

Health Information Retrieval

State of the art report

Carla Teixeira Lopes

Doctoral Programme in Informatics Engineering
Faculdade de Engenharia da Universidade do Porto, Portugal
carla.lopes@fe.up.pt

Supervisor: Maria Cristina Ribeiro

July 2008

Abstract

This report provides an overview of the field of Information Retrieval (IR) in health-care. It does not aim to introduce general concepts and theories of IR but to present and describe specific aspects of Health Information Retrieval (HIR). After a brief introduction to the more broader field of IR, the significance of HIR at current times is discussed. Specific characteristics of Health Information, its classification and the main existing representations for health concepts are described together with the main products and services in the area (e.g.: databases of health bibliographic content, health specific search engines and others). Recent research work is discussed and the most active researchers, projects and research groups are also presented. Main organizations and journals are also identified.

Contents

1	Introduction	1
1.1	Information Retrieval	1
1.1.1	Definition	1
1.1.2	Brief history	1
1.2	Health Information Retrieval	2
2	Health Information	3
2.1	Classification	4
2.2	Bibliographic content	4
2.3	Full-text content	5
2.4	Databases and Collections	6
2.5	Aggregations	7
3	Health Concept Representation	8
3.1	Introduction	8
3.2	Medical Subject Headings	9
3.3	Non-NLM thesaurus	10
3.4	Unified Medical Language System	11
3.5	Ontologies	15
4	Health Specific Search Engines	16
5	Recent Work	17
5.1	Health Information Seeking	17
5.1.1	Consumers	18
5.1.2	Professionals	18
5.1.3	Consumers and Professionals	19
5.2	Indexing	19
5.3	Retrieval	20
5.3.1	Health Information Characteristics	20
5.3.2	Query Expansion	21
5.3.3	Ranking	21
5.3.4	IR Models	22
5.4	Evaluation	22
5.5	User Interfaces and Visualization	23
6	People	23
7	Research Groups	24
8	Research Projects	25
9	Main Organizations	25

1 Introduction

This document describes the state of the art on Health Information Retrieval (HIR), a field of study where Information Retrieval (IR) techniques are applied in the specific domain of Health Information. It starts with a brief introduction to the broader field of IR (Section 1.1) where its definition and a short description of its evolution through the years is given, followed by an introduction to the more specific field of HIR (Section 1.2). Section 2 presents a classification of Health Information in several categories where several examples of products and services are given. In Section 3 the main existing representations for health concepts are described, from terminologies to ontologies. Section 4 enumerates and briefly describes the main health specific search engine and Section 5 describes some of the recent research work developed in the area. In sections 6, 7, 8, 9 and 10 are presented major figures in the area along with its research groups, research projects, main organizations and also journals.

1.1 Information Retrieval

1.1.1 Definition

Information retrieval (IR) is a broad and loosely-defined term [88, 55] that has become accepted with the work published by Cleverdon, Salton, Sparck Jones, Lancaster and others [88]. One of the earliest definitions, that dates from the 1960's, characterizes it as an area concerned with the structure, analysis, organization, storage, searching and retrieval of information [91]. More recently, Baeza-Yates et al. [4] presented a similar definition where IR deals with the representation, storage, organization and access to information items. An even more recent definition presents IR as “finding material (usually documents) of an unstructured nature (usually text) that satisfies an information need from within large collections (usually stored on computers)” [55].

Being an imprecise distinction, it is frequent to distinguish IR from Data Retrieval. Rijsbergen [88] does it in eight different perspectives: matching, inference, model, classification, query language, query specification, items wanted and error response. Baeza-Yates et al [4] also make this distinction defining data retrieval as retrieving all objects that satisfy clearly defined conditions where an erroneous object among the retrieved ones means a total failure. In IR it is not a major problem to identify non-relevant documents in the retrieved set of documents because this field of study deals with information that is not structured and that may be semantically ambiguous (the opposite of Data Retrieval) [4, 55, 2]. Research in database systems is usually associated with the Data Retrieval field.

1.1.2 Brief history

The first manual IR systems appeared with the Sumerians, in the beginning of 3000 BC, when they constructed areas for storing and classifying written materials (cuneiform inscriptions - one of history's oldest writing system) to assist the operation of various social groups [45]. As time went by, inventions like paper, the printed press raised the importance of systems to store, manage and retrieve information. After the invention of computers, in 1945, Vannevar Bush wrote the article “As We May Think” [8] in which he criticizes the artificiality of the epoch's indexing systems. He idealizes a system that operates through association (as the

human mind) and is mechanized so it may be consulted with speed and flexibility. This was the first conception of an automatic IR system. The term Information Retrieval was coined later, in 1950, by Calvin Mooers [62].

Since the late 1950s, the field of IR has evolved through several relevant works such as: H.P. Luhn's work [52]; the SMART system by Salton and his students [92] where some important IR concepts (such as the vector space model and relevance feedback) were developed; Cleverdon's Cranfield evaluation model [11]; Sparck Jones development of idf [40] and the probabilistic models by Robertson [89] and Croft [14, 103]. In 1992, with the beginning of the Text REtrieval Conference (TREC) [65], that provide the necessary infrastructure for large-scale evaluation, allowed the modification of old models/techniques and the proposal of new ones.

With the development of the World Wide Web (Web), the interest in the IR field has spread from information specialists (librarians, the legal community and other) [2, 9, 55] to a broader audience [2]. The larger interest on the area together with the increase in the amount of information available and in users' requirements have contributed to significant progresses on the field that has evolved from a simple document retrieval to a broad range of related areas. Some of these areas are: question-answering, cross-lingual search, topic detection and tracking, summarization, multimedia retrieval and others [2, 9].

Despite the recent and numerous developments, IR is far from being a "solved problem" [9]. On one side, information is being produced more than ever – a report of a study developed in UC Berkeley's School of Information Management and Systems [53] estimates that the amount of new information stored on paper, film, magnetic, and optical media has doubled and that the information on the Web has tripled from 2000 to 2003. On the other side, the ways people produce, search, manage and use information is rapidly evolving.

1.2 Health Information Retrieval

Probably by the same reasons that made the general IR field evolve so much in the last years, research and interest on the application of this field's techniques to the specific domain of health and biomedicine have also grew. The Web and its applications have profoundly changed the availability and ease of access to health information, not only to health professionals but also to consumers.

To health professionals, applications providing an easy access to validated and up-to-date health knowledge are of great importance to the dissemination of knowledge and have the potential to impact the quality of care provided by health professionals. On the other side, the Web opened doors to the access of health information by patients, their family and friends, making them more informed and changing their relation with health professionals. According to Adam Bosworth [6], in a good health system, consumers should be able to discover the most relevant information possible, should have access to health support personalized services and should be able to learn from and educate those in similar health circumstances.

To professionals, one of the main and oldest IR applications is PubMed from the US National Library of Medicine (NLM) that gives access to the world's medical research literature. To consumers, health information is available through different services and with different quality. Lately, the control over and access to health information by consumers has been a hot topic, with plenty government initiatives all over the world that aim to improve

consumer health giving consumers more information and making easier the sharing of patient records. Some of the main IR related companies, like Google and Microsoft are beginning to invest in this area. Google launched, at May 2008, a service entitled Google Health [23] and Microsoft has also launched, at October 2007, a service entitled Microsoft HealthVault [60]. Revolution Health [87] is another service of this type provided recently from Steve Case's (AOL founder) company. Generally, the movement of a consumer-driven healthcare is about giving consumer more power to manage their health. This can be done by giving them access to their health record anywhere/anytime, making easy to find all needed information, providing online tools that enable personalized advice and other approaches.

Currently, the Web is already used in large scale to access health information. A 2006 study [20] concluded that "eight in ten USA internet users go online for health information". They also found that: health searches are as popular as reading blogs or using the Internet to look up a phone number or address; "the typical health information session starts at a search engine, includes multiple sites, and is undertaken on behalf of someone other than the person doing the search"; 53% of health seekers report that most recent health information session had some kind of impact on how they take care of themselves or care for someone else.

The revolution in health care caused by the new information society is also prognosed by Haux et al. [27] in several dimensions. Two of their theses are "Patients and their families will be knowledgeable of the information resources available over the Internet and will make use of them. New services will arise" (they predict the number of access to health web sites will increase by more than 30% and that over 20% of patients will inform themselves about health on the Internet) and that "Knowledge about diseases will be current, comprehensive and internationally available via electronic media, including their availability to patients and their family members ('consumers'). This knowledge will be available in different qualities. Therefore, internationally accredited certification will be available for their contents (e.g. by specialty associations). Knowledge support will partially be integrated in clinical routines." (they predict that over 80% of polled medical knowledge will result from electronic media and that over 80% of the guidelines used routinely in clinical work, will be available electronically).

2 Health Information

Information is a concept hard to define and may be viewed in several ways. It's frequent to see its definition through its comparison to the concepts of data and knowledge. Data consist of the observations or measurements made about the world. Information is the aggregated and organized data that describes a specific situation. Knowledge is what is learned from data and information, accumulated and integrated over time, that can be applied in new situations.

Current times are integrated in the third era of technological evolution, the so-called information era [106]. Never before, so much information has been created and transferred. A 2003 study estimated that new stored information grew about 30% a year between 1999 and 2002 [53]. In fact, information is in the center of our society and it has become an indispensable resource to every society's areas.

In healthcare, information plays a crucial role in professionals activities and consumers attitudes. Two old studies (one from 1966 and other from 1973) predicted that health-care personal spend about one-third of their time handling and using information [41, 54]. According to Hersh [30], it is likely that the time dedicated to managing information in healthcare is as large, or even larger, nowadays.

In this report, the term Health Information is adopted instead of Biomedical Information, because it's broader than the latest, including not only concepts from biological and medical sciences but also information on related areas such as health care (e.g.: health care facilities, manpower and services; health care economics and organizations).

2.1 Classification

Textual Health Information may be classified in patient-specific information and knowledge-based information [30, 96]. The first type relates to individual patients and its purpose is to tell health professionals about the health condition of a patient. It typically comprises the patient's medical record and it may contain unstructured data (e.g.: lab results, vital signs) or free/narrative text (e.g.: radiology report). The second type of classification is related to the information that derives and is organized from observational and experimental research. This information provides health professionals the knowledge acquired in other situations so it may be applied to individual patients or used to conduct further research. Similarly to other types of scientific information, the knowledge-based information can be subdivided in primary information (direct results of original research that appears in papers, journals or other sources) and secondary information (reviews, condensations, summaries of primary literature like books, monographs, review papers, clinical guidelines, health information on web pages and other sources).

Another way of classifying knowledge-based information is to divide it in four subcategories: bibliographic, full-text, databases/collections and aggregations [30, 96]. Each subcategory will be described next along with some of its main examples.

2.2 Bibliographic content

The bibliographic content contains literature reference databases (or bibliographic databases), web catalogs and specialized registries. The distinction between these subcategories is becoming blurry as, for example, literature reference databases started providing links to the referenced literature, moving closer to web catalogs.

Literature Reference Databases These databases catalog books and periodicals and were the original IR databases in the past 1960s, designed to guide the searcher and not to provide information.

