries. Queries from the medical logs have a larger number of MeSH terms and disease terms than AOL-NM. If the errors analysed above could be amended using the query context or session, for example, then a mapping to MeSH could be helpful to detect queries or sessions on medical information.

Going further, we present in Figure 6 the most popular categories for the first level of the MeSH hierarchy. We also show the results obtained by Herskovic et al. [22] for PubMed, in order to compare our findings. We show only the categories that have more than 5% of the queries containing MeSH terms mapped to it.

When Herskovic and colleagues did this experiment, they found that PubMed users were more interested in the category *Chemical and Drugs*. In general, the distributions over the categories for the AOL-Medical,

Table 6: General MeSH Statistics

| Metric | Laypeople | | Experts | | AOL-NM |
|------------------------------------|-----------|-------|---------|-------|--------|
| | AOL-M | HON | TRIP | GM | |
| % of queries containing MeSH terms | 77.87 | 77.81 | 85.96 | 79.02 | 50.51 |
| Mean MeSH Depth | 3.99 | 3.83 | 3.86 | 4.01 | 4.37 |
| Mean MeSH terms per query | 2.14 | 2.19 | 2.78 | 2.07 | 1.12 |
| Mean Disease terms per query | 0.81 | 0.60 | 0.99 | 1.17 | 0.05 |

HON and TRIP search logs are similar. However, differently from PubMed, we found that the users are generally most interested in *Diseases*, and then *Chemicals and Drugs*. The results for GoldMiner show another trend for the second most popular category, focused on anatomy rather than on drugs, likely because radiologists often have to append to their query the part of the body that they are interested in. In its actual version, GoldMiner has a filter for age, sex and modality (e.g CT, X-ray), but it has no filter for body parts. This analysis suggests that it could be interesting to add a filter for body regions as well.




Fig. 6: Popular categories according to MeSH mappings

Last, the four classes to the right of Figure 6 partly explain the high percentage of AOL-NotMedical terms mapped to MeSH terms. Also, the high percentage of these least medical categories, together with low percentage of relevant medical categories, the four classes to the left of Figure 6, can be used as a discriminative feature to distinguish between medical and non-medical logs.

6 Analysing Sessions

From now on, we consider user sessions instead of separate queries. Once more, we study first the user behaviour, then the content of each session.

6.1 Session Characteristics

A series of queries, part of an information seeking activity, is defined as a session. We consider that, after issuing the first query, a user may act in four different ways: (1) repeat exactly the same query, (2) repeat the query adding one or more terms to increase precision, (3) reduce the number of terms to increase recall, or (4) reformulate the query changing some or all the terms used. We ignore the first case because we cannot be sure if a user is really repeating the same query or just changing the result page, as some search engines record the same query as a result of a page change.

Table 7 depicts the changes made by users during the sessions. If during one single session a user adds

a term to the previous query and then changes a few words, we count one action in the row Exp.Ref (for expansion and reformulation – the order is not important). At the end, we divide the number of actions of each row by the total actions in the query log. Hence, Table 7 shows that the most frequent user action is the reformulation alone but it is more likely to happen in search engines targeting laypeople, e.g., 84% of the sessions in the AOL-Health logs and 63% of HON had only reformulations. The last row of Table 7 shows that expert users might be more persistent than laypeople, as more than 10% of the sessions in the professional search engines are composed of every type of action, while in laypeople logs this number is less than a third of this. In the literature, White et al [52] also hypothesise that expert users are more persistent than laypeople.

To better understand the last row of Table 7, we plot the two graphs in Figure 7. The first graph is the cumulative distribution of session length, showing that TRIP has clearly longer sessions, with 20% of the sessions being longer than 5 queries. The second graph shows the last row of Table 7 distributed over different session length. In this graph we can see how TRIP and GoldMiner users tend to perform more actions even for short sessions, as 20% of sessions of length 4 have already done all 3 actions. We also studied the user behaviour when the query repetition is allowed and we found a very similar situation.

Table 7: Aggregated percentages for query modifications along the sessions

| Action | Laypeople | | Experts | | AOL-NM |
|----------------|-----------|-------|---------|-------|--------|
| | AOL-M | HON | TRIP | GM | |
| Expansion | 6.66 | 13.83 | 14.85 | 5.96 | 3.71 |
| Reduction | 1.23 | 2.23 | 4.35 | 9.61 | 0.84 |
| Reformulation | 84.74 | 63.56 | 43.96 | 49.56 | 80.27 |
| Exp. and Red. | 0.37 | 1.29 | 5.09 | 3.54 | 0.57 |
| Exp. and Ref. | 5.43 | 13.90 | 15.27 | 8.28 | 9.66 |
| Red. and Ref. | 1.01 | 2.21 | 5.63 | 12.01 | 2.09 |
| Exp. Red. Ref. | 0.56 | 2.98 | 10.85 | 11.04 | 2.86 |




Fig. 7: The top graph shows the cumulative distribution of sessions length in terms of number of queries. The users of HON and GoldMiner tend to have shorter sessions, while the users of TRIP have longer sessions. The graph below shows the percentage of users that in a single session perform all three actions (expand, reduce and reformulate the previous query) for sessions of different sizes. As we can expect, this percentage increases as the query length increases but it is much higher for the expert users.

