Faculty of Science and Technology – Department of Computer Science
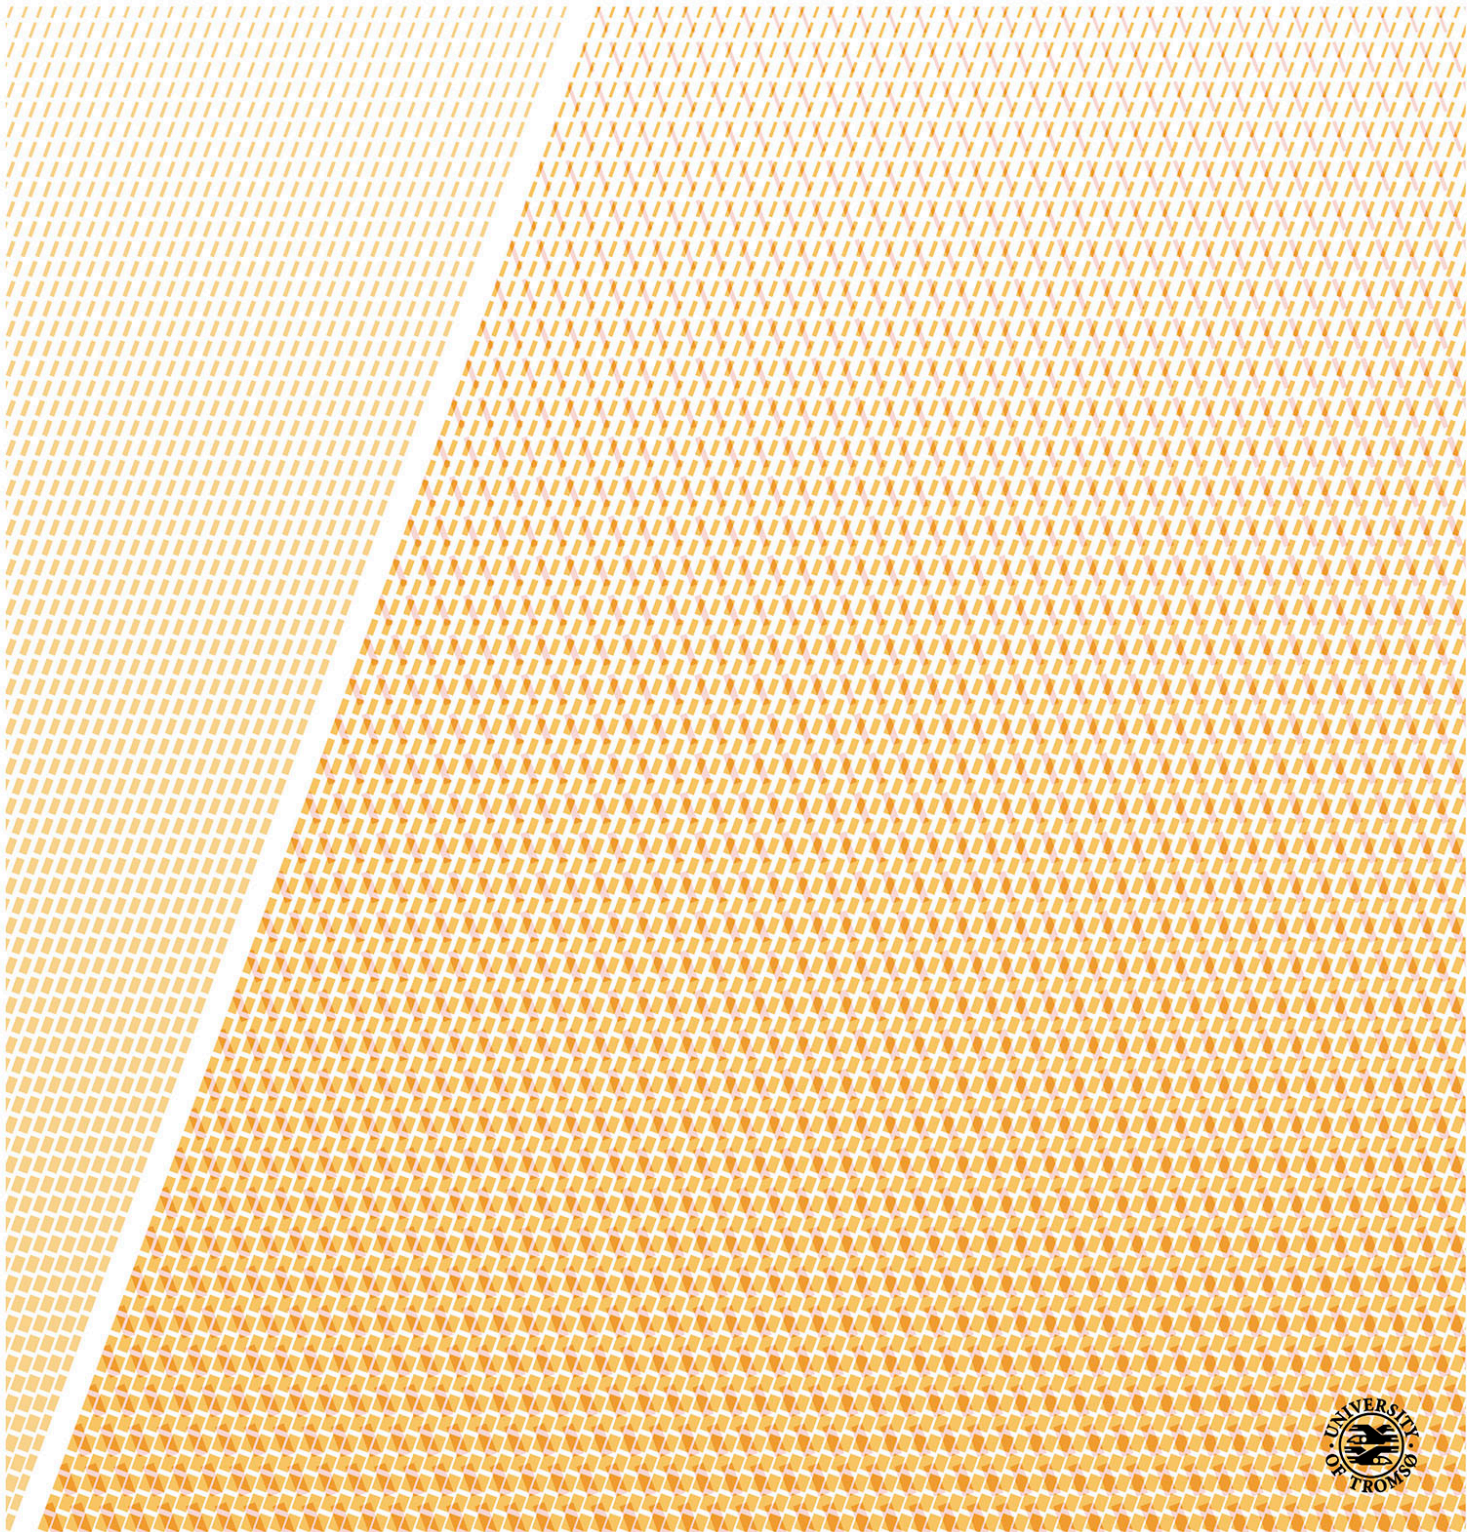
HealthTrust: trust-based retrieval of health social media videos

—

Luis Fernandez-Luque
A dissertation for the degree of Philosophiae Doctor – June 2014

Dedicated to my family, friends, supervisors and colleagues. They made it possible.

Abstract

There is a global trend towards the use of the Internet to search for information about health issues. We have access to a wide range of online health information; especially the so-called social media (e.g. blogs, videos). However, finding good quality resources is not easy in the current context of information overload. Today, very relevant and valuable health social media has to compete in visibility with misleading information such as anti-vaccination and pro-anorexia content. General web information retrieval approaches, such as Google, tend to retrieve popular content that can be misleading or even repulsive. For example, people searching for videos about diabetes foot care will discover that the top videos retrieved by a YouTube search include macabre amputations. Traditional health information retrieval approaches based on quality labels face many scalability challenges. The PhD project described here focuses on the unmet need for better technical solutions for the retrieval of high quality and relevant health social media.

This thesis summarizes nearly six years' work in the field of health social media summarized in ten research papers. I have applied a wide range of research methods such as qualitative research with patients, web-data analysis and literature reviews. My first research challenge was to grasp some understanding of the emerging health social media avalanche where research literature was virtually nonexistent. Secondly, I explored a wide range of technical solutions for the retrieval of relevant and trustworthy health information such as web search-engines, recommender systems and personalized health education systems. Building on the knowledge acquired during the dissertation, I proposed a new trust-based metric called HealthTrust for the retrieval of health social media. HealthTrust is a metric measuring the trustworthiness of the content within a health community and it can be used to rank search results of health social media. The rationale for choosing an approach based on social network analysis within a health community relies on the assumption that health communities have a common shared knowledge about the relevance and trustworthiness of the content and their providers. The HealthTrust algorithm was successfully tested for the retrieval of diabetes social videos.