MEDLINE (Medical Literature Analysis and Retrieval System Online) is probably the best-known and premier bibliographic database of the National Library of Medicine (NLM) [69]. It started being the print publication of *Index Medicus*, a print catalog of the medicine literature, with its first volume published in 1879 [30]. It contains over 16 million references to journal articles of approximately 5200 journals with a subject scope of biomedicine and health. In 2007, 670,000 references were added.

MEDLINE is freely available via PubMed¹ and a search generates a list of citations (including authors, title, source, and often an abstract) to journal articles, an indication of free electronic full-text availability (generally through PubMed Central²) or a link to the website of the publisher or other full text provider. It may also be searched using the NLM Gateway, a single Web interface that searches multiple NLM retrieval systems [70]. Other websites also provide access to MEDLINE, some for free and others for some fee (usually providing value-added services).

Besides MEDLINE, NLM has many other databases and electronic resources (listed in [73]). Their bibliographic databases are organized in three categories: citations to journals and other periodicals since 1966 (accessible through PubMed that is composed by MEDLINE, MEDLINE in-process citations and publisher-supplied citations), citations to books, journals and audiovisual material (available through LOCATORplus³) and citations to journal articles prior to 1966 and scientific meeting abstracts (accessible through NLM Gateway⁴).

In addition to NLM, there are other producers of bibliographic databases, both public and private. Some are produced by other USA National Institute of Health's institutes like the National Cancer Institute (NCI). Some of these bibliographic databases tend to be more focused in specific resource and subject types like CINAHL (Cumulative Index to Nursing and Allied Health Literature⁵) – the major non-NLM database for the nursing field.

Web Catalogs Web catalogs are web pages that contain links to other pages and sites and share many features with traditional bibliographic databases [30]. The number of such catalogs is increasing [96]. Some well-known catalogs are: MedicalMatrix⁶, Hardin.MD⁷, HealthFinder⁸, HON Select⁹, Intute: Health and Life Sciences¹⁰ and MedWeb Plus¹¹.

Specialized Registries Specialized registries differ from literature reference databases because they point to more diverse information resources. This type of information resource may overlap with literature reference databases and web catalogs but, generally, it points to more diverse information resources. One famous specialized registry is the *National Guidelines Clearinghouse* (NGC¹²), produced by the Agency for Healthcare Research and Quality (AHRQ), with information about clinical practice guidelines.

2.3 Full-text content

This subcategory contains online versions of the complete version of periodicals, books and websites. Originally, full-text databases were mainly online versions of journals and it only started to include books with the decrease of the price of computers and the growth of the

¹<http://pubmed.gov>

²<http://www.pubmedcentral.nih.gov>

³<http://locatorplus.gov>

⁴<http://gateway.nlm.nih.gov>

⁵<http://www.cinahl.com>

⁶<http://www.medmatrix.org>

⁷<http://www.lib.uiowa.edu/hardin/md>

⁸<http://www.healthfinder.gov>

⁹<http://www.hon.ch/HONselect>

¹⁰<http://www.intute.ac.uk/healthandlifesciences>

¹¹<http://www.medwebplus.com>

¹²<http://www.guideline.gov/>

Internet that also led to development of web sites.

Most periodicals are nowadays published electronically. Some are published by the company that produces the print version of the journal (e.g.: Elsevier), others by Highwire Press (which provides a website, searching and browsing interfaces and other tools to publishers that haven't moved directly to the online world), some have created their own site and others have exclusively published online. Electronic versions are usually enhanced by additional features such as: easier access, provision of additional data like figures, tables, raw data and images, true bibliographic links. Nowadays, there are publishers that have decided to make their papers freely available on the Web. Some highly visible approaches are BiomedCentral¹³ (BMC), Public Library of Science¹⁴ (PLOS) and PubMed Central¹⁵ (PMC).

Textbooks are also increasingly publishing versions on the Web. These electronic versions allow several additional features over the printed versions: they feature high-quality images; the insertion of multimedia content (e.g.: sound and video) which is impossible in the printed version; the insertion of links to other resources; the inclusion of interactive self-assessment questions and an easier access to book updates. There are already several online versions of well-known books in the health area.

The third type of full-text content is composed of web sites that provide full-text information. This excludes web sites that implement services such as bibliographic databases, online versions of books and other printed material, specialized databases and collections and aggregations of these. The Web has several health full-text information sites that are developed by everyone, from consumers to governments. Some of these websites, that include more than just collections of text (e.g.: interaction with experts, links to other sites) are: Intelihealth¹⁶, Netwellness¹⁷, WebMD¹⁸, eMedicine¹⁹, Medscape²⁰ and Institute for Clinical Systems Improvement guidelines²¹. The first three sites are directed to consumers while the last three are more oriented to health professionals.

2.4 Databases and Collections

This category consists of databases and other specific collections of information. This kind of information is usually housed in database management systems and contains several types of resources like images (from radiology, pathology and other areas), genomics (gene sequencing, protein characterization and others), citations (which link scientific literature) and Evidence Based Medicine (EBM) resources. The dynamic nature of web databases make them more appropriate to some type of content.

Images, an important piece of the healthcare practice, education and research are one of those types. There are several health images databases available on the Web. One of the most

¹³<http://www.biomedcentral.com>

¹⁴<http://www.plos.org>

¹⁵<http://pubmedcentral.gov>

¹⁶<http://www.intelihealth.com/>

¹⁷<http://www.netwellness.com/>

¹⁸<http://www.webmd.com/>

¹⁹<http://www.emedicine.com/>

²⁰<http://www.medscape.com/>

²¹http://www.icsi.org/guidelines_and_more/

famous is the Visible Human Project²², which consists of three-dimensional representations of normal male and female bodies built from cross-sectional slices of cadavers.

Genomics studies the genetic material in living organisms and its research has been evolving rapidly in the last years. One of its main driving forces was the Human Genome Project, led by the National Human Genome Research Institute at the National Institutes of Health (NIH). This project ended in April 2003 with the production of a version of the human genome sequence that is freely available in public databases [64]. Several genomics databases are available across the Web and at the center are those produced by the National Center for Biotechnology Information (NCBI). NCBI's databases are linked among themselves, along with PubMed in the NCBI's Entrez system²³.

The number of citations to a specific scientific work is usually used as an evaluation measure of that work's quality and impact. To track citations in scientific literature there are citation databases that provide linkages to articles that cite others. The best-known citation databases are the Science Citation Index (SCI) and the Social Sciences Citation Index (SSCI) from Thomson Reuters, available through the Web of Science service²⁴.

On the Web there are also databases that hold the principles of EBM and try to eliminate the problems of scattering and fragmentation of the primary literature. These databases provide systematic reviews (e.g.: The Cochrane Database of Systematic Reviews²⁵) or high synthesized synopses of evidence-based information (e.g.: Clinical Evidence²⁶, DynaMed²⁷, PIER²⁸, UpToDate²⁹).

2.5 Aggregations

This last category includes aggregations of the first three categories for all types of users, from consumers to health professionals and scientists. The distinction between this category and some of the above content with several links is blurry but, typically, aggregations have a larger variety of information that serve diverse needs of their users. They are, for example, websites that collect several types of content to generate a coherent resource.

One of the largest aggregated consumer information resource is MedlinePlus, a service of the U.S. NLM and the NIH that is updated daily. It aggregates information from these entities and from other trusted sources on over 750 diseases and conditions. It also has preformulated MEDLINE searches to give access to medical journal articles, information on drugs, an illustrated encyclopedia, a medical dictionary, links to clinical trials, interactive patient tutorials and updated health news [71]. CancerNet is another consumer-oriented service from the National Cancer Institute³⁰ (NCI) that contains information on all disease's aspects. Healthline³¹ is another service focused on the consumers with online health search,

²²http://www.nlm.nih.gov/research/visible/visible_human.html

²³<http://www.ncbi.nlm.nih.gov/Entrez/>

²⁴<http://scientific.thomson.com/products/wos/>

²⁵<http://www.cochrane.org>

²⁶<http://www.clinicalevidence.com>

²⁷<http://www.dynamicmedical.com>

²⁸<http://pier.acponline.org>

²⁹<http://www.uptodate.com>

³⁰<http://www.cancer.gov>

³¹<http://www.healthline.com>

content and navigation features. Kosmix RightHealth³² creates a page for health related topics based on information (in several formats) dispersed across the Internet. The two last consumer aggregations have been recently created and combine this service with others. The first is Microsoft HealthVault³³, a hub of a network of sites, personal health devices and other services to let consumers manage their health. The second is Google Health that has the same goal of giving control to health consumers, storing health information in one central place. Besides allowing the creation of online health profiles, the import of medical records, the searching of doctors and hospitals and the connection to online health services, it also aggregates resources about health topics and issues.

There are also aggregated content more directed to professionals like: MDConsult³⁴ – a service of Elsevier that aggregates medical resources in an integrated way to help health professionals and Merck Medicus³⁵ – developed by Merck, available to all licensed US physicians, which includes resources like Harrison’s Online and MDConsult.

3 Health Concept Representation

3.1 Introduction

The area of IR greatly benefits from the availability of well defined information structures that can be used in the indexing and retrieval processes of IR. Health information is, by its nature, highly detailed [96], where knowledge organization is one of the oldest applications of classification, dating to Aristotle’s effort in biology and formal descriptions [78]. Health concepts’ representation is more challenging than in many domains, due to its levels of precision, complexity, implicit knowledge and breadth of application [96]. However, it is also an area where great efforts have been developed and several representation’s systems have appeared.

Before describing some of the main health concept representation systems, they are classified in three categories with increasing degree of formalism: terminologies, controlled vocabularies/thesaurus and ontologies. A terminology or vocabulary is a list of terms, that are representations of the concepts used in a specific area. When simple relationships among different terms are specified, this representation system becomes a controlled vocabulary or a thesaurus. The relationships are typically of three types: hierarchical (terms are broader or narrower), synonymous, related (terms with relations that are neither hierarchical or synonymous). Ontologies are the most formal category having also logical descriptions that serve to computationally define terms. They must also exhibit internal consistency, acyclic polyhierarchies and computable semantics [96]. Recently there has been an explosion of modern biomedical ontologies [97].

The Health Concept Representation Systems described in this section may be used in several health informatics’ research areas such as Information Retrieval, Natural Language Processing, Semantic Interoperability and Decision Support Systems. In Information Re-

³²<http://www.righthealth.com>

³³<http://www.healthvault.com>

³⁴<http://www.mdconsult.com>

³⁵<http://www.merckmedicus.com>

trieval, they may be used in the indexing process (manual indexing is usually done using thesaurus) and in the retrieval process (e.g.: synonymous relations may be used to improve the expression of the information needs).

This section will start to describe the NLM's thesaurus that is used to index most of the NLM's databases, followed by other non-NLM thesaurus used in the health area. Then, NLM's Unified Medical Language System, together with its 3 knowledge sources: Metathesaurus, Semantic Network and the SPECIALIST Lexicon and Tools are presented. In the end, some of the main health ontologies and two general ontologies are briefly presented.

3.2 Medical Subject Headings

The Medical Subject Headings (MeSH) is the NLM's thesaurus used to index most of the NLM's databases [13]. It has sets of terms naming descriptors that are arranged in both alphabetic and hierarchical structure (which allows searching at various levels of specificity). The 2008 version of MeSH has 24,767 descriptors organized in an eleven-level hierarchy of headings [68]. the MeSH vocabulary files may be downloaded from the NLM site, at no charge, in XML or ASCII format.