6.2 What are the Sessions About?

In this section, we attribute meaning to the users' queries in order to better understand their behaviour in a medical search context. We decide to use the same classes defined in Cartright, White and Horvitz [8]: symptom, cause and remedy, so that a direct comparison can be performed. A difference of their method and ours is that we classify the queries into the semantic types using MetaMap, as done in [26,35,38,39], instead of hand-made rules.

In Figure 8, we show all concepts that have a frequency of at least 5% in any query log. Additionally, we show the type '*Sign and Symptom*' because it is an important concept in our further analysis. We show only these 10 semantic types for a matter of readability, as currently MetaMap recognises 133 semantic types and it is not possible to visualise them all¹¹. The single most common type in all the medical logs is '*Disease and Syndrome*'. As we expect, the top types in AOL-NotMedical are not really related to the medical domain, and the second most common semantic type for GoldMiner is related to parts of the body, as one might expect for radiology queries.

After a meticulous analysis of the semantic meaning assigned for the queries and the experiments described in Section 4.2, we defined the following classification based on the three classes created in [8] (some examples of queries classified for each type are given for a better understanding):

- **Symptom:** Sign or Symptom (*cough; sore; headache; red eyes*), Findings (*stress; testicular cyst*)
- **Cause:** Anatomical Abnormality (*hiatal hernia*), Cell or Molecular Dysfunction (*macrocytosis*), Congenital Abnormality (*scoliosis*), Disease or Syndrome (*diabetes; heart failure*), Experimental Model of Disease (*cancer model*), Injury or Poisoning (*achilles tendon rupture*), Mental or Behavioural Dysfunction (*bipolar disorder*), Neoplastic Process (*lung cancer; tumor*), Pathologic Function (*atypical hyperplasia*)
- **Remedy:** all 28 types belonging to the high-level group Chemicals & Drugs, which includes Clinical Drug (*cough syrup*), Antibiotic (*penicillin*), Pharmacologic Substance (*tylenol; mietamizol*), Amino Acid, Peptide, or Protein (*vectibix; degarilex*), Immunologic Factor (*vaccine; acc antibody*), Vitamin (*quercetin, vitamin B12*), Therapeutic or Preventive Procedure (*treatment; physiotherapy*), etc.

We analyse the most popular semantic types found in the queries and show them in Table 8, together with a direct comparison to Cartright et al. [8]. The largest difference between all four medical logs analysed in this paper and the Cartright et al. results is in the symptom category. For the latter, 63.8% of the sessions are focused on symptoms, while between 5.5% and 9.1% are focused on symptoms in our analysis. The main reason for Cartright's result is linked to the way in which they created their dataset: *keeping only sessions that had at least one query containing a term in a wordlist extracted from a list of symptoms from the Merck*

¹¹ A complete list of all semantic types can be found online: <http://metamap.nlm.nih.gov/SemanticTypesAndGroups.shtml>

medical dictionary. Their preprocessing step therefore explains the fact that most of the sessions were concentrated only on searching for symptoms. Conversely, our analysis reveals that the most common user focus is on causes rather than on symptoms. Also, the second most common focus is on a way to cure a disease. It is important to note that Cartright et al. logs date from 2009, it means they are 3 years younger than AOL, but roughly 3 years older than HON, also suggesting that large divergence found is due to the preprocessing steps and not to an evolution on how the users search.

Once more, GoldMiner presents a different behaviour, we hypothesise that the low number of sessions on remedies is explained by the fact that radiologists are not interested in remedies when searching for images as they are rather in the diagnosis phase. It is interesting to note that searching for causes and remedies in the same session is a very frequent task for medical professionals in the TRIP logs, with 16% of the sessions searching for both remedies and causes.

In Table 9, we show the behaviour modifications along a session. One oscillation is characterised by a transition from one focus type to another and then back to the original type. Originally, this study was made to support the hypothetico-deductive searching process in which a user cyclically searches for a symptom, then a cause and then returns to symptom [8]. The symptom-cause pattern was also found in our experiments, but with a more balanced distribution in relation to the other patterns. Again, the large number of behaviours involving symptoms found in [8] is likely an artefact of how the dataset was constructed. We see that the cause-remedy pattern plays a very important role, especially in the TRIP log, in which this is the most common pattern. Finally, the least frequent pattern found in all four datasets is the symptom-remedy one. The study of the behaviour modification was used in [8] to build a classifier to predict what is the next user action, allowing a search system to support medical searchers by pre-fetching results of possible interest or suggesting useful search strategies.