- .

Contents

List of Figures

List of tables

Key concepts

Trustworthiness: within the context of this dissertation, the term trust refers to the quality of being believed. Trustworthiness can be applied both to a particular content (e.g. online video) or the author of the content. Trustworthiness towards content might vary between users and communities.

Health Social Media: is the health-related use of social media tools (aka web applications) that allows the creation and sharing of user-generated content. The term Social Media is commonly used to refer both the tools and the content. Most popular example of social media is Facebook, but other sharing content sites such as YouTube are also considered social media.

ePatient: the term ePatients is commonly used to refer to health consumers (patients, healthy people, caregivers) that uses the internet for their own personal health purposes. The term has gradually evolved to refer to empowered patients, who take active role in their health taking advantage of new technologies.

eHealth: according to the World Health Organization, eHealth is the use of information and communication technologies (ICT) for health. Examples include treating patients, conducting research, educating the health workforce, tracking diseases and monitoring public health.

Metadata: metadata is commonly referred the "data about data", so the structured description of data. For example, a video file will have metadata with the title of the video, subtitles, encoding, etc.

Information Retrieval: information retrieval is the activity of obtaining information resources from a collection of resources (e.g. videos from YouTube). The retrieval of information is based on a particular information need (e.g. user searching for diabetes videos). The most popular systems for online retrieval information are Web Search Engines and Recommender Systems.

List of papers[1]

RQ1.Paper 1

Gómez-Zúñiga B, **Fernandez-Luque L**, Pousada M, Hernández-Encuentra E, Armayones M. ePatients on YouTube: Analysis of Four Experiences From the Patients' Perspective. Med 2.0 2012;1(1):e1

RQ1.Paper 2

Fernandez-Luque L, Elahi N, Grajales FJ 3rd. An analysis of personal medical information disclosed in YouTube videos created by patients with multiple sclerosis. Stud Health Technol Inform. 2009;150:292-6.

RQ1.Paper 3

S Konstantinidis, **L Fernandez-Luque**, P Bamidis, R Karlsen. The Role of Taxonomies in Social Media and the Semantic Web for Health Education. Methods Inf Med 2013; 52

RQ1.Paper 4

E Gabarron, **L Fernandez-Luque**, M Armayones, A YS Lau. Identifying measures used for assessing quality of YouTube videos with patient health information: A Review of Current Literature. Interact J Med Res 2013;2(1):

RQ1.Paper 5

Syed-Abdul S, **Fernandez-Luque L**, Jian WS, Li YC, Crain S, Hsu MH, Wang YC, Khandregzen D, Chuluunbaatar E, Nguyen PA, Liou DM. Misleading health-related information promoted through video-based social media: anorexia on YouTube. J Med Internet Res. 2013 Feb 13;15(2):e30.

RQ2.Paper 1

Fernandez-Luque L, Karlsen R, Bonander J. Review of extracting information from the Social Web for health personalization. J Med Internet Res. 2011 Jan 28;13(1):e15. doi: 10.2196/jmir.1432.

RQ3.Paper 1

Yom-Tov E, **Fernandez-Luque L**, Weber I, Crain SP. Pro-anorexia and pro-recovery photo sharing: a tale of two warring tribes. J Med Internet Res. 2012 Nov 7;14(6):e151. doi: 10.2196/jmir.2239.

RQ3.Paper 2

Chomutare T, Arsand E, **Fernandez-Luque L**, Lauritzen J, Hartvigsen G. Inferring community structure in healthcare forums. An empirical study. Methods Inf Med. 2013;52(2):160-7. Epub 2013 Feb 8.

RQ4.Paper 1

Fernandez-Luque L, Karlsen R, Melton GB. HealthTrust: Trust-based Retrieval of YouTube's Diabetes Channels, 2011, 20th ACM international conference on Information and knowledge management.

RQ4.Paper 2

Fernandez-Luque L, Karlsen R, Melton GB. HealthTrust: A Social Network Approach for Retrieving Online Health Videos. J Med Internet Res. 2012 Jan 31;14(1):e22.

---

[1]The description of the papers and their relevance to this thesis is provided at the end of the first chapter.

PART I – Summary

1. Introduction

Most people with access to the Internet will search for information regarding their health or their loved ones' health[151,93]. In their search for online information, they will often find social media. Kaplan and Haenlein defined social media as consisting of a "set of Web applications, which allows the creation and exchange of user-generated content"[85]. Thus health social media can be defined as "the application of social media in the health domain". Social media is becoming a popular channel for the dissemination of health information[97,123].For example, more than 500 channels have been created on YouTube by American hospitals, containing thousands of videos [13]. Similarly, the United Kingdom's National Health Service has published more than 500 videos on YouTube [165]. As I explain below, finding high-quality social media is no easy task despite the abundance of content.

This PhD dissertation researches the problem of the health information overload, especially in social media and online videos. I aim to increase our understanding about how health social media content is generated, disseminated and consumed. The case of online videos is of paramount importance, since it is one of the most popular types of online content and has been shown to have great potential for the education of both patients and professionals. My final objective is to design new tools and algorithms to make it easier to find relevant resources for health consumers. This is a particularly important societal challenge since social media is used massively in our society, and as I explain later in the introduction, the quality of health social media can be very heterogeneous.

In this chapter, I provide a complete overview of the work carried out in this dissertation. First of all, the subsection "Background for the research" summarizes the background of the work carried out in this dissertation and it is followed by a brief summary of the "Research Gaps". Secondly, I introduce the "Research Problems and Questions". The section "Research Context" introduces the context of the research. In the following subsections of Research Approach and Research Design, I explain how the research was carried out. Finally, I bring the introduction to a close with a summary of the main contributions and papers.

1.1.Background for the research

The perfect storm

Nowadays, we have access to a huge amount of online health information such as videos, blogs, and web portals. The perfect storm of online content has been catalyzed by the appearance of social media. Kaplan and Haenlein defined social media as consisting of a "set of Web applications, which allows the creation and exchange of user-generated content"[85]. For example, YouTube allows the creation and exchange of videos and Flickr the sharing of photos. Most health agencies, hospitals and healthcare organizations publish content on social media channels. As S. Fox reports, patients and individual healthcare professionals are also creating social media content [151]. This is not surprising if we take into account the popularity of social media channels such as Facebook (2nd most visited web worldwide) and YouTube (3rd most visited web worldwide according to www.alexa.com). As Figure 1 shows, hundreds of American hospitals publish videos on YouTube [13]. A similar trend has also been found in Europe[10]. Not surprisingly, this growth in online health videos is being driven by an increase in demand. In fact, most adults in USA and Europe already access the Internet to search for health information [93,151]. This trend is also found in emerging countries such as China, India and Brazil[109].
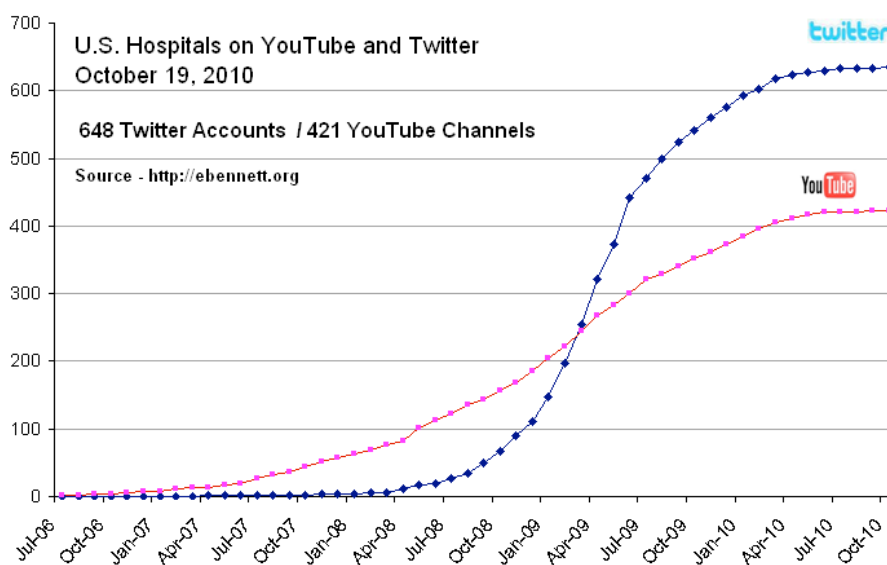
Figure 1: US Hospitals and Social Media (source Ed Bennet 2010 [13])

3

A very complex storm

Health consumers can be overwhelmed by the amount of information and in addition they have to contend with misleading content[120,136].The online health domain is a very complex context for the application of information retrieval techniques. One reason for this is that the concept of "quality" is not clear since it can refer to the technical quality of the video, the medical content or its popularity[60]. What is more, *relevance* which is a traditional information retrieval metric is not trivial in the health domain since it is highly personalized. For example, a video about cooking without sugar will be relevant for most people affected by diabetes but not for people affected by a cystic fibrosis-related diabetes who can eat food with sugar. A third challenge is the appearance of misleading and harmful information[120,136],such as promoting anorexia as a lifestyle. As explained later in my studies about anorexia videos, misleading information can be popular, relevant for the topic, of high quality (e.g. visually appealing) and even contain accurate information. For example, members of the pro-anorexia online community share tips such as taking laxatives which is a dangerous but effective way of losing weight and those sharing such harmful information may be highly reputable within their pro-anorexia community.