The hierarchies in which descriptors are placed are also called trees³⁶. Each Descriptor appears in at least one place in the trees and may appear in as many additional places as may be appropriate. XML MeSH is structured in three levels: Descriptor/Concept/Term. A Descriptor may consist of a class of Concepts, that, correspond to a class of Terms which are synonymous with each other. For example³⁷:

- Cardiomegaly [Descriptor]
 - Cardiomegaly [Concept, Preferred]
 - * Cardiomegaly [Term, Preferred]
 - * Enlarged Heart [Term]
 - * Heart Enlargement [Term]
 - Cardiac Hypertrophy [Concept, Narrower]
 - * Cardiac Hypertrophy [Term, Preferred]
 - * Heart Hypertrophy [Term]

Each Concept has a Preferred Term which is also the name of the Concept and each Descriptor has a preferred Concept. The name of the Descriptor corresponds to the preferred Term of the preferred Concept. The terms in one Concept are not strictly synonymous with terms in another Concept, even in the same record. Additionally, MeSH has two types of relationships: hierarchical and associative [75]. The first type is a crucial component of a thesaurus and is formalized by the MeSH tree structure that represents distinct levels of specificity (terms that are broader or narrower). MeSH descriptors are organized in 16 categories [72] that can be explored through the browser available at ³⁸. The associative

³⁶2008's MeSH list of trees in: <http://www.nlm.nih.gov/mesh/trees2008.html>

³⁷Obtained from: http://www.nlm.nih.gov/mesh/concept_structure.html

³⁸<http://www.nlm.nih.gov/mesh/MBrowser.html>

relationships are often represented by the “see related” cross reference. They can be used to add/suggest terms to a specific search, to point out in thesaurus the existence of other descriptors which may be more appropriate or to point out distinctions made in the thesaurus or in the hierarchical structure of the thesaurus.

Besides the existence of descriptors (also called MeSH headings), MeSH has additional types of vocabulary: qualifiers (or subheadings), check tags, publication characteristics and supplementary concept records. Qualifiers can be attached to descriptors to narrow the focus of a term (e.g.: drug therapy, diagnosis, etiology, surgery). For example, a deficiency of monoamine oxidase is retrieved by the Descriptor Monoamine Oxidase combined with the Qualifier deficiency [77]. There are rules restricting the attachment of certain Qualifier (the allowable qualifiers are mentioned in the field *Allowable Qualifiers* for each term). Check tags are a special class of MeSH descriptors that must be considered routinely for every article (that’s why they’re called check tags) and represent facets like species, gender, human age, historical time periods and pregnancy [74]. Publication characteristics (or types) describes the item being indexed instead of its topic. It has 3 main categories: publication components (e.g. English Abstract), publication formats (e.g. lectures, letter), study characteristics (e.g. clinical trial, meta-analysis). The Supplementary Concept Records allow the indexing with headings from other thesaurus (non-MeSH headings) [74]. Currently, MeSH has 172,000 headings in the Supplementary Concept Records within a separate thesaurus [68].

MeSH’s use is not restricted to indexing NLM’s databases. It is also used by other organizations to index bibliographic content, such as health libraries and the National Guidelines Clearinghouse³⁹ [30].

3.3 Non-NLM thesaurus

In addition to Mesh there are other thesaurus in the health area used to index documents.

CINAHL⁴⁰, a database of nursing and allied health literature, uses the CINAHL Subject Headings, which is based on MeSH and has additional domain-specific terms [7]. EMBASE⁴¹, an european database of biomedical and pharmacological information, has a vocabulary called EMTREE⁴² with features similar to those of MeSH⁴³.

Other common vocabularies in the health area include the Logical Observation Identifier Names and Codes (LOINC) that provides a universal code system for reporting laboratory and other clinical observations in electronic messages [58], HL7 vocabulary tables⁴⁴ that identify, organize and maintain coded vocabulary terms used in HL7 messages and the National Drug Code Directory⁴⁵.

A last vocabulary worth notice is the Consumer Health Vocabulary⁴⁶, developed in an open source and collaborative initiative, that aims to help bridge the communication gap

³⁹<http://www.guideline.gov>

⁴⁰<http://www.cinahl.com>

⁴¹<http://www.embase.com>

⁴²<http://www.info.embase.com/entree/about>

⁴³Comparison at: http://www.info.embase.com/entree/about/entree_mesh_comparison.pdf

⁴⁴<http://www.hl7.org/Special/committees/Vocab/vocab.htm>

⁴⁵<http://www.fda.gov/cder/ndc>

⁴⁶<http://www.consumerhealthvocab.org>

between consumers and health professionals. With this vocabulary, technical terms like “myocardial infarction” may be translated into lay language like “heart attack”. In practice it can be used to improve IR systems, to help lay-people read and understand health-related information and others. This vocabulary has an online browser⁴⁷ and is organized by concepts and its associated terms.

3.4 Unified Medical Language System

The Unified Medical Language System (UMLS) started at the NLM, in 1986, by the hands of its Director, Donald Lindberg, as a “long-term research and development project” [50]. This project aimed at reducing barriers to the application of computers to the health area and more specifically to the effective retrieval of machine-readable information [50, 35]. Two of such barriers are the variety of ways to express a same concept in different vocabularies and the diffusion of useful information among different systems. In fact, the medical informatics field is characterized by a large diversity of vocabularies developed for specific applications (e.g.: epidemiological systems, medical expert systems, indexing literature, codes for billing and procedures and others). The lack of a common language barred the interoperability of the applications that used these vocabularies and was a motivation to the development of the UMLS.

The UMLS consists of three knowledge-sources that can be used separately or together. One is the Metathesaurus that has more than one million biomedical concepts from over 100 sources (including MeSH), another is the Semantic Network with 135 broad categories and 54 relationships between categories, the last one is the SPECIALIST Lexicon and Tools which has lexical information and programs for language processing [44]. The uses of these knowledge-sources can be very diverse (e.g. information retrieval, natural language processing, automated indexing, thesaurus construction, electronic health records and others). Each knowledge-source will be described in more detail in this section.

The UMLS is made available, at no cost, as 3 separate sets of relational files that are distributed on DVD or can be downloaded from the UMLS Knowledge Source Server (⁴⁸) and 3 tools: the MetamorphoSys (UMLS installation wizard and customization tool included in each UMLS release – ⁴⁹), the RRF Subset Browser (to find a term within a customized Metathesaurus subset or any vocabulary in the Rich Release Format (RRF) format) and the LVG (set of tools and data that are distributed with the UMLS as part of the SPECIALIST system).

Metathesaurus The UMLS Metathesaurus is a multi-source (thesauri, classifications, code sets, and lists of controlled terms used in patient care, health services billing, public health statistics, indexing and cataloging biomedical literature, and/or basic, clinical, and health services research [76]), multi-purpose and multi-lingual thesaurus of health related concepts, their various terms and the relationships among them. It is called a Metathesaurus as it transcends the thesauri, vocabularies and classifications it covers [76].

The list of the source vocabularies included in the 2008AA UMLS’s release is available in <http://www.nlm.nih.gov/research/umls/metaa1.html>. The Metathesaurus isn’t a new,

⁴⁷<http://samwise1.partners.org/vocab>

⁴⁸<http://umlsks.nlm.nih.gov/>

⁴⁹<http://www.nlm.nih.gov/research/umls/meta6.html>

single and unified standard vocabulary [36] and it doesn't contain logical assertions across terms from different vocabularies. Rather, it establishes conceptual linkages between its source vocabularies preserving their meanings, concept names and relationships.

In the Metathesaurus, synonymous terms are clustered into a concept with a unique identifier (CUI). Each term, identified by a unique identifier (LUI), is a normalized name and may have several strings (identified by SUI), which represent the terms' lexical variants in the source vocabularies. Each string is associated with one or more atoms (identified by AUI) that represent the concept name in the source. For example (from [44]), the concept Headache (C0018681) has the following structure:

- headache (L0018681)
 - headaches (S1459113)
 - * headaches (A1412439) – BI
 - Headache (S0046854)
 - * Headache (A2882187) – SNOMED
 - * Headache (A0066000) – MeSH
- cranial pain (L1406212)
 - Cranial Pain (S1680378)
 - * Cranial Pain (A1641293) – MeSH
- cephalgia head pain (L0290366)
 - HEAD PAIN CEPHALGIA (S0375902)
 - * HEAD PAIN CEPHALGIA (A0418053) – DxP

If two different terms have different meanings (e.g. cold) they are assigned the same LUI that stays associated with different CUI (e.g. cold temperature, common cold, cold sensation). The same can happen with strings and concepts.

The Metathesaurus is distributed in two formats: Original Release Format (ORF) and Rich Release Format (RRF). The first is centred on the concept level. The second appeared later and is a source-centric view that supports source transparency (it has the atom level that represents the original source information and also new data fields and files to capture source specific identifiers and mappings). The access to the Metathesaurus can be made through the UMLS KS or the RRF Browser in MetamorphoSys.

Semantic Network The Semantic Network is an upper-level ontology in the health field [10] composed of 135 *Semantic Types*, which may be assigned to Metathesaurus' concepts and 54 *Semantic Relationships*, a set of relations that may hold between Semantic Types. The Semantic Types are the nodes in the Network and the Relationships are the links. It is provided in a relational table format and in a unit record format [66].

Semantic Types are organized into 2 hierarchies: *Entity* and *Event* and its current scope is very broad, allowing the semantic categorization of a wide range of terminology [66]. Each Metathesaurus's concept is assigned at least one semantic type (the most specific type

available in the hierarchy). Instead of adding semantic types to the Network to encompass an object in the most appropriate categories, concepts that don't belong in a granularity level, must be associated with a Type of an upper level. For example, the Semantic Type *Manufactured Object* has two child nodes: *Medical Device* and *Research Device*. If an object is neither a medical device nor research devices, it is simply assigned to the more general type *Manufactured Object* [66].

Semantic Relationships may be hierarchical or associative (non-hierarchical). The *isa* link is the primary link in the Network and is the one that establishes the hierarchy of types and relations (e.g.: animal *isa* organism and treats *isa* affects). The set of associative are grouped into five major categories (which are also relations): *physically related to*, *spatially related to*, *temporally related to*, *functionally related to* and *conceptually related to*. Whenever possible, relations are defined between the highest level semantic types and, generally, are inherited by all the children of those types. The relations do not necessarily apply to all instances of concepts that have been assigned to the semantic types that are the nodes of that link. If it makes no sense, the inheritance of relations may also be blocked to a single or all of the children of the semantic types that is links. For example, “*conceptual part of* links *Body System* and *Fully Formed Anatomical Structure*, but it should not link *Body System* to all the children of *Fully Formed Anatomical Structure*, such as *Cell* or *Tissue*” [66]. A portion of the UMLS Semantic Network can be seen in ⁵⁰.

SPECIALIST Lexicon and Tools The SPECIALIST Lexicon and Tools has two main components: the lexicon and the tools. The Lexicon is a general English lexicon of common words that includes many biomedical terms and was developed to provide information to the SPECIALIST Natural Language Processing (NLP) System [67]. The Tools are programs that process terms. Together they pre-process terms before their introduction in the Metathesaurus and they are very useful to NLP applications in health domain.