7 User Classification

We have seen in Sections 5 and 6 that experts and laypeople use different search strategies. In this section, we take advantage of these differences to build an automatic classifier that can assist search systems, exploring the user domain knowledge.

The expertise inference can be directly applied by a search engine to tailor the results shown, e.g. boosting easy to read documents for laypeople [49,40,11], or search aids, such as query suggestions to match the

searcher expertise. Also, the search strategies employed by experts could be used to support non-experts in learning more about domain resources and vocabulary [52].

In order to take advantage of the user domain expertise, it is necessary to be able to identify whether a user is an expert or not. We employed a Random Forest classifier¹² to solve this binary classification problem, since it is a well-known machine learning technique and has shown to be suitable for this task before [10,38].

The classifier relies upon a set of features to take its decision between the two modelled classes: expert or layperson. We list 14 features proposed in this work in Table 10, and group them into two sets: (1) user behaviour features, and (2) medicine-related features. The first set is made from the analysis of Sections 5.1 and 6.1, while the second one covers Sections 5.2 and 6.2.

Table 10: Features used in the expertise classification task. Two groups were created using the features discussed in the previous sections of this work.

| Group | Feature | Explanation |
|--------------------------|---|--|
| User Behaviour Features | AvgCharPerQuery AvgTermsPerQuery | Average number of characters and terms used by the user in each query |
| | AvgQueryPerSession AvgTimePerSession | The average number of queries and time per session. |
| | AvgExpansions AvgReductions AvgReformulation | Compares the <i>i</i> -th query to (<i>i</i> -1)-th query and counts the expansions, reductions and reformulations made |
| | AvgSymptomsPerQuery AvgCausesPerQuery AvgRemediesPerQuery AvgOtherTypePerQuery | The average number of symptoms/causes/remedies/none of them per query |
| | PercQueriesWithMeSH | Percentage of queries that could be mapped to any MeSH concept |
| Medical Related Features | AvgMeSHPerQuery | Average number of MeSH concepts identified in all queries |
| | AvgMeSHDepth | The average depth of all identified concepts |

To form our dataset, we merged the users from AOL-Medical and HON logs into the laypeople class, and the users from TRIP and GoldMiner logs into the expert class. As noted in Section 5.1, the number of users from TRIP logs is considerably larger than the other logs, therefore we repeat the sampling made to generate Table 4, and 94,812 users are created for the layperson class and 90,090 for the expert one.

We performed a ten-fold cross-validation experiment and present the results in Table 11. We employed as the baseline a simple classifier that always outputs the positive class, which could reach an F_1 of 67.8%. The

¹² The Random Forest classifier is based on the python machine learning module scikit-learn (<http://scikit-learn.org/>). Hyper-parameters were optimised using a grid-search approach.




Fig. 8: The top most frequently used semantic types (frequency in percentage). Many of the most used types are aggregated to study the user focus described in Table 8

Table 8: User focus when searching for medical content in a single session.

| Intent | Laypeople | | Experts | | AOL-NotMedical | Cartright et al.[8] |
|--------------------|-------------|------|---------|-----------|----------------|---------------------|
| | AOL-Medical | HON | TRIP | GoldMiner | | |
| None | 34.0 | 40.4 | 16.8 | 21.2 | 82.9 | 3.9 |
| Symptom | 9.1 | 6.3 | 5.5 | 6.4 | 3.9 | 63.8 |
| Cause | 24.3 | 20.9 | 26.0 | 58.2 | 3.3 | 5.3 |
| Remedy | 14.7 | 16.2 | 17.4 | 3.3 | 7.5 | 1.1 |
| Symptom and Cause | 6.8 | 6.1 | 7.2 | 6.4 | 0.5 | 22.6 |
| Symptom and Remedy | 2.1 | 2.6 | 4.5 | 0.5 | 0.9 | 2.0 |
| Cause and Remedy | 7.1 | 5.0 | 15.9 | 3.0 | 0.8 | 0.4 |
| All three | 1.9 | 2.5 | 6.7 | 1.0 | 0.2 | 0.8 |

next two rows of Table 11 show the classification performance when using only a single group of features. Clearly for our experiments, the user behaviour features were more important than the medical ones: while the medical features marginally improved the F_1 score for each class, the user behaviour features could reach an improvement of 14% over the baseline for detecting experts. The last row of Table 11 shows the performance of the classifier using all features. We highlight an improvement of more than 20% over the baseline for both classes.

The Random Forest classifier also allows us to compute the Gini importance score for each feature. This value (from 0.0 to 1.0) is higher when the feature is more important, indicating how often a particular feature was selected for a split in a random forest, and how large its overall discriminative value was for the classification problem under study. We show in Figure 9 all

the features according to the Gini importance score. Befitting the results of Table 11, the most important features were predominantly user behaviour features.