YouTube is a good example of the "perfect storm" of Health Social Media. The video sharing platform YouTube was created in 2006. It has gradually become the third most visited webpage worldwide and the biggest repository of videos. YouTube's global audience has motivated many healthcare actors to publish content on that platform including the World Health Organization, the UK´s National Health Service, the New England Journal of Medicine, patients, medical associations and individuals. The content provided by those actors does, however, not solve the issue of quality since misleading videos are reportedly highly ranked when searching for specific topics.

The symbiosis of trust and relevance is the key

The ubiquitous concepts in the online health debate are "reliability", "credibility", "reputation" and "trust" of content. As explained in Chapter 5, these inter-related concepts are highly complex but in most cases refer to the reputation acquired by an academic degree or professional license, and the reputation built within a community (e.g. the most reputable dermatologist in the area). Information and trust are a crucial part of health education since the "message" (e.g. healthy eating) is more or less credible depending on the trust one has in the messenger. In the context of health social media, the role of online

communities in building trust is crucial. In fact, online health communities are complex structures where different stakeholders and sub-communities co-exist and influence each other[28]. To retrieve online information, one has to consider the relevance of the information and the trustworthiness of the messenger.

The web information retrieval approach

In the current context of information overload, more information does not necessarily make it easier to find relevant health information. Not surprisingly, most health consumers use general web search engines to find health information[151]. As explained in this dissertation, traditional web information retrieval tends to retrieve highly popular content which in many cases is bogus and misleading health information[54]. The popularity of YouTube extends to the health domain, and it represents a prime example of the problems of finding high-quality social media. In fact, despite the advanced information retrieval algorithms developed by YouTube, which is owned by Google, there is great concern about the quality of health videos on that platform where it is common to find highly ranked videos promoting anorexia as a lifestyle or lobbying against vaccination [86,136,152].As I will explain in this dissertation, you can find very good and trustworthy videos on YouTube but they are often less visible than misleading and harmful videos. Given this situation, I decided to focus on studying why misleading content is often ranked highly when using standard retrieval tools such as YouTube Search. The main challenges of using general web search engines for health social media are:

- Non-relevant content is sometimes highly ranked due to its popularity outside the health-domain. For example, a video about a singer who happens to be diabetic tends to be highly ranked despite its minimal value for patients with diabetes.

- Misleading information is also highly ranked. There are communities which promote misleading information, e.g. anti-vaccination, and also bogus content promoting fake cures for incurable diseases.

The web health information retrieval approach

Traditional web health information retrieval has mainly been based on quality certification or seals, combining both manual and semi-manual approaches[47]. A very common approach has been the promotion of quality labels given by third party organizations such as Health on the Net (HON). However, these approaches that involve

the manual review of content by health professionals is too costly to be applied in the exploding health social media context[14]. Their vision of quality focuses mainly on the medical trustworthiness and reliability of the content and its provider. Some of the limitations of that approach are:

- Online Health Certifications and quality seals, such as HON, are often binary: a very complete health portal may have the same seal as a poorly edited blog.

- Certifications are traditionally given to the site but not to its content: a blog may have the quality seal but not each blog post.

- Most quality seals ignore non-medical parameters such as complexity of the language, joyfulness, technical quality, etc.

Other approaches

The problem of finding relevant information for patients is not only addressed as an information retrieval challenge. The domain of personalized health education has been dealing for many years with the recommendation of online health information that is personalized according to the unique needs of patients[102]. Traditional personalized health education is based on structured metadata of the profiles and educational resources, while in the social media user profiles and metadata of resources are heterogeneous and incomplete.

1.2.Research Gaps

First Research Gap: the lack of knowledge about Health Social Media

This PhD project started in 2007 when the explosion of social media was yet to come. YouTube and Facebook were two and three years old respectively. Therefore, the first research gap was the lack of knowledge in the field of Health Social Media. In 2008 (the year I published my first paper focusing on YouTube) only 13 papers about YouTube were indexed in PubMed (medical research database), while in 2012 a total of 85 papers were indexed. Hence, to understand the challenges of finding trustworthy and relevant health social media was a research challenge per se, especially regarding health videos.
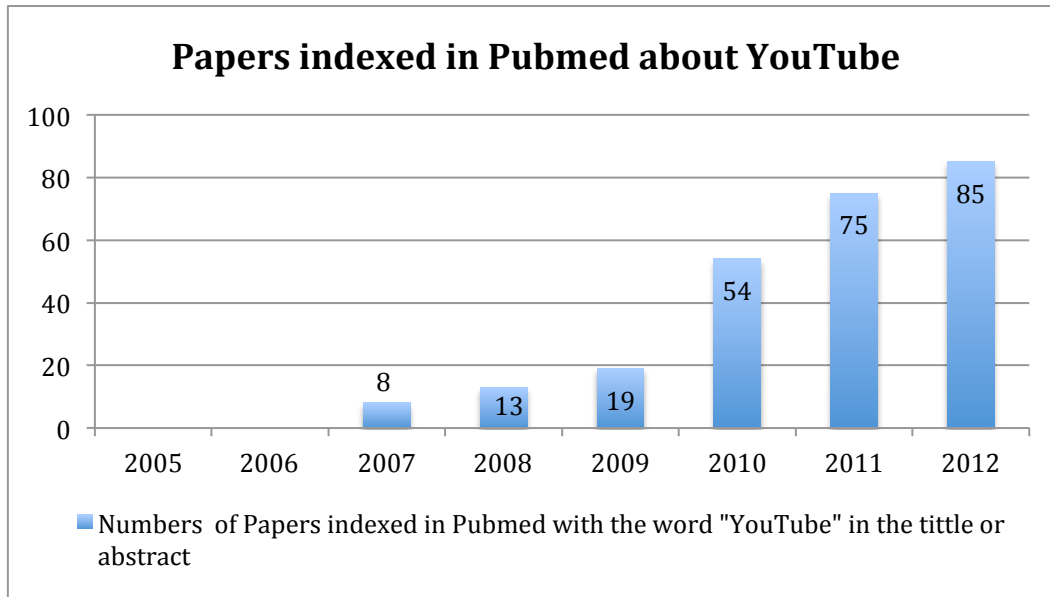
Figure 2: Publications about YouTube in PubMed

<u>Second Research Gap: The lack of information about technical solutions for finding health social media</u>

As I explained in the background sub-section, a wide range of possible technical solutions exist for facilitating the retrieval of relevant trustworthy health social media. Many of those advanced techniques have been successfully applied in the area of social media. Although the quality of health content retrieved by web search engines has been contested, traditional web information retrieval research has not addressed the health domain as a study case. Consequently, very few of those techniques have been applied to the context of health social media and none of them, to our knowledge, to the retrieval of health social videos. Thus in this dissertation we had to face the research gap in the identification of technical approaches for the retrieval of health social media.

<u>Third Research Gap: Trust-based metrics for the retrieval of health social media based on social network analysis.</u>

It is well known that trust and credibility are key factors when finding online health information. Trust in a health website can be affected by multiple factors, such as accessibility and style, and not merely the accuracy of the content [143]. In addition, online communities of patients are known to be very good at finding trustworthy information[45]. Part of the success of online health communities lies in the creation of a social network where influence and trust is built on solid links between users. One of the main research gaps addressed in this dissertation is how to extract trust-based metrics from online

communities using social network analysis. A related research challenge is how to apply those metrics to improve the retrieval of health social media. Although that approach is not entirely new outside the health domain, very few studies on social network analysis have applied to online health communities [24,28,104]and even fewer have focused on the retrieval of health social media.

1.3.Research Questions

This dissertation addresses the lack of research in the area of health social media, particularly videos, from an information retrieval point of view. The overall problem of finding trustworthy online health videos has raised the following research questions.

How can computing techniques support the retrieval of trustworthy health social media?

This broad research question can be narrowed down and divided into the following secondary research questions.

- What are the challenges of finding health social media and videos in particular?
- Which are the technical solutions for modeling health social media?
- How can Social Network Analysis be used to extract information about the characteristics of health social media?
- Can trust-based metrics improve the retrieval of social videos about diabetes?

The first two research questions can be grouped as background research within this dissertation. The third research question deals with experiments related to social network analysis of health social media. The latter uses a metric derived from social network analysis for retrieving health social media. The research performed to answer **RQ1** and **RQ2** was necessary to understand the context of the more experimental research of this dissertation addressed in **RQ3** and **RQ4**.The research conducted to answer **RQ1** and **RQ2** was necessary to identify the most experimental research of this dissertation addressed in **RQ3** and **RQ4**.