The SPECIALIST Lexicon's entries record the syntactic (how the words are put together), morphological (inflection, derivation and compounding) and orthographic (spelling) information of each term/word. Lexical items may be composed by more than one term, if it is how it is presented in English, medical dictionaries or thesauri, or if it is an expansion of generally used acronyms and abbreviations.

Each unit lexical record has attributes (which are called *slots*) and values (which are called *fillers*) and is delimited by braces ({}). The unit lexical records for “anaesthetic” (one as a noun entry and the other as an adjective entry) will help to show some of the features of the Lexicon and the characteristics of lexical records (example from [67]). The noun lexical record for “anaesthetic” is:

```
{
base=anesthetic
spelling_variant=anaesthetic
entry=E0330018
cat=noun
variants=reg
variants=uncount
}
```

⁵⁰http://www.nlm.nih.gov/research/umls/META3_Figure_3.html

Every lexical record has the slot *base* that indicates the base form of the term. Optionally, it may also have one or more spelling variants expressed by slots *spelling_variant*. The next slot (*entry*) contains the unique identifier (EUI) of the record that has an “E” followed by seven digit numbers. The slot *cat* exists in every record and indicates the syntactic category of the entry (e. g.: noun, adjective, verb). The *variants* slot indicates the inflectional morphology of the entry. In the example given, these slots indicate that “anesthetic” is a count noun which undergoes regular English formation (“anaesthetics”). In the adjective lexical record:

```
{
base=anesthetic
spelling_variant=anaesthetic
entry=E0330019
cat=adj
variants=inv
position=attrib(3)
position=pred stative
}
```

the *variants* slot has the filler *inv* that indicates that the adjective “anesthetic” doesn’t form a comparative or superlative. The first *position* slot indicates that the adjective is attributive and appears after color adjectives in the normal adjective order. The second *position* slot indicates this adjective can appear in predicate position.

Other slots indicate the complementation (e. g.: in verbs if it is a intransitive or transitive verb), derivation (e. g.: adjective/noun – red/redness) and spelling variants (e. g.: British-American variants – centre/center) of each entry. For more detailed information see [67].

Lexical entries are independent of semantics, representing only a spelling-category pair. If different senses have the same spellings and syntactic category, they are represented by a single lexical entry in the Lexicon.

The Lexical Tools is a set of three programs implemented in Java: a lexical variant generator (lvg), a word index generator (Wordind) and a normalizer (NORM). They are designed to help dealing with the high degree of variability in natural language words (e. g.: treat, treats, treated, treating) and even in the order of words in “multi-word” terms.

The lvg program performs lexical transformations of input words. It consists of several flow components that can be combined (for example, the flow *i* simply generates inflectional variants and the flow *l:i* generates the same inflectional variants but in lowercase).

The WordInd breaks strings into words and produces the Metathesaurus word index (MRXW). The use of this program before searching in the word index assures congruence the words to be looked up and the word index. The program outputs one line for each word found in the input string (e.g.: for the input string *Heart Disease, Acute*, are returned three lines for the three words: *heart, disease* and *acute*). The output words are always presented in lowercase.

The NORM program generates normalized strings that are used in the normalized string index (MRXNS). This program must be used before MRXNS can be accessed and it is a selection of LVG transformations (in fact, it is the lvg program with the N pre-selected flow option). The normalization process involves stripping possessives (e.g.: *Hodgkin’s diseases*,

NOS – Hodgkin diseases, NOS), removing stop words (e.g.: *Hodgkin diseases, NOS – Hodgkin diseases*), lower-casing each words (e.g.: *Hodgkin diseases, – hodgkin diseases*), replacing punctuation with spaces (e.g.: *hodgkin diseases, – hodgkin diseases*), breaking a string into its constituent words/uninflecting (e.g.: *hodgkin diseases – hodgkin disease*) and sorting the words in alphabetic order (e.g.: *hodgkin disease – disease hodgkin*).

For more information on the SPECIALIST lexicon, lexical variant programs, and lexical databases, see [57].

3.5 Ontologies

An initial definition of ontology referred to the set of primitive entities that describes and models a specific knowledge domain and should reflect its underlying reality [25]. In computer science, ontology means an organization of concepts in domains, exhibiting internal consistency, acyclic polyhierarchies and computable semantics [10].

A health ontology aims to study classes of health significant entities such as substances (e.g.: mitral valve), qualities (e.g.: diameter of the left ventricle) and processes (e.g.: blood circulation) [96]. In theory an ontology is different from a terminology because it is concerned with the definition of health classes and the relations among them, while the last just collects names of health entities. In practice, this distinction is less obvious and often the products developed fall in between terminologies and ontologies as they are more than lists of terms but do not necessarily meet the requirements of formal organization [96].

Ontologies are increasingly playing an important role in medical informatics research [63] (e.g.: natural language processing, interoperability, SemanticWeb) where they also act as an enabling resource in several applications. Some of the main health ontologies and two other general ontologies will be presented next.

GALEN (Generalised Architecture for Languages, Encyclopedias, and nomenclatures in medicine) was the name of an European Union project (1992-1999) that illustrated how medical concepts could be represented as a formal ontology and how this could be used in practical applications [84]. One of this project's core features is the Common Reference Model, an ontology that aims to represent "all and only sensible medical concepts", whose access is made through OpenGALEN⁵¹. GALEN provides the blocks required for describing terminologies and a mechanism for combining concepts. For example, it has explicit representations for *adenocyte* and for *thyroid gland* and instead of having one for *adenocyte of thyroid gland*, it has an indication that these concepts can be combined. GALEN has an hierarchies of categories and a "rich hierarchy of associative relationships used to define complex structures" [96].

As said before, when describing the Semantic Network component of the UMLS, this component serves as a basic, high-level ontology for the biomedical domain [56]. Its semantic types are used to categorize UMLS concepts that, in turn, must be assigned to, at least, one semantic type.

The Systematized Nomenclature of Medicine Clinical Terms (SNOMED CT) is the most complete biomedical terminology developed in native description logic formalism [96]. It has a good concept coverage with over 361,800 concepts (as of July 2004), organized in

⁵¹<http://www.opengalen.org>

eighteen independent hierarchies [59]. Each concept is described by several characteristics such as unique identifiers, parent(s), name(s) and role(s) or semantic relation(s). It is freely available as part of the UMLS.

Cyc⁵² is a general ontology that is built around a core of more than 1 million hand-coded assertions that capture “common sense” knowledge. Groups of assertions that share a common set of assumptions (e.g.: domain, level of detail, time interval) are called *microtheories*. One of these *microtheories* is OpenCyc, the upper level and publicly available part of the Cyc with 6,000 concepts and 60,000 assertions on these concepts [96].

WordNet [19] is an electronic lexical database used by applications of natural language processing and information retrieval. It is a general English resource composed of set of synonyms (synset) with different structures for nouns, verbs, adjectives and adverbs. This resource may be used in conjunction with other concept representations like UMLS, that don't include all word-level synonyms and permutations in health concepts, providing another component to medical concept representation and retrieval. The latest version of WordNet (3.0) contains 82,115 noun synsets categorized into hierarchies that are not based on a formal ontology theory. For example, *vasoconstriction* is not formally related to the health domain. Instead it is related to *constriction* what emphasizes the physical mechanism rather than pathology. WordNet is also freely and publicly available for download.

In [96], the representation of blood in the several systems is analysed, showing the differences among them and the richness of ontologies when compared to taxonomies. Blood is a complex case because it has two different superordinates: tissue and body substance. This is highlighted by the differing representations that raise compatibilities' issues among ontologies. The richness of ontologies is also emphasized by the additional knowledge about blood through concept's properties and through the associative relations to other concepts (in taxonomies the arrangement of concepts is only made in is-a hierarchies).

4 Health Specific Search Engines

This section presents some health specific search engines with special characteristics that may be useful in this domain. Three of them are still in a beta version (Healia, MEDgle and Medstory).

Healia⁵³ is a health search engine directed to consumers developed over more than four years with an award from the National Cancer Institute. According to their creators, its main differences to other search engines are: high quality search results (through a quality index they've developed to the analysis of health content's quality), availability of “Personal Search” filters that allows the finding of more personally relevant information, use of health thesaurus to guide search, availability of two interfaces (a simpler one without filters and a standard one with filters).

MEDgle⁵⁴ is a MEDical GLobal Electronic computer generated search to general medical conditions through symptoms, diagnoses, physicians, drugs and medical procedures. In addition to the traditional text boxes of search engines, its interface is also graphic and has

⁵²<http://www.cyc.com/>

⁵³<http://www.healia.com>

⁵⁴<http://www.medgle.com>

several filter options. Its results present not only external links but also a brief summarization of web contents about the searched topic.

Fealth⁵⁵ is the *Advanced Search* product of Findica, an information and technology enterprise that develops products and services to Health Care Companies.

Medstory⁵⁶ is a Microsoft health search engine directed to consumers (Health tab) and health professionals (Research tab). It uses information from other users' searches to help the user refine and guide his search, delivering more than the standard list of web pages. It offers high-level health categorizations (e.g.: Drugs & Substances, Conditions, Procedures), different among consumers and professionals, that allow the refinement of searches. In addition, it offers another type of high-level categorizations, according to the type of resource returned (e.g.: Web, News/Media, Audio/Video, Clinical Trials, Research Articles).

Textmed⁵⁷ is a search engine to find information about several types of medical entities (e.g.: disease, drug, chemical, organism). It currently has 88,545 entities obtained from the analysis of about 15 million Medline abstracts. In each results' page it presents the juxtapositions (co-occurrence terms) for the searched term and its relationships to other entities. Finally it presents the three top articles with that entity.

5 Recent Work

Latest times have been of great developments in HIR. This can be confirmed by the attention given by major IR related companies like Google and Microsoft, by the number of recent created search engines (a large number of them is still in beta version) and by the large number of publications. Globally, research has been dedicated to find more efficient ways to adjust queries to user needs, to index health data, to present search results, to rank results (e.g. based on readability, content) and also to analyze information needs and behaviours.

This section describes the most recent research directions in HIR, organized by high-level research topics.

5.1 Health Information Seeking

With the increase of the availability of health information on the Web and the engagement of users in health information seeking in the Web, studies of information needs and behaviors in HIR proliferate. Usually, research in this area involves not only informatics but also areas like information science, psychology and medicine. These studies produce results useful to more technical studies that aim to improve all the retrieval process, to contribute to the understanding of the Internet's influence on health behaviors and, ultimately, in the health care system. Some research papers focus on the consumer, some on the health professionals and others on global aspects of the retrieval process, done either by consumers or by health professionals.

⁵⁵<http://www.fealth.com>

⁵⁶<http://www.medstory.com>

⁵⁷<http://www.textmed.com>

5.1.1 Consumers

In the consumer arena, a large number of studies are dedicated to characterize consumers' health information needs and their behavior. These studies are usually based on the application of surveys and log analysis. The Pew Internet & American Life Project⁵⁸, that explores the impact of the Internet in several aspects of society, has already published two reports that characterize online health search: one of the year 2000 [21] and other of the year 2006 [20]. Rains [83] discuss the perceptions of the Web's use to seek health information gathered from the Health Information National Trends Survey's findings. These studies are centered in the USA citizen and there are others that characterize the behaviors of others countries' citizens. Two recent examples focus on portuguese [93] and greek citizens [26].