Table 11: Classification results: compared to the baseline, the Random Forest classifier using all the features can reach an improvement of 26% when detecting experts.

| Classifier | Pos. Class | Acc. | Prec. | Rec. | F1 |
|---------------------------------------|------------|------|-------|-------|------|
| Baseline | Layp. | 51.3 | 51.3 | 100.0 | 67.8 |
| | Exp. | 48.7 | 48.7 | 100.0 | 65.5 |
| Random Forest User Behaviour Feat. | Layp. | 75.7 | 76.3 | 76.4 | 76.3 |
| | Exp. | 75.7 | 75.1 | 75.0 | 75.0 |
| Random Forest Medical Features | Layp. | 67.1 | 67.4 | 69.5 | 68.5 |
| | Exp. | 67.1 | 66.8 | 64.6 | 65.7 |
| Random Forest All Features | Layp. | 83.5 | 84.1 | 83.6 | 83.9 |
| | Exp. | 83.5 | 82.8 | 83.4 | 83.1 |

Table 9: Cycle Sequence along a single session

| Pattern | Interaction | Laypeople | | Experts | | Cartright et al.[8] |
|----------------|------------------------|--------------------------------|------|---------|-----------|---------------------|
| | | AOL-Medical | HON | TRIP | GoldMiner | |
| | | Sessions with oscillations (%) | | | | |
| Symptom-Cause | Symptom→Cause→Symptom | 19.2 | 15.6 | 13.2 | 22.7 | 51.4 |
| | Cause→Symptom→Cause | 19.9 | 18.8 | 14.5 | 35.3 | 38.4 |
| Symptom-Remedy | Symptom→Remedy→Symptom | 8.2 | 11.8 | 10.8 | 4.1 | 5.1 |
| | Remedy→Symptom→Remedy | 8.1 | 14.2 | 11.6 | 3.8 | 2.7 |
| Cause-Remedy | Cause→Remedy→Cause | 18.2 | 18.4 | 24.8 | 20.3 | 1.5 |
| | Remedy→Cause→Remedy | 26.4 | 21.2 | 25.1 | 13.8 | 0.9 |




Fig. 9: Feature importance according to the Gini importance score generated by the Random Forest classifier. The error bars represent the standard deviation from the mean value for each feature.

8 Discussion

We presented the analysis of four different query logs divided into five datasets. We discuss each of the initial research questions in the next subsections, with the coverage of the third research question (relation to previously published results), covered in each subsection.

8.1 MetaMap and Short Queries

This study relies on the accuracy of MetaMap to enable the intent of the searchers to be identified. As MetaMap was designed for annotation of documents and not queries, we did an evaluation of its performance for short queries. Using an existing dataset of 10,000 manually annotated queries [35], we evaluated MetaMap on two of the categories used in this paper:

cause and remedy. The category symptom was not evaluated as it is not included in the dataset used. It is found that MetaMap can annotate the cause category with an F1 of 78% and the Remedy category with an F1 of 72%. While these values are not directly comparable to other results published, they correspond to the level of accuracy measured for related tasks: MetaMap was shown to map disease concepts in queries with an F1 of 70% [36], and a mapping into five classes in [39] on 1000 queries was done with an F1 of 70%. Most importantly, inter-annotator agreement for the manual annotation of the query corpus in [36] was 73%. This demonstrates that the results obtained by annotating the queries by MetaMap are at the same level as those obtained by manual annotation, implying that the MetaMap annotations are sufficiently accurate for this study.

8.2 How is Search Conducted for Medical Content?

This section covers the behaviour of the users when searching for medical information. Analyses were done both at the level of individual queries and of sessions. It was found that the mean terms per query and mean chars per query were higher for experts than for laypeople with a small effect (measured by Cohen's d value). This supports the small effect also detected for these characteristics by White et al. [52]. Moving toward sessions, we found also longer sessions in both terms of mean queries per session and mean time spent per session, with a small effect, as detected in White et al. Although White et al. studied search logs from a general purpose commercial search engine, for which assumptions had to be made about the behaviour of users in order to detect experts and laypeople, we were able to find very similar effect size, including the same importance ranking for the four characteristics measured. These small effects were sufficient to successfully train a classifier to predict expert and layperson classes in [52] and also in this paper.

When analysing the user behaviour in terms of sessions, we conclude that experts are more persistent than laypeople, as more than 10% of the sessions in the professional search engines were composed of all possible query modification actions (expansion, reduction, reformulation). This was also found in White's work, where they noted that sessions conducted by domain experts were generally longer than non-expert sessions and that domain experts consistently visited more pages in a session. Alternatively, longer sessions could mean that experts are struggling to find relevant information. It supports the current efforts of the information retrieval community to help experts finding scientific material to improve their clinical decisions [43]. It would be interesting to study if the increase of expertise of laypeople can change their user behaviour over time as suggested by Wildemuth [55], but this will likely require years of search engine logs.