The research in this dissertation has used **several study cases**. The selection of **anorexia** as a case-study was due to the problem of the sub-communities promoting anorexia as a lifestyle, consequently making that case-study one of the most complex and interesting. Also, chronic conditions such as **diabetes** and **multiple sclerosis** have been used in

several studies. People with chronic conditions are prime examples of health information seekers, who also use the Internet to socialize. Although the findings of those study cases cannot be automatically generalized to the overall health social media context they represent a fair representation of the main problems.

RQ1: What are the characteristics of health social videos?

RQ1 deals with the characterization of the problems of finding trustworthy health social media. In other words, **RQ1** focuses on research to understand the context of health social media (e.g. metadata, users' motivations). **RQ1** has been very challenging due to the immaturity of the research in this area, especially concerning online videos. This research question is sub-divided into the following sub-questions:

- **RQ1.1**: Does the online community influence the motivation of people with chronic conditions to publish videos about their health?
- **RQ1.2**: Do health videos contain relevant medical vocabulary in their textual metadata?
- **RQ1.3**: What are the quality features of online health videos?
- **RQ1.4**: Do misleading and informative online videos on the topic of anorexia have different characteristics?

RQ2: Are there technical solutions for modeling health social media?

RQ2 deals with the identification of computing techniques that can be used to address the modeling of health social media. Modeling of users and content is a crucial part of the information retrieval process.

RQ3: How can Social Network Analysis be used to extract information about the characteristics of health social media?

I established the third research question to explore how social network analysis can be used to characterize online communities and trust within these communities. Online diabetes communities were selected as the case-study since it is one of the most common chronic diseases. The selection of anorexia as a case-study was due to the problem of the sub-communities promoting anorexia as a lifestyle. This question is divided into the following questions:

- **RQ3.1**: Can social network analysis be used to infer the misleading nature of

social photos in the case of anorexia?
- **RQ3.2**: Do the most centric member on diabetes online social networks have different health characteristics, such as experience living with the disease?

RQ4: Can trust-based metrics improve the retrieval of social videos about diabetes?

RQ4 is designed to analyze how a trust metric derived from social network analysis could be used to improve the retrieval of online diabetes videos. The goal of this research question is to study the possibility of using metrics based on social network analysis to improve the retrieval of diabetes videos.
- **RQ4.1**: Can a metric of trustworthiness within a health community be used to retrieve relevant trustworthy providers of diabetes videos?
- **RQ4.2**: Can a metric of trustworthiness within a health community be used to search for relevant trustworthy diabetes videos?

1.4.Research Context

I carried out my research while working at Norut, which is a multidisciplinary applied research institute located in Tromsø. I am working in the ICT department of the research institute with around 10 people, nearly half of the work is done on eHealth projects. I am also an active member, as a student, of the Computer Science Department of the University of Tromsø, especially in the Medical Informatics and Telemedicine group.

This dissertation started in 2007 when social media was just becoming popular and almost no research had been conducted in health social media. For example, YouTube was created in 2005 and now it is the third most visited website worldwide and hundreds of hospitals publish content on YouTube. This is just an example of how fast the field of health social media is evolving. Another additional problem is the multidisciplinary understanding required to have a comprehensive knowledge of this field where the borders between research domains such as Computing, Social Science and Health are hazy. So a major problem in this PhD was the lack of knowledge about how health social media is created, disseminated and used. This problem was addressed by a strong partnership with fellow researchers from multiple disciplines as explained below.

The papers included in this dissertation are the result of a long research project in a very rapidly evolving and immature field. This challenge was counterweighed by an

exponential adoption of health social media that has facilitated access to data and publications due to the unprecedented interest in the field. As an example of the interest in the topic I have been regularly invited to interviews on th radio, printed media, and keynotes at conferences. This interest in our research is partially responsible for the extra-time needed for the dissertation.

The work performed in this dissertation took place with the collaboration of multiple research groups such as:

IMIA Social Media Working Group (International): from the beginning of my dissertation I started to collaborate with a group of researchers and practitioners interested in health social media, who ultimately created the IMIA Social Media Working Group. Collaboration with this group has been crucial for the progress of the dissertation since I had continuous conversations with fellow researchers about my research. Most of the papers presented in this dissertation contain co-authors from the group.

Medical Informatics and Telemedicine Group (UiTø, Norway): I am an affiliated member of the MI&T group at the Computer Science Department of the University of Tromsø led by Prof. Gunnar Hartvigsen. This group provided the academic medical informatics angle in my work.

Open Distributed Systems Group (UiTø, Norway): I am an affiliated member of the open distributed system group at the Computer Science Department of the University of Tromsø, where my supervisor, Associate Professor Randi Karlsen, is a member. My supervisor was crucial in providing the computing angle to my work, especially regarding Information Retrieval.

ITACA-TSB (Polytechnic University of Valencia, Spain): Dr. Vicente Traver is the leader of the research group TSB at the ITACA Institute in the Polytechnic University of Valencia. He has been my mentor for nearly a decade and during this PhD project he has helped me as a very active co-supervisor with lively discussions about eHealth and Health Social Media. We have published three books about the topic in Spanish during my PhD.

GroupLens and Institute of Health Informatics (Univof Minessota, USA): During my stay in Minnesota I collaborated with Prof. Joseph Kostan of GroupLens. This group is a reference in social computing and recommender systems. They helped me understand methodological aspects of Information Retrieval research. In addition, I collaborated with

Prof. Genevieve M. from the Institute of Health Informatics. She guided me in the process of applying for an IRB (Institutional Review Board) and the recruitment of patients.

Children's Health Informatics Program (Harvard Medical School, USA): During my stay abroad in the USA I was hosted in the group of Assis. Professor Kenneth Mandl. That stay was crucial for a better understanding of the health domain and also the challenges of deploying web applications for health consumers.

Diabetes Hands Foundation (USA): the Diabetes Hands Foundation and its president Manny Hernandez were crucial for the success of this project. They facilitated the recruitment of patients within the online community TuDiabetes and in lengthy conversations helped me understand the problems that people with diabetes face when searching for health information.

Psinet Research Group –Open University of Catalunya (Spain): this group of the Open University of Catalunya was crucial to help me understand the psychological factors that patients face when seeking online health information or peer-support. Dr. Manuel Armayones from that group has been collaborating with me for most of the PhD project.

Microsoft Research/Yahoo Research (Israel): The collaboration with Dr. Elad Yom-Tov (now working at Microsoft Research and previously at Yahoo) on RQ3.Paper1 was very important since his expertise in online data mining allowed me to acquire more knowledge about sub-communities' network dynamics.

1.5.Research Funding

The funding for this project was provided by multiple research projects. The main source of funding has been the Tromsø Telemedicine Laboratory (a Centre for Research-based Innovation co-funded by the Research Council of Norway, project 174934), more specifically the projects pEducator and MyHealthService. In addition, we had some funding from the HealthTrust project (funded by TromsøForskningsstiftelse). Last but not least we have used internal funding at Norut for several activities (e.g. research assistants).

1.6.Research Approach

1.6.1. Research Design

As explained in Figure 3 below, in collaboration with fellow researchers I conducted eight different studies to gain knowledge to answer all the different research questions. These studies led to 10 publications which are included in this dissertation.
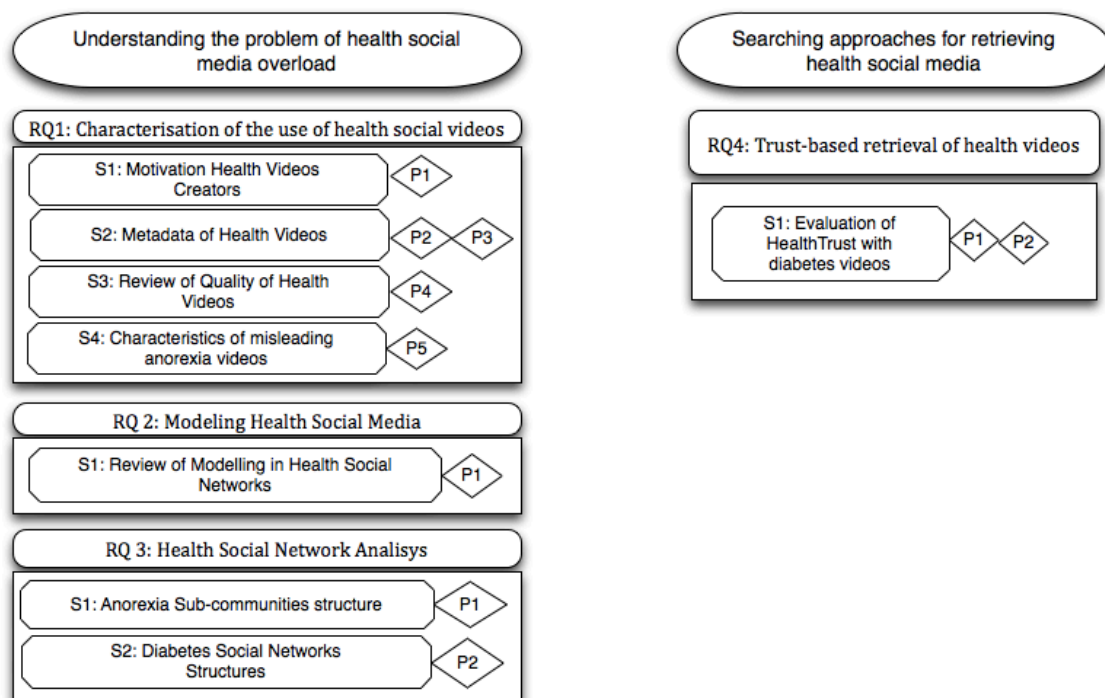
Figure 3: Study Design (P= paper, S=Study, RQ=Research Question)