There are also studies that analyze the use patterns of specific health information retrieval systems like one that, through the analysis of MEDLINEplus logs, studies the stability of user queries overtime [94]. Others are dedicated to explore the user's information search process (e.g.: explore student's search process and outcome in Medline to write an essay for an EBM class [38]) and to investigate the influence of technical knowledge and cognitive abilities on health information seeking [95]. Yoo's paper [109] has a more theoretical approach, attempting to explain how and why middle-aged women use health web sites based on two theories (the uses and gratifications approach from mass communication research and the theory of planned behavior from social psychology). An earlier and more comprehensive paper is the one from Cline [12] that covers aspects like: potential benefits of health information seeking on the Internet, information quality problems, criteria for evaluating online health information and design features problems (like disorganization and technical language).

5.1.2 Professionals

There are also studies that seek to understand the information needs of health professionals. Revere [86] does a literature review on these needs, trying to answer the following questions: "What are the information needs of public health professionals?", "In what ways are those needs being met?", "What are the barriers to meeting those needs?" and "What is the role of the Internet in meeting information needs?". A Gonzalez-Gonzalez's paper [22], analyzes spanish primary care physicians' information needs that arise in office practice and their information seeking patterns to satisfy these needs (in and after the consultations). Twose has a paper [104] with the same goals but a different approach that included the combination of usage statistics from a web portal (that allowed the access to a library's electronic resources), self-report and observational data collected during an offered course. Other papers are more dedicated to the study of information seeking behaviors, such as [28] where Hemminger, through a census survey, analysed the information seeking behavior of academic scientists of basic science and medical science departments; and one that determines how good is Google to lead doctors to a correct diagnosis [101].

Still directed to professionals, other lines of research include the study of: information retrieval systems' impact on professionals' performance in answering clinical questions (in [107] this is done with a pre/post intervention experimental design and several types of health

⁵⁸<http://www.pewinternet.org>

professionals), the assessment of the effectiveness of information retrieval systems for professionals (in [29] it is defined a framework that is used to assess Medline articles of evaluation studies), changes in information behavior after certain changes like the introduction of a clinical librarian service [105] and collaborative information seeking behavior in the context of medical care [85].

5.1.3 Consumers and Professionals

There are also general studies that are neither directed to consumers or professionals. Some papers are more theoretical, globally studying health information seeking behavior ([47] examines scientific literature from 1982 to 2006 on this subject to capture the concept's characteristics and discuss operationalizations, antecedents, and outcomes), others explore certain aspects in health information seeking behavior like user navigation [24], education disparities [51] and changes in user needs [1].

Some studies are more focused on web aspects. Spink's [98] report findings from an analysis of health queries to different web search engines, providing insights into health querying and suggesting implications of the use of web search engines for health information seeking. Another paper [17] studies the prevalence of health-related searches on the Web, based on the proportion of pages on the web containing the search string and the word *health*.

5.2 Indexing

Indexing in HIR is a research area than can largely benefit from the diversity and quantity of Health Concept Representations (see Section 3). In fact, a significant number of papers about this topic use at least one of these representations. In a Douyere's paper [15] the MeSH thesaurus is adapted to the broader field of health Internet resources and is used together with the Dublin Core metadata format to catalogue and index French health resources. In [33], MeSH is also combined with a well-established method for extraction of domain terms in the development of an automatic term extraction method for indexing large medical collections such as MEDLINE. Houston et al. [34] explore the use of concept spaces (automatically generated thesauri, where concepts are represented as nodes and relationships as weighted link, with associative memory that allows new paradigms for knowledge discovery and document searching) in HIR. They evaluated and compared the use of terms suggested by MeSH, UMLS MetaThesaurus and the automatic generated thesauri with document collection's terms. No statistically significant differences among the thesauri were found and there was almost no overlap of relevant terms suggested by different thesauri what suggests that recall could be significantly improved using a combined thesaurus approach.

There are also papers dedicated to the development of Health Concept Representations. Zeng is a researcher with a large work on this field, being the coordinator of the Consumer Health Vocabulary Initiative⁵⁹. One of her papers [111] describes the development of computerized methods to mark up Web content. Another paper [115] from a different author also presents research on segmenting and labeling HTML medical journal articles through

⁵⁹<http://www.consumerhealthvocab.org/>

a hidden markov model approach. In [43] it is suggested the use of social bookmarking (such as the tags from CiteULike) as an additional health concept representation and a way to discover documents not yet indexed in on-line databases.

5.3 Retrieval

5.3.1 Health Information Characteristics

One of the greatest concerns on HIR is related to the quality of published information quality and possible the interpretations given by non-experts. The analysis of health information quality (e.g. accuracy, timeliness, accessibility) and its adjustment to the user (e.g. readability) are popular lines of research in HIR. The evaluation of these two factors may be used in the results' ranking. Some of the studies in this area are described next.

In [5] it was done an accessibility, quality and readability evaluation of health information on breast cancer, depression, obesity and childhood asthma available in English and Spanish. The accessibility of search engines was assessed using a structured search experiment, the content's quality was evaluated by physicians using structured implicit review and the reading grade level was assessed using an established method (Fry Readability method). They found that coverage of key information is poor but the accuracy is generally good and that high reading levels are necessary to comprehend health contents.

Yan et al. [108] express the need of new computational models of readability to rank results in information retrieval systems as traditional readability formulas are too generalist. Facing this necessity, they propose a concept-based model of text readability that takes into account textual genres of a document and domain specific knowledge in three major readability formulas.

Miller et al. [61] also state that traditional readability formulas are not targeted to specific domains like health as they ignore the use of specialized vocabulary. In their paper they propose a naïve Bayes classifier for three levels of health terminology specificity (consumer, health learner, health professional) created with the lexicon of a medical corpus. This classifier attained an accuracy of 96% and was applied to consumer health web pages. Only 4% of pages were classified as consumer ones, while all the others were included at the professional level. Miller was the second author of a recently published paper [49] that also describes the evaluation of the naïve Bayes classifier that was compared with readability formulas and the readability assessment of an expert and a consumer. The classifier indicated that documents were at a lower level of readability difficulty than the readability formulas. A previous paper from Leroy and other authors [48] compared four types of documents: easy and difficult WebMD documents, patient blogs, and patient educational material. They found that it is possible to simplify many documents based on terminology in addition to sentence structure (but this can still be insufficient for difficult documents).

In [42] the males and females' familiarity with terms of three types of health topics (male-specific, female-specific and gender-neutral) is evaluated. It was found that males were more familiar with neutral and male-specific topics and that females has no topic effect. In face of these results, the tailoring of health readability formulas to target populations is also discussed. Rosemblat et al. [90] have done an exploratory study to analyze the relevance of readability's predictors in the consumer health domain, based on expert judgment and to

characterize expert ratings' patterns across the various predictors. They concluded that the development of health readability tools may require the modification of existing measures (e.g. including health-related vocabulary) and the addition of new predictive features.

Zeng is the first author of three papers related to health consumer terminology. The first [110] studied the characteristics of consumer terminology used in HIR through the log analysis of two consumer web sites and patients' interviews. They concluded that there are significant mismatches between consumer and information source terminologies. In [114] it was created a method to measure the familiarity of medical terms and a predictive model for familiarity, based on term occurrence in text corpora and reader's demographics. In the third paper [113] the authors developed a systematic methodology using text analysis and human review to assign *consumer-friendly* names to UMLS's concepts. The evaluation of this method was done applying a questionnaire to consumers and the results suggested this methodology is useful in the development of consumer health vocabularies.

5.3.2 Query Expansion

Is often happens that query terms are related to terms used to index documents but are not indexing terms. This motivates the development of techniques of query expansion. These methods are used to improve precision in search results with alternative/additional query terms (synonyms or other semantic relationships). This technique is quite used, probably due to the large number of Health Concept Representations available.

In another Zeng's paper [112], a tool to assist people in health-related query construction was developed. The suggested terms were selected based on their semantic distance to the original query (calculated through co-occurrences in medical literature and log data and also through semantic relations in medical vocabularies). They concluded that semantic-distance-based query recommendations can help consumers with query formulation during HIR.

In [18] it is proposed a query expansion method through text analysis of non-topical terms in Web documents. They define the concepts of topical terms (TT) as the terms that represent the subject content of documents (e.g.: breast cancer), non-topical terms (NTT) are, usually, terms that occur before or after topical terms to represent a specific aspect of the subject (e.g. 'about' in 'about breast cancer') and semi-topical terms (STT) that are terms that normally do not occur alone, being used in conjunction with topical terms to narrow or further specify the subject – they are normally domain-specific (e.g.: 'risk of' in 'risk of breast cancer'). The defined method to query expansion is based on the use of NTT and STT in conjunction with TT.

Ide et al. [39] describe the algorithms used in a search engine with query expansion and probabilistic relevancy ranking, evaluated using data and standard evaluation methods from the 2003 and 2006 TREC Genomics Track.

5.3.3 Ranking

The ranking algorithms used in a IR system, responsible of adjusting the position of each result in the returned results' list, can also be used to enhance these systems (more specifically, the precision at first returned results). Price is a co-author with, at least, two publications

about approaches to improve ranking in HIR. In [80] the authors present a system that ranks results according to the likely quality of page health contents. In [79] the approach to improve the results' ranking is different. The authors model queries as relationships between concepts and try to match these relations with the ones existing between documents. The analysis of users' browsing behavior is also a typical approach to the development of ranking algorithms (e.g.: [3]) .

5.3.4 IR Models

The definition of IR models is another topic seen in HIR papers. The use of semantic information is typical in the most recent proposed models. This is the case of the model proposed by Price et al. in [81] and [82]. In this model, the authors describe the content of documents in domain-specific collections using semantic components ("segments of text about a particular aspect of the main topic of the document that may not correspond to structural elements in the document" [82]), complementary to full text and keyword indexing. In the first paper [81], the authors introduce the model, present the results of its application to the representation of clinical questions in the medical domain and present ways to use the model for retrieval. In the second paper [82], the authors present experimental evidence that the model enhances the retrieval of domain-specific documents in response to real users' realistic queries.

In [32] an IR model, implemented in MedSearch (a Medline retrieval system), that discovers similarities between documents containing semantically similar but not necessarily lexically similar terms is presented. In [16] it is presented a methodology to build and enhance an ontology in health domain through semantic elements extraction. The extracted information from Web documents is then summarized, indexed and stored in the database for an implemented information retrieval system.

A very recent paper [37] presents an information seeking model to represent human search expertise that may allow the development of an intelligent search agent that generates adaptive search strategies based on the human search expertise. The model described is hierarchical and multi-level where each level represents a problem space traversed during the search process and a layer of knowledge required to a successful search.

5.4 Evaluation

The evaluation of an IR system isn't a simple task, involving frequently human intervention (in the judgement of relevance or in users' studies). Another approach involves the use test collections (set of documents and questions) and less human intervention. An example are the ones used in TREC⁶⁰. The development of such collections is the target of some papers. For example, in [31], Hersh et al. have developed a test collection to assess visual and textual methods in biomedical image retrieval.