8.3 What are the Users Searching for?

The investigation of what the users search for led us to conclusions that are significantly different from results published in the literature. In both of our analyses, the one based on the MeSH hierarchy and the one based on semantic types, we observed that users are more concerned with diseases rather than symptoms, converse to what Cartright et al. [8] found. This difference is large: Cartright et al. found that searches for symptoms occur in 63.8% of the sessions, while our results showed that symptoms were only in 5.5% to 9.1% of sessions, depending on the search engine. In our analysis, the cause category appeared most often, in 20.9% to 58.2% of the sessions, depending on the search engine. This large difference is likely due to the fact that Cartright et al. had to make assumptions about the characteristics of a medical query in order to extract medical queries from the logs of a general purpose search engine, whereas we used search logs from domain-specific medical search engines in three of the four cases. This allowed us to make the very strong assumption that users will always enter medical queries into these search engines. Understanding what users are searching for is an essential step towards providing more relevant search results.

We also identified patterns supporting the hypothetico-deductive searching processes, especially for the cause-remedy component, in which both laypeople and experts cycle through searching for causes and remedies in sessions so as to discover potential treatments for a disease. Finally, we found that TRIP users, mainly users falling into our expert class, use the hypothetico-deductive method very often, in more than 60% of their

sessions, versus less than 25% for AOL and HON. This supports the hypothesis that experts have much more complex information needs, which are not well addressed by the current search systems [43].

An interesting kind of search in the medical domain is the one for self-diagnosis purposes [17], which often arises before consulting a medical professional (or to help the decision to consult). Previous research has shown that exposing people with no or scarce medical knowledge to complex medical language may lead to erroneous self-diagnosis and self-treatment and that access to medical information on the Web can lead to the escalation of concerns about common symptoms (e.g., cyberchondria) [53]. Also current commercial search engines are far from being effective in answering such queries [60], presenting on average only 3 highly relevant results in the top 10 results. In the same manner that experts can assist non-experts in detecting credible content on the Web [44], a search system capable of inferring user expertise can learn about the decisions taken by experts to better support non-experts. In the case of self-diagnosis, the symptom-cause cycle in the expert search logs can be explored to provide query suggestions for non-experts.

After consulting a medical professional, non-experts often query about a disease or about a treatment that was recommended to them [17]. When literally copying-and-pasting the complex terms into a search box, they are presented with documents that are potentially as complex as their queries [20]. Inferring user medical knowledge can help matching non-experts with the suitable documents for them even for complex queries, significantly diminishing harmful situations and misunderstandings.

8.4 What are the Most Useful Features to Infer User Expertise?

We grouped the features collected in this work into two distinct sets: user behaviour features and medicine-related features. Judging by the experimental results, the user behaviour features are indispensable; and, while the medicine-related features alone were not very effective, they showed to provide large gains in all metrics, when combined with the user behaviour ones.

When analysing the features through the Gini importance coefficient, the average MeSH depth was considered the best medical feature by the classifier, a feature that was also highly ranked in [38]. For the user behaviour features, the main four metrics analysed here were also important in [52], while features based on query modification in a session seemed not to be well used by the classifier.

9 Conclusion

In this paper, we conducted a detailed study of medical information search behaviour through query logs. We studied how users search for medical documents, as well as what they search for. Results were compared to those in published studies analysing search logs in the medical domain. Almost all recent studies about the behaviour of searchers looking for medical information have been based on the search logs of a large commercial general purpose search engine. This paper performs the important task of reproducing these studies as far as possible on search logs from other search engines to find out to what extent these results can be supported or not. An important difference with this study compared to published studies is the use, in three of the four cases, of domain-specific medical search engines targeted at either experts or laypeople, meaning that we have very strong priors about who is using the search engines and what they are searching for. This avoids assumptions that have to be made in order to extract medical queries or extract expert or laypeople queries from the search log of a general purpose search engine.

Our results support those published in the literature for the following outcomes: (1) It is possible to distinguish between medical experts and laypeople based on search behaviour characteristics; (2) experts issue more queries and modify their queries more often, meaning that they can be either more persistent than laypeople or that their information need is more complex and more difficult to reach.

A large difference with respect to what is published in the literature was found for what the users are searching for. Our analysis showed that diseases were the focus of the largest number of sessions (20.9%–58.2%), as opposed to symptoms (63.8% in [8]). We suggest that this difference is mainly due to the criteria used to extract medical queries from the search logs of a general purpose search engine, which skewed the results toward symptoms. This result suggests that the occurrence of Cyberchondria [53] is less prevalent, especially on domain-specific medical search engines. A further result from this study that is potentially useful for search systems is the study of features for distinguishing experts and laypeople, showing that although the behavioural features were the most discriminative ones, the combination of behavioural features with medicine-related features reached the best results.