In order to address **RQ1** about the current challenges of finding health social media and videos we performed mainly two types of studies: literature reviews and qualitative content analysis. We performed **RQ1.Study 1** consisting of the analysis of videos from patients who were sharing about their disease on YouTube. The methodology of this study was based on the use of qualitative techniques under the supervision of a panel of psychologists. Since the metadata of the health videos was crucial for information retrieval we performed a couple of studies (**RQ1.Study 2**) to understand the characteristics of metadata from health videos (e.g. comments and semantics). In addition, **RQ1.Study 3** includes literature reviews on aspects such as the quality of health videos. An important

aspect of this research is the understanding of the human aspects that stimulate the creation of health videos. Finally, I decided to conduct a more detailed study (**RQ1.Study 4**) in the case of anorexia related videos on YouTube. The rationale of this study was to gain a better understanding of the differences between harmful videos (e.g. pro-anorexia videos) and more informative videos.

RQ2 deals with the understanding of which technical solutions can be used for modeling health social media content and users. These technical solutions include a wide range of technologies such as collaborative filtering, social network analysis, modeling and personalization. **RQ2.Study 1** consists of a literature review of how metadata from health social networks can be used to extract information to enhance information retrieval. The literature review was done with the collaboration of an Information Retrieval expert (my supervisor) and an anthropologist working in a public health organization. The goal was to address and discuss technical, health and ethical aspects related to the modeling of health social media.

In **RQ3,** we performed several studies using social network analysis to characterize online communities. **RQ3.Study 1** explored the sub-community interaction in the case of anorexia where two distinguished sub-communities were interacting and one was promoting a harmful view of the disease. **RQ3.Study 2** focused on the study of the characteristics of diabetes communities, including the influence of the different types of members.

The results from **RQ1** and **RQ2** were crucial for designing the trust-based metric HealthTrust, which was designed and tested under **RQ4. Study 1**. HealthTrust is a metric based on social network analysis designed to identify the trustworthiness of health social media. In RQ4.Study 1 the HealthTrust metric was evaluated to search for diabetes videos, using the search results for diabetes videos on YouTube as baseline.

Table 1 List of Studies

RQ No. Study No.	Purpose	Papers
RQ1.S1	To understand the motivations of patients publishing online videos based on video-interviews of expert users.	RQ1.Paper1
RQ1.S2	To characterize health video metadata. In particular, comments on videos about multiple sclerosis, and use of medical terms in surgery videos.	RQ1.Paper2 RQ1.Paper3
RQ1.S3	Based on research literature, to understand different quality features of online health videos.	RQ1.Paper4
RQ1.S4	To understand better the differences between untrustworthy and trustworthy YouTube videos in the case of anorexia.	RQ1.Paper5
RQ2.S1	Using a literature review, to understand which metadata form health social networks can be extracted to enhance information retrieval.	RQ2.Paper1
RQ3.S1	To understand the structure and dynamics of anorexia communities within social media platforms.	RQ3.Paper1
RQ3.S2	To understand the structure and dynamics of diabetes communities within social media platforms.	RQ3.Paper2
RQ4.S1	HealthTrust: to study how social network metrics can be applied to search diabetes videos.	RQ4.Paper1 RQ4.Paper2

1.6.2. Research Methods

Interdisciplinary research is at the core of the Medical Informatics research in which this PhD is framed. The health domain is very complex and includes societal, health and psychological aspects. Research on the application of informatics in the health domain is inevitably multidisciplinary. The multidisciplinary nature of this dissertation has been addressed by the collaboration with people from different disciplines across the different studies. Among others, we have collaborated with psychologists, healthcare professionals, anthropologists, public health researchers, social scientists, computer scientists and end

users. The following table describes the research methods and data sources of the different studies in this dissertation.

Table 2: List of Studies and Research Methods

RQ No. Study No.	Research Method
RQ1.S1	Content analysis of video interviews of 4 patients members of YouTube that were transcribed, analyzed and categorized. (RQ1.Paper 1)
RQ1.S2	1) Extraction from YouTube of 7,047 comments from 769 videos published by (self-reported) patients with multiple sclerosis. A final random selection of 320 comments was analyzed regarding the disclosure of private health information. (RQ1.Paper2) 2) Extraction and Natural Language Processing of 64,367 tags from 4,307 YouTube videos about surgery. Tags were analyzed to determine the prevalence of standardized medical vocabulary (i.e. SNOMED CT). (RQ1. Paper3)
RQ1.S3	Systematic literature review (185 abstracts retrieved, 13 papers selected) and classification of quality feature in online health videos. (RQ1.Paper4)
RQ1.S4	Extraction of data from 7,583 anorexia-related videos, classification and categorization of 140 videos, followed by statistical analysis (e.g. ANOVA multivariable) to identify differences between informative and misleading videos about anorexia. (RQ1.Paper5)
RQ2.S1	Descriptive literature review of papers from different domains about the extraction of information from health social networks. (RQ2.Paper1)
RQ3.S1	Data extraction of 543,891 Photos and over 3 million social links from 753 users publishing about anorexia on the photo-sharing site Flickr. Web Data Mining and Social Network Analysis of photo-sharing communities about Anorexia. (RQ3.Paper1)
RQ3.S2	User profile extraction and Social Network Analysis of online diabetes communities with a total of 140,000 registered users and 1.6 million posts. (RQ3.Paper2)
RQ4.S1	Information retrieval evaluation of diabetes videos/channels retrieved using the HealthTrust metric. Evaluation based on ratings from medical experts (162 video ratings) and patients (427 video ratings), using the results retrieved by YouTube as a benchmark. (RQ4.Paper1 and RQ4.Paper2).

1.7.Claimed Contributions

1.7.1. Contributions of thesis

The primary objective of this thesis is to study how to apply Information Retrieval techniques to improve the retrieval of high-quality health videos within the context of the social web. The following list contains a summary of the main contributions.

As we mention in the discussion, the entire challenge of online health information retrieval cannot realistically be addressed in one dissertation. Thus, it is not possible to automatically generalize all my contribution. However, the study cases were carefully designed to have a reasonable representation of the most significant challenges (e.g., information seeking in chronic diseases, topics with a high prevalence of misleading information).

Contribution 1: Increased knowledge in the use of health social videos

Due to the incipiency of the field of health social media when this PhD project started, our background studies constitute a clear contribution to the eHealth field. I contributed with new knowledge about motivational aspects of generators of content, quality of health video metadata, popularity features of misleading content, etc.

Contribution 2: Increased knowledge of the challenges related to the modeling of health social media and videos in particular

This dissertation is one of the first studies to identify and analyze the technical challenges of modeling health social media.

Contribution 3: Social network analysis of online health communities

In this dissertation the network dynamics of health communities have been studied in order to better understand the influences within the communities. This contribution was based on studies on diabetes and anorexia communities. I participated in research that provided new algorithms for clustering anorexia sub-communities based on social network analysis.

Contribution 4: Social Network analysis to infer quality of health social media

> The approach of using social network metrics to infer trust within health social networks and apply it to the retrieval of diabetes videos is a clear contribution to the field. I designed, developed and tested algorithms for the retrieval of online diabetes videos.