An example of a user centered study is the one presented in [46] where a comparative usability evaluation of an automated text summarization system and three search engines is done. The evaluation involved audio and video recording of subject interactions with the

⁶⁰<http://trec.nist.gov>

interfaces. Another paper in which evaluation is done with human intervention is one from Tang et al. [102]. In this paper, human assessors were used to do the relevance judging according to a scheme preciously developed. The goal was to compare the performance of domain-specific health and depression search engines with Google on both relevance of results and quality of advice.

5.5 User Interfaces and Visualization

The development of user interfaces and aspects of information visualization in IR systems is another area with a significant quantity of research work. In the health area, with developments in concept representation models, it also became popular the use of this information to improve the way results are presented to the user.

In [100] is presented a system that implements a concept-based visualization of the results that, according to a user study, is less suitable for searching specific information and more suitable to the exploration of mostly unknown data. In [15], a paper already cited in this report, a terminology is developed from the MeSH thesaurus and metadata elements and is used in several tasks, one of them being the visualization and navigation through the concept hierarchies. Stapley et al. [99] have built a system for retrieving and visualizing co-occurrences of gene names in Medline abstracts. From the co-occurrence data is built a graph where nodes are genes and edge lengths are a function of the co-occurrence of the two genes in the literature.

6 People

The most notorious person in HIR is William Hersh⁶¹, M.D. and Professor and Chair of the Department of Medical Informatics & Clinical Epidemiology in the School of Medicine at Oregon Health & Science University. His research focuses on the development and evaluation of IR systems for health practitioners and researchers. He has a several publications in this field, is the author of a book with an Information Retrieval's health perspective [30], the author of Health Informatics' book chapters (like [96, 10]) and the editor of other Health Informatics' books (like [10]).

Another person with a significant research activity in HIR is Qing Treitler Zeng⁶², an Assistant Professor of the Decision Systems Group in Harvard Medical School and a Research Associate at the Brigham and Women's Hospital. One of her research interests is the semantic knowledge-based information retrieval and presentation. One of her latest research projects, in which she assumes the role of leader, is the Consumer health vocabulary⁶³ which aims to build, in an open source and collaborative way, a vocabulary that links everyday words about health with technical terms used by professionals. One of this project's applications is in the area of Information Retrieval as stated in [112], one of her many publications in the HIR field.

⁶¹<http://medir.ohsu.edu/hersh/>

⁶²<http://dsg.harvard.edu/qzeng/>

⁶³<http://www.consumerhealthvocab.org/>

Susan L. Price (no personal webpage was found) is thought to be a PhD Student at the Department of Computer Science of the Portland State University, supervised by Professor Lois Delcambre. She has been working on the use of semantic components to enhance retrieval of domain-specific information (being health the domain selected as a case study) and is the author of papers with Lois Delcambre, Marianne Lykke Nielsen of the Royal School of Library and Information Science in Denmark and William Hersh.

Other authors with some publications in the area are David Howard Hickman⁶⁴ with jointed works with William Hersh; Angelos Hliaoutakis⁶⁵, a MSc student in Computer Engineering that has been publishing in the HIR area; Thanh Tin Tang⁶⁶ and Madhu C. Reddy⁶⁷ which has interest in Collaborative information behavior and has been developing research work in the area of healthcare.

There are also authors with a less straight connection to the specific area of HIR and more connected to the broader fields of eHealth and Health Informatics. However, having a straight connection to the health area and published work in HIR, they will be mentioned next. Gunther Eysenbach⁶⁸ is an Associate Professor at the Department of Health Policy, Management and Evaluation at the University of Toronto and one of the editors of [96]. Susannah Fox⁶⁹ is a Pew Internet & American Life Project associate director connected to the Internet's impact on health care with several published reports⁷⁰. Steve Rains⁷¹ has also some published papers in the area of Health Information Seeking Behavior. Edward H. Shortliffe⁷² is a well known name in the broad area of health informatics with specific interests in decision-support systems, integrated workstations for clinicians and web-based information dissemination.

Three last authors worth mentioning, more connected to the development of ontologies in the health area, are James J. Cimino⁷³ that is one of the editors of [96], Mark Musen⁷⁴ and Alan Rector⁷⁵.

7 Research Groups

Only two research groups on the specific field of HIR were found and the second one is related to the fields of health and IR but is not directly related to HIR. The first is the

⁶⁴http://www1.va.gov/pshsrd/docs/CV_Hickam_02.htm and <http://www.bio-computing.org/showauthor.php?surname=Hickam&initials=DH>

⁶⁵<http://www.softnet.tuc.gr/angelos/>

⁶⁶http://www.informatik.uni-trier.de/~ley/db/indices/a-tree/t/Tang:Thanh_Tin.html and <http://libra.msra.cn/authordetail.aspx?id=1316585&query=web+search>

⁶⁷<http://faculty.ist.psu.edu/reddy/Research.htm>

⁶⁸<http://gunther-eyenbach.blogspot.com/>

⁶⁹http://www.pewinternet.org/PPF/a/104/about_staffer.asp

⁷⁰<http://www.pewinternet.org/PPF/c/5/topics.asp>

⁷¹http://datamonster.sbs.arizona.edu/communication/faculty/each_detail.php?option=1&detail=57&mtitle=Core%20Faculty

⁷²<http://www.dbmi.columbia.edu/shortliffe>

⁷³<http://www.dbmi.columbia.edu/ciminoj/>

⁷⁴http://med.stanford.edu/profiles/Mark_Musen/

⁷⁵<http://www.cs.manchester.ac.uk/research/publications/byauthor/Rector/>

Centre for Health Information Management Research⁷⁶, at the University of Sheffield, more connected to the study of Health Information Needs and Behaviors. The second one is the W3C Semantic Web Health Care and Life Sciences Interest Group⁷⁷ which aims to develop, promote and support the use of Semantic Web technologies in health care and life sciences.

8 Research Projects

The two most active persons in the area have four research projects mentioned next. William Hersh has three projects: the TREC Genomics Track retrieval of scientific literature on genomics⁷⁸, ImageCLEFmed⁷⁹ that is part of the Cross Language Evaluation Forum⁸⁰ (CLEF) for medical image retrieval and OHSUMED⁸¹, a test collection of a subset MEDLINE references, created to assist in information retrieval research.

The fourth project is led by Qing Zeng and is entitled Consumer Health Vocabularies⁸², a project dedicated to the development of vocabularies that link everyday words about health to professional terms used by health care professionals.

9 Main Organizations

Some of the main organizations in HIR, some more connected to the health area than others, are listed below.

- American Library Association – <http://www.ala.org>
- American Medical Informatics Association – <http://www.amia.org>
- American Society for Information Science and Technology – <http://www.asis.org>
- American Society for Indexing – <http://www.asindexing.org>
- Association for Computer Machinery, Special Interest Group on Information Retrieval – <http://www.sigir.org>
- Medical Library Association – <http://www.mlanet.org>
- National Library of Medicine – <http://www.nlm.nih.gov>
- Special Libraries Association – <http://www.sla.org>

⁷⁶<http://www.shef.ac.uk/chimr/research/>

⁷⁷<http://www.w3.org/2001/sw/hcls/>

⁷⁸<http://ir.ohsu.edu/genomics/>

⁷⁹<http://ir.ohsu.edu/image/>

⁸⁰<http://www.clef-campaign.org/>

⁸¹<http://ir.ohsu.edu/ohsumed/>

⁸²<http://www.consumerhealthvocab.org/>

10 Main Journals

Some of the main journals where HIR's papers can be published are listed below. Some of them are more related to the Computer Science field, some to the Information Science field and others to the health and health informatics field.

- ACM Transactions on Information Systems – <http://www.acm.org/tois>
- Artificial Intelligence in Medicine – http://www.elsevier.com/wps/find/journaldescription.cws_home/505627/description#description
- British Medical Journal – <http://www.bmj.com>
- Computers in Biology and Medicine – <http://www.elsevier.com/locate/inca/351>
- Health Information and Libraries Journal – <http://www.blackwellpublishing.com/journal.asp?ref=1471-1834>
- Journal of Biomedical Informatics – <http://www.academicpress.com/jbi>
- Journal of the American Medical Informatics Association – <http://www.jamia.org>
- Journal of the American Society for Information Science and Technology – <http://www.asis.org/jasist.html>
- Journal of the Medical Library Association – <http://www.mlanet.org/publications/jmla/>
- Informatics for Health and Social Care – <http://www.tandf.co.uk/journals/tf/14639238.html>
- Information Processing and Management – <http://www.elsevier.com/locate/infoproman>
- Information Retrieval – <http://www.springer.com/computer/database+management+%26+information+retrieval/journal/10791>
- International Journal of Medical Informatics – http://www.elsevier.com/wps/find/journaldescription.cws_home/506040/description#description
- Medical & Biological Engineering & Computing – <http://www.springer.com/engineering/biomedical+eng/journal/11517>
- Medical Decision Making – <http://mdm.sagepub.com/>
- Methods of Information in Medicine – <http://www.schattauer.de/index.php?id=704>

References

- [1] Anne Adams and Ann Blandford. Digital libraries' support for the user's 'information journey'. In *JCDL '05: Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries*, pages 160–169, New York, NY, USA, 2005. ACM Press.
- [2] James Allan, Jay Aslam, Nicholas Belkin, Chris Buckley, Jamie Callan, Bruce Croft, Sue Dumais, Norbert Fuhr, Donna Harman, David J. Harper, Djoerd Hiemstra, Thomas Hofmann, Eduard Hovy, Wessel Kraaij, John Lafferty, Victor Lavrenko, David Lewis, Liz Liddy, R. Manmatha, Andrew Mccallum, Jay Ponte, John Prager, Dragomir Radev, Philip Resnik, Stephen Robertson, Roni Rosenfeld, Salim Roukos, Mark Sanderson, Rich Schwartz, Amit Singhal, Alan Smeaton, Howard Turtle, Ellen Voorhees, Ralph Weischedel, Jinxi Xu, and Chengxiang Zhai. Challenges in information retrieval and language modeling. *SIGIR Forum*, 37(1):31–47, 2003.
- [3] Ioannis Anagnostopoulos and Ilias Maglogiannis. Monitoring browsing behaviour and search services evolution adaptation with a capture-recapture Internet-based programming technique: A case-study over medical portals. *Information Services and Use*, 27(3):105–122, 2007.
- [4] Ricardo Baeza-Yates and Berthier Ribeiro-Neto. *Modern Information Retrieval*. Addison Wesley, May 1999.
- [5] G. K. Berland, M. N. Elliott, L. S. Morales, J. I. Algazy, R. L. Kravitz, M. S. Broder, D. E. Kanouse, J. A. Muñoz, J. A. Puyol, M. Lara, K. E. Watkins, H. Yang, and E. A. McGlynn. Health information on the internet: accessibility, quality, and readability in english and spanish. *JAMA : the journal of the American Medical Association*, 285(20):2612–2621, 2001.
- [6] Adam Bosworth. Putting health into the patient's hands - consumerism and health care. Opening Plenary Session and Keynote Address, May 2007.
- [7] S. H. Brenner and E. J. Mckinin. CINAHL and MEDLINE: a comparison of indexing practices. *Bulletin of the Medical Library Association*, 77(4):366–371, October 1989.
- [8] Vannevar Bush. As We May Think. *The Atlantic Monthly*, 1(176):101–108, July 1945.
- [9] Jamie Callan, James Allan, Charles L. A. Clarke, Susan Dumais, David A. Evans, Mark Sanderson, and Chengxiang Zhai. Meeting of the MINDS: an information retrieval research agenda. *SIGIR Forum*, 41(2):25–34, December 2007.
- [10] H. Chen, Sherrylynn S. Fuller, and Charles P. Friedman. *Medical Informatics: Knowledge Management and Data Mining in Biomedicine (Integrated Series in Information Systems)*. Springer, June 2005.
- [11] Cyril Cleverdon. The Cranfield tests on index language devices. *Aslib Proceedings*, 19(6):173–194, 1967.