One of the limitations of this study is the lack of clickthrough information, which would have allowed us to perform a more detailed analysis of search behaviour. A further limitation is that MetaMap can only annotate

English text. Laypeople in particular prefer to query in their own language, as is clear from the high number of non-English queries that were removed from the HON search logs for this study. There is certainly a vast amount of work to be done for supporting such a query analysis for languages other than English, in particular due to the lack of such detailed language resources for many languages. MeSH on the other hand exists for many languages and mapping tools do exist. Still, detecting language of very short queries is not easy to do, so a multilingual scenario has many additional challenges.

The results of our analysis can be used to better understand the users through building detailed user profiles based on user behaviour in order to provide users with documents and query suggestions suited to their level of expertise. We can also identify new features for improving a search engine, such as the suggestions arising from this analysis to add a filter or facets for body regions to the GoldMiner search engine.

By using logfiles of several domain-specific medical search engines, this paper explores complementary information to most analyses of medical log files that either use general search engine logs or PubMed logfiles. This allows us to obtain information on user groups in a different way compared to general search engines where assumptions have to be made that can influence the analysis of the group behaviour.

Acknowledgements This research was partly funded by the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 257528 (KHRESMOI), partly funded by Horizon 2020 program (H2020-ICT-2014-1) under grant agreement n° 644753 (KCONNECT), and partly funded by the Austrian Science Fund (FWF) project n° I1094-N23 (MUCKE)

References

1. A. R. Aronson. Effective mapping of biomedical text to the UMLS Metathesaurus: the MetaMap program. pages 17–21, 2001.
2. A. R. Aronson, O. Bodenreider, H. F. Chang, S. M. Humphrey, J. G. Mork, S. J. Nelson, T. C. Rindfleisch, and W. J. Wilbur. The NLM Indexing Initiative. pages 17–21, Lister Hill National Center for Biomedical Communications (LHNCBC), National Library of Medicine, Bethesda, MD 20894, USA., 2000.
3. A. R. Aronson and F. Lang. An overview of metamap: historical perspective and recent advances. *JAMIA*, 17(3):229–236, 2010.
4. A. R. Aronson and T. C. Rindfleisch. Query expansion using the UMLS Metathesaurus. *Proceedings of the AMIA Annual Symposium*, pages 485–489, 1997.
5. S. K. Bhavnani. Domain-specific search strategies for the effective retrieval of healthcare and shopping information. In *CHI '02 Extended Abstracts on Human Factors*