Table 3: List of key findings

No.	Key finding - Contribution	Paper
K1.1	Social interaction is one of the main driving forces behind those publishing videos about their health.	RQ1.Paper1
K1.2	Textual metadata can be of very heterogeneous quality, but still contains a lot of relevant health information for modeling	RQ1.Paper2 RQ1.Paper 3
K1.3	The quality of health videos is a multidimensional concept. Reliability of the content and its provider are very important quality criteria according the literature.	RQ1.Paper4
K1.4	Common popularity metrics, such as favoring ratio, correlate negatively with the trustworthiness of anorexia-related videos.	RQ1.Paper5
K2.1	Most technical solutions for modeling social media will have shortcomings in the health domain due to text analysis complexities, privacy issues and popular but harmful content. Link and Social Network Analysis is promising but has not been studied in detail in the health domain.	RQ2.Paper1 **RQ1.**Paper
K3.1	On a photo-sharing site, the best predictors of users belonging to the misleading sub-community promoting anorexia are social network metrics. Tag-based classification was less accurate.	RQ3. Paper1
K3.2	Most experienced members on online diabetes communities have longer vital experience with the disease.	RQ3.Paper2
K4.1	In diabetes online communities the most reputable members are those with more experience with diabetes.	RQ4.Paper1
K4.2	The HealthTrust metric based on Social Network Analysis to infer quality of health videos performs well for filtering misleading content compared to YouTube searches.	RQ4.Paper2

1.7.2. Dissemination and Exploitation

Publicly funded research is a contract with society and I am therefore of the firm belief that it is our moral obligation to disseminate as much as possible among the public who support our work with their taxes. In addition, my research affects a significant part of our society who search for online health information and also use social media.

I have published other nine papers in areas related to this dissertation[1,53,55,95,124,134,142], but they do not address the specific research questions of the dissertation (e.g. white papers, opinion letters). Those papers are part of the dissemination of my work since they helped to increase awareness and scientific discussion about the topic. For example, I have participated in brief communications about online health and social media that have been published in The Lancet [1] and the British Medical Journal [142].

My research has been featured in major newspapers in Spain [39] and India [84], and also in two health-related newspapers in Spain [12,127], and radio interviews in Norway and Spain. In addition to press features, I have been collaborating in five books about Health Social Media in Spanish [9,61,139,156,157] and Norwegian [82]. I have co-edited two translations to Spanish of books about ePatients [35,61] and co-edited three books about health social media in Spanish [56,156,158]. More than 1000 copies of those printed books have been distributed, and over 20,000 copies of digital editions have been downloaded.

I have been invited to give keynotes at several conferences about topics related to my dissertation, such as the *Annual Norwegian National Gynecology and Obstetrics Conference* (2012), the *Spanish Congress of Patients with Cancer* (2013), the *Spanish Congress of Patients with Chronic Diseases* (2014), the open session of the *Spanish Medical Informatics Conference* (2014) and Keynote at the *Taiwanese Annual Medical Informatics Conference* (2014). Most of my slides from conferences are available in Slideshare (available at htttp://slideshare.net/luis.luque), which have accumulated thousands of views. I have also been involved in the Scientific Program Committees of several international conferences such as (IEEE CBMS, Medicine 2.0, IEEE BHI). I was also chairman for 3 years of the conference Salud 2.0 Euskadi, a conference host in Bilbao (Spain) about Health Social Media.

I believe that research without technology transfer is less valuable for society. During my PhD, I had the great help of excellent students from the University of Sevilla (Spain) who

got very interested in the topic. They decided to become entrepreneurs in the area of Health Social Media. I supported them in the process both financially and advising about the complexities of the eHealth business. Since the end of 2011, the company has provided expert IT consultancy services in the domain of mobile and social apps for health. The company employs six people and has clients in more than five countries.

1.8.Description of Papers

RQ1.Paper 1: Gómez-Zúñiga B, **Fernandez-Luque L***, Pousada M, Hernández-Encuentra E, Armayones M. *ePatients on YouTube: Analysis of Four Experiences From the Patients' Perspective*. Med 2.0 2012;1(1):e1

- Relevance: this paper presents a study on the motivations of patients publishing videos on YouTube. The relevance for this PhD dissertation lies in the understanding of the human factors of the main problem addressed.
- My contribution: I had the original idea of the study and was responsible for the data collection. I contributed with the analysis of the data discussion of the results.
- Quality indicator (Cited by 3 [2]): this is a new spin-off journal from JMIR, the top journal in the field of eHealth.

RQ1.Paper 2: **Fernandez-Luque L***, Elahi N, Grajales FJ 3rd. *An analysis of personal medical information disclosed in YouTube videos created by patients with multiple sclerosis*. Stud Health Technol Inform. 2009;150:292-6.

- Relevance: this paper shows results on the study of metadata from the videos' comments. It was interesting to find a significant amount of comments with personal health information.
- My contribution: I led the study. Other authors helped in the analysis of the comments and discussion of the results.
- Quality indicators (Cited by 26): the paper was accepted by the conference Medical Informatics Europe, the leading medical informatics conference in Europe.

RQ1.Paper 3:S Konstantinidis, **L Fernandez-Luque***, P Bamidis, R Karlsen. *The Role of Taxonomies in Social Media and the Semantic Web for Health Education*. Methods Inf Med 2013; 52

---

[2] "Cited by" based on Google Scholar metrics (January 2014)

- <u>Relevance</u>: in this paper we studied the semantics of metadata (i.e. tags) in online videos about surgery. The relevance of this study lies in the possibility of using semantic information to model health videos.
- <u>My contribution</u>: I was not the leader of the paper, but I contributed to the study design, data collection and discussion especially regarding health videos.
- <u>Quality indicator (Cited by 3)</u>: Methods Inf Med is the oldest journal in the medical informatics field, it was founded in 1962. It has an impact factor of 1,472 (2010)

RQ1.Paper 4: E Gabarron, **L Fernandez-Luque***, M Armayones, A YS Lau. *Identifying measures used for assessing quality of YouTube videos with patient health information: A Review of Current Literature*. Interact J Med Res 2013;2(1):e6

- <u>Relevance</u>: in this paper we studied the research literature to identify which quality features are to be considered in online health videos.
- <u>My contribution</u>: as co-author I participated in the study design, data collection and manuscript writing. I was the author bringing the technical perspective to the study.
- <u>Quality indicator (Cited by 6)</u>: this is a new spin-off journal from JMIR, the top journal in the field of eHealth. It is already indexed in PubMed but does not yet have an impact factor.

RQ1. Paper 5: Syed-Abdul S, **Fernandez-Luque L[3]***, Jian WS, Li YC, Crain S, Hsu MH, Wang YC, Khandregzen D, Chuluunbaatar E, Nguyen PA, Liou DM. *Misleading health-related information promoted through video-based social media: anorexia on YouTube*. J Med Internet Res. 2013 Feb 13;15(2):e30.

- <u>Relevance</u>: in this study we explored the reliability of YouTube videos about anorexia. We found a significant amount of misleading videos promoting anorexia as a lifestyle and some of the shortcomings of the results retrieved from YouTube.
- <u>My contribution</u>: I co-led the paper with researchers from the Tapei Medical University (Taiwan). My main contribution was with the study design, data collection and the interpretation of the results.
- <u>Quality indicators (Cited by 10)</u>: this paper is published in JMIR, which is the leading eHealth journal with a JCR IF of 4.7 (2010).

---

[3]First authorship shared among the first three authors

RQ2.Paper 1: **Fernandez-Luque L***, Karlsen R, Bonander J. *Review of extracting information from the Social Web for health personalization.* J Med Internet Res. 2011 Jan 28;13(1):e15. doi: 10.2196/jmir.1432.

- <u>Relevance</u>: this paper provides a complete review of different techniques for extracting information from health social networks.
- <u>My Contribution</u>: as leading author I coordinated the whole study and paper writing. Other authors provided the socio-ethical and information retrieval perspective, while I focused on health social media.
- <u>Quality indicators (Cited by 31)</u>: this paper is published in JMIR, which is the leading eHealth journal with aJCR IF of 4.7 (2010). In addition, this paper has been cited 18 times according to Google Scholar (February 2013).

RQ3.Paper 1: Yom-Tov E, **Fernandez-Luque L***, Weber I, Crain SP. Pro-anorexia and pro-recovery photo sharing: a tale of two warring tribes. J Med Internet Res. 2012 Nov 7;14(6):e151. doi: 10.2196/jmir.2239.
- <u>Relevance</u>: in this study we explored the interaction within pro-anorexia and anti-anorexia communities, showing that health communities are inter-related and highly complex structures.
- <u>My contribution</u>: I participated in the original idea of the study, discussion of the results, manuscript writing. I was the only author specialized in the health and eHealth domain.
- <u>Quality indicator (Cited by 6)</u>: this paper is published in JMIR, which is the leading eHealth journal with a JCR IF of 4.7 (2010).

RQ3.Paper 2: Chomutare T, Arsand E, **Fernandez-Luque L***, Lauritzen J, Hartvigsen G. *Inferring community structure in healthcare forums. An empirical study*. Methods Inf Med. 2013;52(2):160-7. Epub 2013 Feb 8.

- <u>Relevance</u>: in this study we explored the structure of online diabetes communities. The relevance for this dissertation is due to the need to have a better understanding of health social networks.
- <u>My contribution</u>: I supported the leading author in the study design and the interpretation of the results.
- <u>Quality indicator (Cited by 5)</u>: Methods Inf Med is the oldest journal in the medical informatics field, it was founded in 1962. It has an impact factor of 1,472 (2010)

RQ4.Paper 1:Fernandez-Luque L*, Karlsen R, Melton GB. *HealthTrust: Trust-based Retrieval of YouTube's Diabetes Channels*, 2011, Proceedings of the 20th ACM International Conference on Information and Knowledge Management.