- [12] R. J. W. Cline and K. M. Haynes. Consumer health information seeking on the Internet: the state of the art. *Health Educ. Res.*, 16(6):671–692, December 2001.
- [13] M. H. Coletti and H. L. Bleich. Medical subject headings used to search the biomedical literature. *Journal of the American Medical Informatics Association : JAMIA*, 8(4):317–323, 2001.
- [14] Bruce W. Croft and David J. Harper. Using probabilistic models of document retrieval without relevance information. *Journal of Documentation*, 35(4):285–295, 1979.
- [15] Magaly Douyere, Lina F. Soualmia, Aurelie Neveol, Alexandrina Rogozan, Badisse Dahamna, Jean-Philippe Leroy, Benoit Thirion, and Stefan J. Darmoni. Enhancing the MeSH thesaurus to retrieve French online health resources in a quality-controlled gateway. *Health Information and Libraries Journal*, 21(4):253–261, 2004.
- [16] Tran Dung and Wataru Kameyama. Ontology-based information extraction and information retrieval in health care domain. *Data Warehousing and Knowledge Discovery*, pages 323–333, 2007.
- [17] G. Eysenbach and Kohler. What is the prevalence of health-related searches on the World Wide Web? Qualitative and quantitative analysis of search engine queries on the internet. *AMIA Annu Symp Proc*, pages 225–229, 2003.
- [18] Rahmatollah Fattahi, Concepción S. Wilson, and Fletcher Cole. An alternative approach to natural language query expansion in search engines: Text analysis of non-topical terms in web documents. *Information Processing & Management*, 44(4):1503–1516, July 2008.
- [19] Fellbaum. *WordNet: An Electronic Lexical Database (Language, Speech, and Communication)*. The MIT Press, May 1998.
- [20] S. Fox. Online health search 2006. Technical report, Pew Internet & American Life Project, 2006.
- [21] S. Fox and L. Rainie. The online health care revolution: How the Web helps Americans take better care of themselves. Technical report, The Pew Internet & American Life Project, 2000.
- [22] A. I. González-González, M. Dawes, J. Sánchez-Mateos, R. Riesgo-Fuertes, E. Escortell-Mayor, T. Sanz-Cuesta, and T. Hernández-Fernández. Information needs and information-seeking behavior of primary care physicians. *Annals of family medicine*, 5(4):345–352, 2007.
- [23] Google. Google health. Available from: <https://www.google.com/health> [cited 2008-07-11].
- [24] L. Graham, T. Tse, and A. Keselman. Exploring user navigation during online health information seeking. In *AMIA Annu Symp Proc. 2006*, pages 299–303, Lister Hill Center, National Library of Medicine, NIH, DHHS, Bethesda, MD, USA., 2006.

- [25] Tom Gruber. Ontology. Available from: <http://tomgruber.org/writing/ontology-definition-2007.htm> [cited 2008-07-10].
- [26] Daphne Halkias, Nicholas Harkiolakis, Paul Thurman, and Sylva Caracatsanis. Internet usage for health-related purposes among greek consumers. In *ICCOMP'07: Proceedings of the 11th WSEAS International Conference on Computers*, pages 281–289, Stevens Point, Wisconsin, USA, 2007. World Scientific and Engineering Academy and Society (WSEAS).
- [27] R. Haux, E. Ammenwerth, W. Herzog, and P. Knaup. Health care in the information society - A prognosis for the year 2013. *International journal of medical informatics*, 66(1-3):3–21, November 2002.
- [28] Bradley M. Hemminger, Dihui Lu, K. T. L. Vaughan, and Stephanie J. Adams. Information seeking behavior of academic scientists. *J. Am. Soc. Inf. Sci. Technol.*, 58(14):2205–2225, December 2007.
- [29] W. R. Hersh and D. H. Hickam. How well do physicians use electronic information retrieval systems? a framework for investigation and systematic review. *JAMA*, 280(15):1347–1352, October 1998.
- [30] William R. Hersh. *Information Retrieval*. Springer, December 2002.
- [31] William R. Hersh, Henning Muller, Jeffery R. Jensen, Jianji Yang, Paul N. Gorman, and Patrick Ruch. Advancing biomedical image retrieval: Development and analysis of a test collection. *J Am Med Inform Assoc*, 13(5):488–496, September 2006.
- [32] Angelos Hliaoutakis, Giannis Varelas, Euripides Petrakis, and Evangelos Milios. Med-Search: A retrieval system for medical information based on semantic similarity. *Research and Advanced Technology for Digital Libraries*, 4172/2006:512–515, 2006.
- [33] Angelos Hliaoutakis, Kalliopi Zervanou, Euripides G. M. Petrakis, and Evangelos E. Milios. Automatic document indexing in large medical collections. In *HIKM '06: Proceedings of the international workshop on Healthcare information and knowledge management*, pages 1–8, New York, NY, USA, 2006. ACM Press.
- [34] Andrea L. Houston, Hsinchun Chen, Bruce R. Schatz, Susan M. Hubbard, Robin R. Sewell, and Tobun D. Ng. Exploring the use of concept spaces to improve medical information retrieval. *Decis. Support Syst.*, 30(2):171–186, December 2000.
- [35] B. L. Humphreys, D. A. Lindberg, H. M. Schoolman, and G. O. Barnett. The Unified Medical Language System: an informatics research collaboration. *J Am Med Inform Assoc*, 5(1):1–11, 1998.
- [36] Betsy L. Humphreys and Peri L. Schuyler. The Unified Medical Language System: moving beyond the vocabulary of bibliographic retrieval. *High performance medical libraries: Advances in information management for the virtual era*, pages 31–44, 1993.

- [37] Peter W. Hung, Stephen B. Johnson, David R. Kaufman, and Eneida A. Mendonca. A multi-level model of information seeking in the clinical domain. *Journal of Biomedical Informatics*, 41(2):357–370, April 2008.
- [38] S. Huuskonen and P. Vakkari. Students’ search process and outcome in medline in writing an essay for a class on evidence-based medicine. *Journal of Documentation*, 64(2):287–303, 2008.
- [39] N. C. Ide, R. F. Loane, and D. Demner-Fushman. Essie: a concept-based search engine for structured biomedical text. *J Am Med Inform Assoc*, 14(3):253–263, 2007.
- [40] Karen S. Jones. A statistical interpretation of term specificity and its application in retrieval. *Journal of Documentation*, 28(1):11–21, 1972.
- [41] R. A. Jydstrup and M. J. Gross. Cost of information handling in hospitals. *Health services research*, 1(3):235–271, 1966.
- [42] Alla Keselman, Lisa Massengale, Long Ngo, Allen Browne, and Qing Zeng. The effect of user factors on consumer familiarity with health terms: Using gender as a proxy for background knowledge about gender-specific illnesses. *Biological and Medical Data Analysis*, pages 472–481, 2006.
- [43] Margaret E. I. Kipp. Tagging for health information organisation and retrieval. In *JCDL '07: Proceedings of the 2007 conference on Digital libraries*, page 485, New York, NY, USA, 2007. ACM.
- [44] Rachel Kleinsorge and Jan Willis. Unified Medical Language System (UMLS) basics. Presentation, September 2007.
- [45] Kotinos. Libraries - history of the private, royal, imperial, monastic and public libraries. Available from: <http://www.libraries.gr/> [cited 2008-07-10].
- [46] A. W. Kushniruk, M. Y. Kan, K. Mckeown, J. Klavans, D. Jordan, M. Laflamme, and V. L. Patel. Usability evaluation of an experimental text summarization system and three search engines: implications for the reengineering of health care interfaces. In *AMIA Annu Symp Proc. 2002*, pages 420–424, Department of Mathematics and Statistics, York University, Toronto, Ontario M3J 1P3, Canada., 2002.
- [47] S. D. Lambert and C. G. Loiselle. Health information seeking behavior. *Qualitative health research*, 17(8):1006–1019, October 2007.
- [48] G. Leroy, E. Eryilmaz, and B. T. Laroya. Health information text characteristics. In *AMIA Annu Symp Proc. 2006*, pages 479–483, Claremont Graduate University, Claremont, California, USA., 2006.
- [49] G. Leroy, T. Miller, G. Rosemblat, and A. Browne. A balanced approach to health information evaluation: A vocabulary-based naïve bayes classifier and readability formulas. *Journal of the American Society for Information Science and Technology*, 59(9), July 2008.

- [50] D. A. Lindberg, B. L. Humphreys, and A. T. Mccray. The Unified Medical Language System. *Methods of information in medicine*, 32(4):281–291, August 1993.
- [51] D. Lorence and H. Park. Study of education disparities and health information seeking behavior. *Cyberpsychology & behavior : the impact of the Internet, multimedia and virtual reality on behavior and society*, 10(1):149–151, February 2007.
- [52] H. P. Luhn. A statistical approach to mechanized encoding and searching of literary information. *IBM Journal of Research and Development*, 1(4):309–317, October 1957.
- [53] Peter Lyman and Hal R. Varian. How much information. Available from: <http://www.sims.berkeley.edu/how-much-info-2003> [cited 2008-07-10], 2003.
- [54] J. J. Mamlin and D. H. Baker. Combined time-motion and work sampling study in a general medicine clinic. *Medical care*, 11(5):449–456, 1973.
- [55] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. *Introduction to Information Retrieval*. Cambridge University Press, July 2008.
- [56] A. T. Mccray. An upper-level ontology for the biomedical domain. *Comparative and Functional Genomics*, pages 80–84, 2003.
- [57] A. T. McCray, S. Srinivasan, and A. C. Browne. Lexical methods for managing variation in biomedical terminologies. *Proc Annu Symp Comput Appl Med Care*, pages 235–239, 1994.
- [58] C. J. Mcdonald, S. M. Huff, J. G. Suico, G. Hill, D. Leavelle, R. Aller, A. Forrey, K. Mercer, G. Demoor, J. Hook, W. Williams, J. Case, and P. Maloney. LOINC, a universal standard for identifying laboratory observations: a 5-year update. *Clinical chemistry*, 49(4):624–633, April 2003.
- [59] Genevieve B. Melton, Simon Parsons, Frances P. Morrison, Adam S. Rothschild, Marianthi Markatou, and George Hripcsak. Inter-patient distance metrics using SNOMED CT defining relationships. *J. of Biomedical Informatics*, 39(6):697–705, December 2006.
- [60] Microsoft. Microsoft HealthVault. Available from: <http://www.healthvault.com/> [cited 2008-07-10].
- [61] Trudi Miller, Gondy Leroy, Samir Chatterjee, Jie Fan, and Brian Thoms. A classifier to evaluate language specificity of medical documents. In *HICSS '07: Proceedings of the 40th Annual Hawaii International Conference on System Sciences*, Washington, DC, USA, 2007. IEEE Computer Society.
- [62] Calvin E. Mooers. Coding, information retrieval, and the rapid selector. *American Documentation*, 1(4):225–229, 1950.
- [63] M. A. Musen. Medical informatics: searching for underlying components. *Methods of information in medicine*, 41(1):12–19, 2002.