- in *Computing Systems*, CHI EA '02, pages 610–611. ACM, 2002.
6. C. Boyer, V. Baujard, and A. Geissbuhler. Evolution of Health Web certification through the HONcode experience. *Stud Health Tech Inform*, 169:53–7, 2011.
 7. D. J. Brenes and D. Gayo-Avello. Stratified analysis of AOL query log. *Inf. Sci.*, 179(12):1844–1858, 2009.
 8. M.-A. Cartright, R. W. White, and E. Horvitz. Intentions and attention in exploratory health search. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*, SIGIR '11, pages 65–74, New York, NY, USA, 2011. ACM.
 9. J. Cohen. *Statistical Power Analysis for the Behavioral Sciences (2nd Edition)*. Routledge, 2 edition, July 1988.
 10. M. J. Cole, J. Gwizdka, C. Liu, N. J. Belkin, and X. Zhang. Inferring user knowledge level from eye movement patterns. *Information Processing & Management*, 49(5):1075 – 1091, 2013.
 11. K. Collins-Thompson, P. N. Bennett, R. W. White, S. de la Chica, and D. Sontag. Personalizing web search results by reading level. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management*, CIKM '11, pages 403–412, New York, NY, USA, 2011. ACM.
 12. D. Demner-Fushman, S. M. Humphrey, N. C. Ide, R. F. Loane, J. G. Mork, P. Ruch, M. E. Ruiz, L. H. Smith, W. J. Wilbur, and A. R. Aronson. Combining resources to find answers to biomedical questions. In *Proceedings of The Sixteenth Text REtrieval Conference, TREC 2007, Gaithersburg, Maryland, USA, November 5-9, 2007*, 2007.
 13. J. C. Denny, J. D. Smithers, R. A. Miller, and A. Spickard. "Understanding" medical school curriculum content using KnowledgeMap. *Journal of the American Medical Informatics Association*, 10(4):351–362, 2003.
 14. S. Duarte Torres, D. Hiemstra, and P. Serdyukov. Query log analysis in the context of information retrieval for children. In *Proceeding of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 847–848, New York, July 2010. ACM.
 15. G. B. Duggan and S. J. Payne. Knowledge in the head and on the web: using topic expertise to aid search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, pages 39–48, 2008.
 16. Eurobarometer. European citizens' digital health literacy. Technical report, European Commission, November 2014.
 17. S. Fox. Health topics. Technical report, The Pew Internet & American Life Project, February 2011.
 18. S. Fox and M. Duggan. Health online 2013. Technical report, The Pew Internet & American Life Project, January 2013.
 19. D. Gayo-Avello. A survey on session detection methods in query logs and a proposal for future evaluation. *Inf. Sci.*, 179(12):1822–1843, May 2009.
 20. L. Goeuriot, L. Kelly, W. Li, J. Palotti, P. Pecina, G. Zuccon, A. Hanbury, G. J. F. Jones, and H. Müller. ShARe/CLEF eHealth Evaluation Lab 2014, Task 3: User-centred Health Information Retrieval. In *Working Notes for CLEF 2014 Conference, Sheffield, UK, September 15-18, 2014.*, pages 43–61, 2014.
 21. D. He and A. Göker. Detecting session boundaries from web user logs. In *Proceedings of the BCS-IRSG 22nd annual colloquium on information retrieval research*, pages 57–66, 2000.
 22. J. Herskovic, L. Tanaka, W. Hersh, and E. Bernstam. A Day in the Life of PubMed: Analysis of a Typical Day's Query Log. *Journal of the American Medical Informatics Association*, 14(2):212–220, 2007.
 23. V. Hollink, T. Tsirikika, and A. P. de Vries. Semantic search log analysis: A method and a study on professional image search. *Journal of the American Society for Information Science and Technology*, 62(4):691–713, 2011.
 24. I. Hsieh-Yee. Effects of search experience and subject knowledge on the search tactics of novice and experienced searchers. *Journal of the Association for Information Science and Technology*, 1993.
 25. R. Islamaj Dogan, G. C. Murray, A. Névél, and Z. Lu. Understanding PubMed® user search behavior through log analysis. *Database*, 2009, Jan. 2009.
 26. A. S. Jadhav, A. P. Sheth, and J. Pathak. Online information searching for cardiovascular diseases: An analysis of mayo clinic search query logs. *Studies in Health Technology and Informatics*, pages 702–706, 2014.
 27. B. Jansen, A. Spink, and I. Taksai. *Handbook of research on Web log analysis*. Information Science Reference - IGI Global Publishing, Hershey, PA, 2008.
 28. B. J. Jansen and A. Spink. How are we searching the world wide web?: a comparison of nine search engine transaction logs. *Information Processing & Management*, 42(1):248–263, Jan. 2006.
 29. B. J. Jansen, A. Spink, J. Bateman, and T. Saracevic. Real life information retrieval: a study of user queries on the web. *SIGIR Forum*, 32(1):5–17, Apr. 1998.
 30. R. Jones and K. L. Klinkner. Beyond the session timeout: Automatic hierarchical segmentation of search topics in query logs. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management*, CIKM '08, pages 699–708, New York, NY, USA, 2008. ACM.
 31. M. Kritz, M. Gschwandtner, V. Stefanov, A. Hanbury, and M. Samwald. Utilization and perceived problems of online medical resources and search tools among different groups of european physicians. *Journal of Medical Internet Research*, Jun 2013.
 32. E.-M. Lacroix and R. Mehnert. The US National Library of Medicine in the 21st century: expanding collections, nontraditional formats, new audiences. *Health Information and Libraries Journal*, 19(3):126–132, 2002.
 33. M. Lui and T. Baldwin. Langid.py: An off-the-shelf language identification tool. In *Proceedings of the ACL 2012 System Demonstrations*, ACL '12, pages 25–30, Stroudsburg, PA, USA, 2012. Association for Computational Linguistics.
 34. E. Meats, J. Brassey, C. Heneghan, and P. Glasziou. Using the Turning Research Into Practice (TRIP) database: how do clinicians really search? *Journal of the Medical Library Association*, 95(2):156–63, 2007.
 35. A. Névél, R. I. Dogan, and Z. Lu. Semi-automatic semantic annotation of pubmed queries: A study on quality, efficiency, satisfaction. *Journal of Biomedical Informatics*, 44(2):310–318, 2011.
 36. A. Névél, W. Kim, W. J. Wilbur, and Z. Lu. Exploring two biomedical text genres for disease recognition. In *Proceedings of the Workshop on Current Trends in Biomedical Natural Language Processing*, BioNLP '09, pages 144–152, Stroudsburg, PA, USA, 2009. Association for Computational Linguistics.
 37. NLM. *UMLS Reference Manual*. Bethesda (MD): National Library of Medicine (US), Sept. 2009.
 38. J. Palotti, A. Hanbury, and H. Muller. Exploiting health related features to infer user expertise in the medical domain. In *Proceedings of WSCD Workshop on Web Search and Data Mining*. John Wiley & Sons, Inc., 2014.