- Relevance: in this short paper we reported an evaluation of the trust-based HealthTrust metric for the retrieval of diabetes channels from YouTube.
- My contribution (Cited by 5): I coordinated the whole study and paper writing. Other authors helped in the discussion, recruitment of patients and the application for the Institutional Review Board.
- Quality indicator: this short paper was accepted in CIKM 2011, which is one of the top conferences in Information Retrieval according to the Australian Government ranking of conferences (CORE A Level, data available at http://www.core.edu.au/). CKIM has very low acceptance rate (30% for short papers).

RQ4.Paper 2:**Fernandez-Luque L***, Karlsen R, Melton GB. *HealthTrust: A Social Network Approach for Retrieving Online Health Videos*. J Med Internet Res. 2012 Jan 31;14(1):e22. doi: 10.2196/jmir.1985.

- Relevance: this full paper evaluated the HealthTrust metric for the retrieval of social videos about diabetes.
- My contribution: I coordinated the whole study and paper writing. Dr. Melton helped with more clinical discussion, the recruitment of patients and the submission of the IRB (Institutional Review Board) at the University of Minnesota. Associated Prof. Randi provided advice with the algorithm aspects.
- Quality indicator (Cited by 13): this paper is published in JMIR, which is the leading eHealth journal with a JCR IF of 4.7 (2010).

1.9.Thesis Structure

This dissertation is organized as follows. Section 2: "A real life example: finding trustworthy health social media" provides a real life case to understand the perspective of the information seeker. The following sections (from Section 3 to Section 6) describe the different parts of my research designed to answer the research questions. Finally these sections are completed with the discussion of the dissertation (e.g. summary of key findings, limitations, future work) and conclusions. The Appendix includes the papers.

2. A real life example: finding trustworthy health social media

In this section, I will explain a personal experience of searching for health information, in order to reflect the real meaning of quality, trust and relevance for a person who is in need of health information.

In the summer of 2011, I got the nicest surprise ever when we found out we were expecting twins. Shortly after the ultrasound the doctors measured the twins and found a discrepancy in their size, something that is very common in identical twins. However, there is always the possibility that the discrepancy is due to a rare condition called TTTS (Twin-to-Twin transfusion syndrome). The doctor in charge told us we had to be referred to a bigger hospital 1000 kilometers away to have the final diagnosis, because in the event of a confirmed diagnosis, my wife would have to undergo intra-uterus surgery in Germany. In addition, she advised us not to search on the Internet.

As we reached home I started to search on the Internet for information, as most parents will do nowadays. First of all, I found a very clear explanation of the possible problem in MedlinePlus (the health portal of the US Government). Identical twins share the placenta and sometimes blood flow from one fetus can be diverted to the other, leading to growth retardation in one and potentially fatal heart problems in the other. Since they share the placenta the sudden death of one fetus could kill or harm the other. The prognosis without surgery was very grim, and the treatment consisted of dividing the placenta using laser surgery. My search for information on hospital websites and research literature was very unpleasant, they tended to focus on the most serious cases and in many cases the information was old and therefore inaccurate. Suspiciously, some hospitals which performed the laser surgery had worrisome and outdated information about TTTS followed by the recommendation to visit them.

My search then turned to social media. Searching for images is powerful and effective, the search for images in Google for TTTS turned out to be a mixed experience, as shown in Figure 4. There were very good illustrations explaining the physiology of the syndrome and also cute photos of healthy twins. However, among the top ranked images there were also photos of placentas and dead fetuses used for medical education. These images, although relevant, could have a hugely devastating impact on parents who are worried enough as it is; luckily my wife was not into searching.
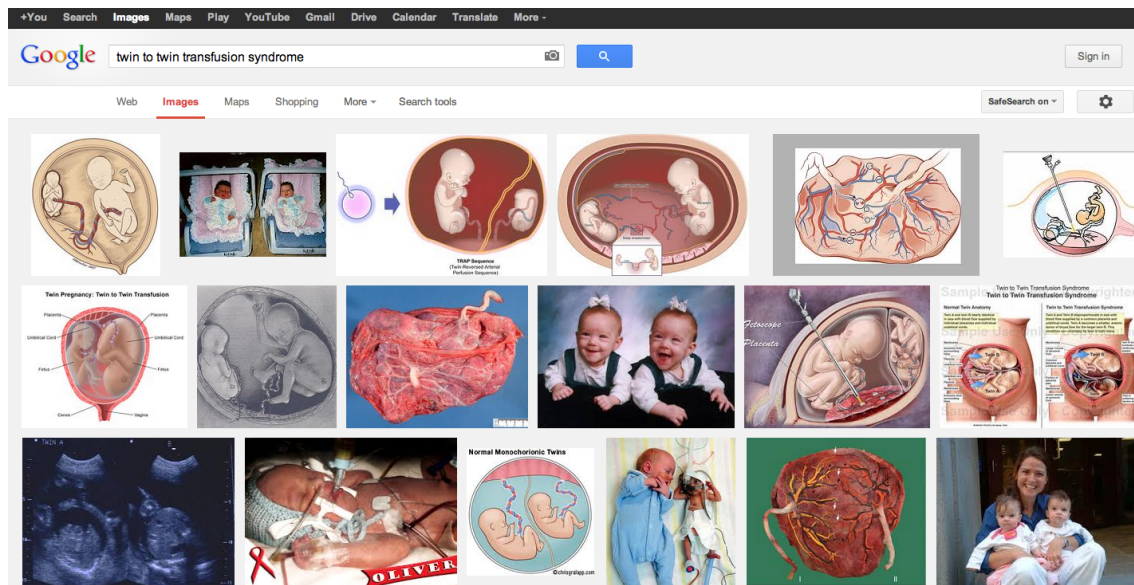
Figure 4: Search results in YouTube Image for "Twin to Twin transfusion syndrome"

After that, I started to search for videos on YouTube. The experience also produced mix feelings. The most common videos were ones made by parents who had lost one or two of their twins, although nowadays with surgery most cases have a better prognosis. "In memoriam" videos were very common, despite in most cases the prognosis being good. If resolved positively, twins do not suffer serious consequences and there is therefore not such a need to build up a community as people with chronic conditions might do. Sadly, sometimes the recommendation made by YouTube for those "in memoriam" videos were very inappropriate (e.g. "baby born without brain turns 2"). These related videos were in fact relevant if we consider the topic to be fetal abnormalities, but from a parent's point of view they were wrong suggestions. On YouTube, I found some videos that explained TTTS for new parents brilliantly in a couple of minutes. If our doctor had suggested that video to us, we would have saved a lot of time. As shown in Figure 5, sometimes very informative videos of a high technical quality (e.g. audio, video) were very hard to find due to the lack of proper descriptions because they had very low-quality metadata. The best videos about TTTS were all made by children's hospitals in the US that are forming a community and linking between them.
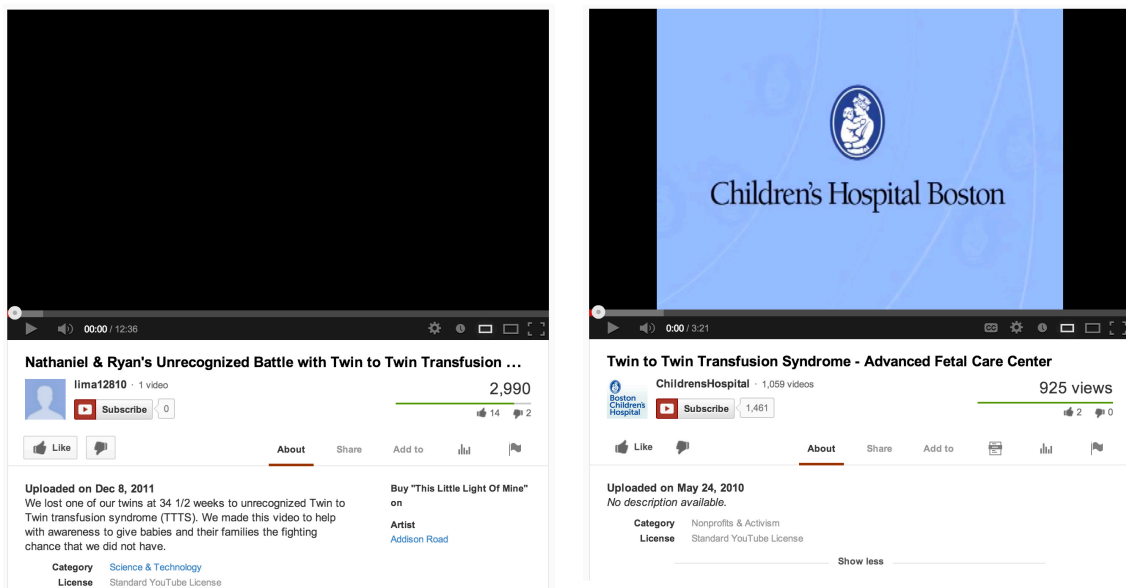
Figure 5: Videos about TTTS with different quality description (metadata)

As informed eParents and without much help from our local hospital, we managed to speed up the process of diagnosis, which is urgent since surgery (if necessary) needs to be done as quickly as possible. At the reference hospital 1000 kilometers from home we found out that we did not have any problem with TTTS but they did diagnose something else and, again, the doctor advised us not to search on the Internet. Today we enjoy two healthy, cute twin girls.