- [64] NHGRI. The human genome project. Available from: <http://www.genome.gov/10001772> [cited 2008-07-10].
- [65] NIST. Text retrieval conference (trec). Available from: <http://trec.nist.gov/> [cited 2008-07-10].
- [66] NLM. 2008AA documentation - semantic network. Available from: <http://www.nlm.nih.gov/research/umls/meta3.html> [cited 2008-07-10].
- [67] NLM. 2008AA documentation - specialist lexicon and lexical tools. Available from: <http://www.nlm.nih.gov/research/umls/meta4.html> [cited 2008-07-10].
- [68] NLM. The medical subject headings vocabulary. Available from: <http://www.nlm.nih.gov/pubs/factsheets/mesh.html> [cited 2008-07-10].
- [69] NLM. MedLine. Available from: http://www.nlm.nih.gov/databases/databases_medline.html [cited 2008-07-10].
- [70] NLM. MEDLINE fact sheet. Available from: <http://www.nlm.nih.gov/pubs/factsheets/medline.html> [cited 2008-07-10].
- [71] NLM. Medline plus. Available from: <http://medlineplus.gov> [cited 2008-07-10].
- [72] NLM. MeSH tree structures. Available from: http://www.nlm.nih.gov/mesh/intro_trees2007.html [cited 2008-07-10].
- [73] NLM. NLM databases & electronic resources. Available from: <http://www.nlm.nih.gov/databases/index.html> [cited 2008-07-10].
- [74] NLM. Principles of MEDLINE subject indexing. Available from: <http://www.nlm.nih.gov/bsd/disted/mesh/indexprinc.html> [cited 2008-07-10].
- [75] NLM. Relationships in medical subject headings (MeSH). Available from: <http://www.nlm.nih.gov/mesh/meshrels.html> [cited 2008-07-10].
- [76] NLM. UMLS metathesaurus fact sheet. Available from: <http://www.nlm.nih.gov/pubs/factsheets/umlsmeta.html> [cited 2008-07-10].
- [77] NLM. Use of MeSH in online retrieval. Available from: http://www.nlm.nih.gov/mesh/intro_retrieval2008.html [cited 2008-07-10].
- [78] P. Pellegrin. *Aristotle's Classification of Animals: Biology and the Conceptual Unity of the Aristotelian Corpus*. Berkeley: University of California Press., 1986.
- [79] S. L. Price and L. M. Delcambre. Using concept relations to improve ranking in information retrieval. In *AMIA Annu Symp Proc. 2005*, pages 619–623, Department of Medical Informatics and Clinical Epidemiology, Oregon Health & Science University, USA., 2005.

- [80] S. L. Price and W. R. Hersh. Filtering web pages for quality indicators: an empirical approach to finding high quality consumer health information on the World Wide Web. In *AMIA Annu Symp Proc. 1999*, pages 911–915, Division of Medical Informatics and Outcomes Research, Oregon Health Sciences University, Portland, USA., 1999.
- [81] Susan L. Price, Lois M. Delcambre, and Marianne L. Nielsen. Using semantic components to express clinical questions against document collections. In *HIKM '06: Proceedings of the international workshop on Healthcare information and knowledge management*, pages 9–16, New York, NY, USA, 2006. ACM Press.
- [82] Susan L. Price, Marianne L. Nielsen, Lois M. L. Delcambre, and Peter Vedsted. Semantic components enhance retrieval of domain-specific documents. In *CIKM '07: Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, pages 429–438, New York, NY, USA, 2007. ACM.
- [83] Stephen A. Rains. Perceptions of traditional information sources and use of the World Wide Web to seek health information: Findings from the health information national trends survey. *Journal of Health Communication*, 12(7):667–680, 2007.
- [84] A. L. Rector and W. A. Nowlan. The GALEN project. *Computer methods and programs in biomedicine*, 45(1-2):75–78, October 1994.
- [85] Madhu C. Reddy and Patricia R. Spence. Collaborative information seeking: A field study of a multidisciplinary patient care team. *Inf. Process. Manage.*, 44(1):242–255, January 2008.
- [86] Debra Revere, Anne M. Turner, Ann Madhavan, Neil Rambo, Paul F. Bugni, Annmarie Kimball, and Sherrilynne S. Fuller. Understanding the information needs of public health practitioners: A literature review to inform design of an interactive digital knowledge management system. *Journal of Biomedical Informatics*, 40(4):410–421, August 2007.
- [87] Revolution Health Group LLC. Revolutionhealth. Available from: <http://www.revolutionhealth.com/> [cited 2008-07-10].
- [88] C. Rijsbergen. *Information Retrieval*. Butterworths, London, 1979.
- [89] S. E. Robertson and Sparck K. Jones. Relevance weighting of search terms. *Journal of the American Society for Information Science*, 27(3):129–146, 1976.
- [90] G. Roseblat, R. Logan, T. Tse, and L. Graham. Test features and readability: Expert evaluation of consumer health text. In *MEDNET*, 2006.
- [91] G. Salton. *Automatic Information Organization and Retrieval*. McGraw-Hill, New York, NY, USA, 1968.
- [92] G. Salton. *The SMART Retrieval System—Experiments in Automatic Document Processing*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1971.

- [93] Silvina Santana and Sousa A. Pereira. Da utilização da internet para questões de saúde e doença em Portugal. *Acta Med Port*, 20:47–57, 2007.
- [94] A. Scott-Wright, J. Crowell, Q. Zeng, D. Bates, and R. Greenes. Analysis of information needs of users of MEDLINEplus, 2002 - 2003. In *AMIA Annu Symp Proc. 2006*, pages 699–703, Decision Systems Group, Brigham & Women’s Hospital, Harvard Medical School, Boston, MA, USA., 2006.
- [95] Joseph Sharit, Mario A. Hernández, Sara J. Czaja, and Peter Pirolli. Investigating the roles of knowledge and cognitive abilities in older adult information seeking on the web. *ACM Trans. Comput.-Hum. Interact.*, 15(1):1–25, May 2008.
- [96] Edward H. Shortliffe, Leslie E. Perreault, Gio Wiederhold, and Lawrence M. Fagan, editors. *Medical Informatics: Computer Applications in Health Care and Biomedicine (Health Informatics)*. Springer, April 2003.
- [97] B. Smith and C. Rosse. The role of foundational relations in the alignment of biomedical ontologies. *Medinfo*, 11(Pt 1):444–448, 2004.
- [98] 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 and libraries journal*, 21(1):44–51, March 2004.
- [99] B. J. Stapley and G. Benoit. Biobibliometrics: information retrieval and visualization from co-occurrences of gene names in medline abstracts. *Pacific Symposium on Biocomputing. Pacific Symposium on Biocomputing*, pages 529–540, 2000.
- [100] H. Stuckenschmidt, F. van Harmelen, A. de Waard, T. Scerri, R. Bhogal, J. van Buel, I. Crowlesmith, C. Fluit, A. Kampman, J. Broekstra, and E. van Mulligent. Exploring large document repositories with RDF technology: The DOPE project. *IEEE Intelligent Systems*, 19(3):34–40, May 2004.
- [101] H. Tang and J. H. Ng. Googling for a diagnosis—use of google as a diagnostic aid: internet based study. *BMJ*, 333(7579):1143–1145, December 2006.
- [102] Thanh T. Tang, Nick Craswell, David Hawking, Kathy Griffiths, and Helen Christensen. Quality and relevance of domain-specific search: A case study in mental health. *Information Retrieval*, 9(2):207–225, March 2006.
- [103] Howard Turtle and Bruce B. Croft. Evaluation of an inference network-based retrieval model. *ACM Trans. Inf. Syst.*, 9(3):187–222, July 1991.
- [104] Claire Twose, Patricia Swartz, Edward Bunker, Nancy K. Roderer, and Kathleen B. Oliver. Public health practitioners information access and use patterns in the Maryland (USA) public health departments of Anne Arundel and Wicomico Counties. *Health Information and Libraries Journal*, 25(1):13–22, March 2008.

- [105] C. Urquhart, J. Turner, J. Durbin, and J. Ryan. Changes in information behavior in clinical teams after introduction of a clinical librarian service. *Journal of the Medical Library Association : JMLA*, 95(1):14–22, January 2007.
- [106] Frank Webster. *Theories of the Information Society (The International Library of Sociology)*. Routledge, May 2002.
- [107] J. I. Westbrook, E. W. Coiera, and A. S. Gosling. Do online information retrieval systems help experienced clinicians answer clinical questions? *J Am Med Inform Assoc*, 12(3):315–321, 2005.
- [108] Xin Yan, Dawei Song, and Xue Li. Concept-based document readability in domain specific information retrieval. In *CIKM '06: Proceedings of the 15th ACM international conference on Information and knowledge management*, pages 540–549, New York, NY, USA, 2006. ACM.
- [109] Eun-Young Yoo and Louise S. Robbins. Understanding middle-aged women’s health information seeking on the web: A theoretical approach. *J. Am. Soc. Inf. Sci. Technol.*, 59(4):577–590, February 2008.
- [110] Q. Zeng, S. Kogan, N. Ash, R. A. Greenes, and A. A. Boxwala. Characteristics of consumer terminology for health information retrieval. *Methods Inf Med*, 41(4):289–298, 2002.
- [111] Q. T. Zeng and J. Crowell. Semantic classification of consumer health content. In *MedNet*, 2006.
- [112] Q. T. Zeng, J. Crowell, R. M. Plovnick, E. Kim, L. Ngo, and E. Dibble. Assisting consumer health information retrieval with query recommendations. *Journal of the American Medical Informatics Association : JAMIA*, 13(1):80–90, 2006.
- [113] Q. T. Zeng, T. Tse, J. Crowell, G. Divita, L. Roth, and A. C. Browne. Identifying consumer-friendly display (CFD) names for health concepts. In *AMIA Annu Symp Proc. 2005*, pages 859–863, DSG, Brigham and Women’s Hospital, Harvard Medical School, Boston, MA, USA., 2005.
- [114] Qing Zeng, Eunjung Kim, Jon Crowell, and Tony Tse. A text corpora-based estimation of the familiarity of health terminology. *Biological and Medical Data Analysis*, pages 184–192, 2005.
- [115] Jie Zou, Daniel Le, and George R. Thoma. Structure and content analysis for HTML medical articles: a hidden markov model approach. In *DocEng '07: Proceedings of the 2007 ACM symposium on Document engineering*, pages 199–201, New York, NY, USA, 2007. ACM.