39. J. Palotti, V. Stefanov, and A. Hanbury. User intent behind medical queries: An evaluation of entity mapping approaches with metamap and freebase. In *Proceedings of the 5th Information Interaction in Context Symposium, IiX '14*, pages 283–286. ACM, 2014.
40. J. Palotti, G. Zuccon, L. Goeuriot, L. Kelly, A. Hanbury, G. J. F. Jones, M. Lupu, and P. Pecina. ShARe/CLEF eHealth Evaluation Lab 2015, Task 2: User-centred Health Information Retrieval. In *Working Notes for CLEF 2015 Conference, Toulouse, France, September 8-11, 2015.*, 2015.
41. G. Pass, A. Chowdhury, and C. Torgeson. A picture of search. In *Proceedings of the 1st international conference on Scalable information systems, InfoScale '06*, New York, NY, USA, 2006. ACM.
42. W. Pratt and M. Yetisgen-Yildiz. A study of biomedical concept identification: Metamap vs. people. In *AMIA Annual Symposium Proceedings*, volume 2003, pages 529–533. American Medical Informatics Association, 2003.
43. K. Roberts, M. Simpson, D. Demner-Fushman, E. Voorhees, and W. Hersh. State-of-the-art in biomedical literature retrieval for clinical cases: A survey of the TREC 2014 CDS Track.
44. J. Schwarz and M. Morris. Augmenting web pages and search results to support credibility assessment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11*, pages 1245–1254, New York, NY, USA, 2011. ACM.
45. C. Silverstein, H. Marais, M. Henzinger, and M. Moricz. Analysis of a very large web search engine query log. *SIGIR Forum*, 33(1):6–12, Sept. 1999.
46. F. Silvestri. Mining query logs: Turning search usage data into knowledge. *Foundations and Trends in Information Retrieval*, 4(1:2):1–174, Jan. 2010.
47. A. Spink, Y. Yang, J. Jansen, P. Nykanen, D. P. Lorence, S. Ozmutlu, and H. C. Ozmutlu. A study of medical and health queries to web search engines. *Health Information & Libraries Journal*, 21(1):44–51, Mar. 2004.
48. T. Tsirikka, H. Müller, and C. Kahn Jr. Log analysis to understand medical professionals' image searching behaviour. In *Medical Informatics Europe*, 2012.
49. T. M. Walsh and T. A. Volsko. Readability assessment of internet-based consumer health information. *Respiratory care*, 53(10):1310–1315, 2008.
50. L. Wang, J. Wang, M. Wang, Y. Li, Y. Liang, and D. Xu. Using Internet Search Engines to Obtain Medical Information: A Comparative Study. *Journal of Medical Internet Research*, 14(3):e74, May 2012.
51. M. Weeber, H. Klein, A. R. Aronson, J. G. Mork, L. T. W. de Jong-van den Berg, and R. Vos. Text-based discovery in biomedicine: the architecture of the dad-system. In *Proceedings of the AMIA Symposium*, pages 903–907, 2000.
52. R. W. White, S. T. Dumais, and J. Teevan. Characterizing the influence of domain expertise on web search behavior. In *Proceedings of the Second ACM International Conference on Web Search and Data Mining, WSDM '09*, pages 132–141, New York, NY, USA, 2009. ACM.
53. R. W. White and E. Horvitz. Cyberchondria: Studies of the escalation of medical concerns in web search. *ACM Transactions on Information Systems*, 27(4):23:1–23:37, Nov. 2009.
54. R. W. White and E. Horvitz. Studies of the onset and persistence of medical concerns in search logs. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval, SIGIR '12*, pages 265–274, New York, NY, USA, 2012. ACM.
55. B. M. Wildemuth. The effects of domain knowledge on search tactic formulation. *Journal of the Association for Information Science and Technology*, 55(3):246–258, Feb. 2004.
56. X. Yan, R. Y. Lau, D. Song, X. Li, and J. Ma. Toward a semantic granularity model for domain-specific information retrieval. *ACM Transactions on Information Systems*, 29(3):15:1–15:46, July 2011.
57. P. Younger. Internet-based information-seeking behaviour amongst doctors and nurses: a short review of the literature. *Health information and libraries journal*, 27(1):2–10, Mar. 2010.
58. X. Zhang, M. Cole, and N. Belkin. Predicting users' domain knowledge from search behaviors. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '11*, pages 1225–1226. ACM, 2011.
59. Y. Zhang. Searching for specific health-related information in medlineplus: Behavioral patterns and user experience. *Journal of the Association for Information Science and Technology*, 65(1):53–68, 2014.
60. G. Zuccon, B. Koopman, and J. Palotti. Diagnose this if you can: On the effectiveness of search engines in finding medical self-diagnosis information. In *Advances in Information Retrieval*, pages 562–567. Springer, 2015.