3. Health Social Media and Online videos

Related papers are: RQ.Paper 1, RQ1.Paper 2, RQ.Paper 3, RQ1.Paper 4 and RQ1. Paper5.

RQ1: **What are the characteristics of health social videos?**

- **RQ1.1**: Does the online community influence the motivation of people with chronic conditions to publish videos about their health?

- **RQ1.2**: Do health videos contain relevant medical vocabulary in their textual metadata?

- **RQ1.3**: What are the quality features of online health videos?

- **RQ1.4**: Do misleading and informative online videos on the topic of anorexia have different characteristics?

As explained in Subsection 3.1, online videos are one of the most popular types of content. YouTube is the largest video repository and it is also a social network. Videos are a very efficient channel for health education and communication, and not surprisingly thousands of health institutions and individuals publish health videos on YouTube. However, dozens of studies have reported the difficulties of finding reliable health videos on YouTube.

In the health domain, quality is not just a matter of how useful or attractive information is, because bad information can ultimately be fatal. In this dissertation, I confronted lack of knowledge regarding the following topics:

- The reasons and motivation behind those who create videos.
- A definition of what makes a "good" health social video, and which features are related to quality.
- The characteristics of misleading health videos, such as those promoting anorexia as a lifestyle.
- The characteristics and quality of health videos' metadata.

In order to address those research questions, I designed and participated in a set of experiments. This research was conducted with a wide and multidisciplinary set of collaborators spanning many different areas of expertise. An important study was related to the motivations of users who publish online videos about their health. That study (see

Subsection 3.2) provides insights into community aspects and human factors of health social videos. Subsection 3.3 reports on a review study of the quality features of online health videos. That study was complemented by a study on the differences between misleading and informative videos related to anorexia, see Subsection 3.5. Finally, Subsection 3.5 reports on two studies aimed at characterizing the metadata of YouTube health videos, in particular the disclosure of personal health information and the use of medical thesauri.

3.1. Social Media and Health Videos

3.1.1. What is Social Media?

Since the origins of humanity we have been exploring new ways to communicate. Communication is an inherent part of mankind and most historical revolutions have come about through disruptive innovations to the way we communicate. Marshall McLuhan in his book "Understanding Media: The extensions of Man", explained that media has evolved technologically from the spoken word to manuscripts, print and more recently to the electronic media[110]. Each technical evolution has resulted in an easier production and distribution of media. The emergence of the Internet and the social web in particular has taken the reach of the media to unprecedented levels. Nowadays, somebody can use a mobile phone to make a video and share it instantly on a social network such as Facebook or YouTube, reaching a massive global audience within minutes. What is more, that video can be enriched by comments of thousands of viewers transforming the original media into a massive dialog. According to Wikipedia, social media refers to "the use of web-based and mobile technologies to turn communication into interactive dialogue". Another definition was proposed by Andreas Kaplan and Michael Haenlein who define social media as "a group of Internet-based applications that build on the ideological and technological foundations of Web 2.0, which allows the creation and exchange of user-generated content."[85].

The current explosion of social media has been made possible due to usability and technical improvements of the web. In the mid-90s, creating a website or web content in general was a complex task which only a few could do. In the early 2000s that changed with the emergence of the Web 2.0 when creating online content became easier and more

accessible to the layperson. The mass adoption of user-generated content was crucial for the transformation of the Internet into a massive repository of content and social networks as it is today. One of the most interesting by-products concerns massive cooperation of which Wikipedia, which has been widely used in the health domain, is an iconic example. For example, more than 900 contributors have written the Diabetes Type II page in Wikipedia.

In addition, the web is continuously evolving and there is a tendency towards the merging of media across different social networks via the use of APIs and eventually semantic web technologies also known as the Web 3.0. For example, a video made on a mobile device can be enriched with metadata about the geo-location where the video was made and shared on YouTube and Facebook. Using the YouTube API the description of the video can be semantically extracted and used to provide information from another system such as the medical research database PubMed in order to provide recommendations of medical research papers. In addition to the information overload, we have an emerging context of connectivity overload.

3.1.2. What is Health Social Media?

Health is a very important part of our lives and not surprisingly a significant part of global GDP (Gross Domestic Product) is dedicated to healthcare: ca. 9% in Spain and ca. 18% in the USA (data available at http://data.worldbank.org/). This partially explains why nearly every technical innovation is rapidly adapted to the health domain. The Social Media is not an exception and it is being adopted rapidly in the health domain despite the lack of research to support the safety and efficacy of the application of social media. Experiences have been reported in the health domain in virtually every type of social media, from the use of Twitter to stream surgery live to the creation of a blog by a patient with a chronic condition.

One of the first attempts to conceptualize health social media was made by Dr. Gunther Eysenbach with the creation of the Medicine 2.0 term (see Figure 6). His definition of Medicine 2.0 lays emphasis on the new principles boosted by social media, such as enhanced participation, openness, collaboration and apomediation[49]. The concepts surrounding health social media, Medicine 2.0 and Health 2.0 are in continuous evolution

as studied by Van De Belt et al[11].The rapid evolution of social media technologies is one of the main challenges for research in this area. Conversely, that rapid evolution and adoption emphasize the need for more research in the area.
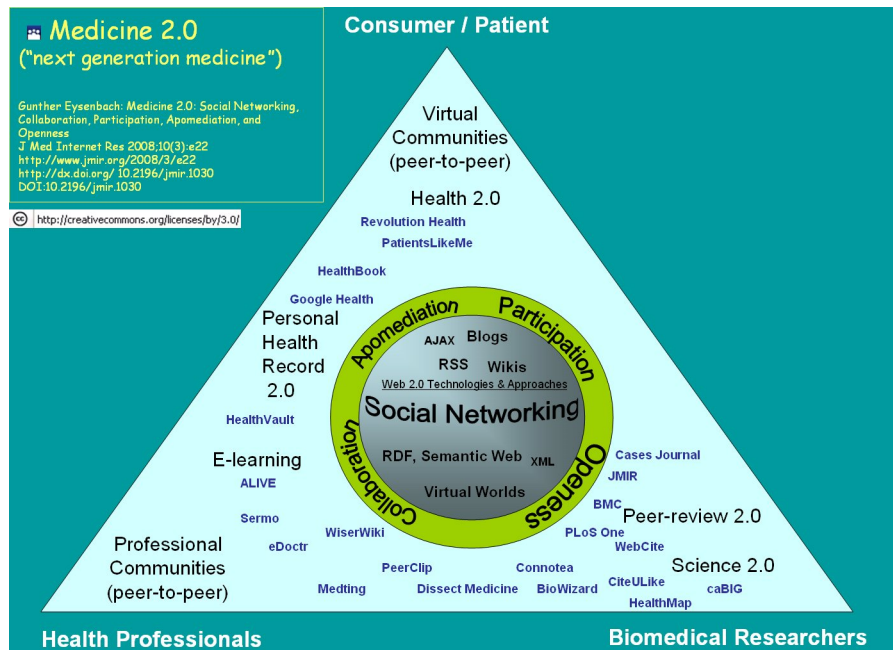
Figure 6: Diagram of Medicine 2.0: Social Networking, Collaboration, Participation, Apomediation and Openness[49]

Although it is virtually impossible to categorize health social media, some attempts have been made based on the functionalities currently supported by most social media tools. The Social Media Working Group of the International Medical Informatics Association (official website: http://imiasocialmedia.wordpress.com/) stratified social media tools across eleven categories: 1) Social Networking (e.g., Facebook, MySpace); 2) Professional Networking (e.g., LinkedIn), 3) Microblogs (e.g., Twitter), 4) Blogs (e.g., Wordpress, Blogger); Forums/Listservs (e.g., Google Groups, Equidad); Photo Sharing (e.g., Flckr, Picasa); Thematic Networks (e.g., Sermo); 10) Collaborative filtering (e.g., Delicious) and 11) Others (e.g., Second Life, Social Games).

3.1.3. Retrieving Health Videos in YouTube

Founded in 2005, YouTube is the leading free video-sharing site that allows people and organizations to find, view and share videos. In the USA, watching online videos has been found to be one of the main activities of Internet users[105]. Online videos are often

embedded within a social network, as shown in Figure 7. YouTube is a social network and community providing new opportunities to connect, collaborate, create and disseminate videos. Users can add comments (either textual or with a video) to somebody else's video, therefore creating a dialog between users and a community around the video content.
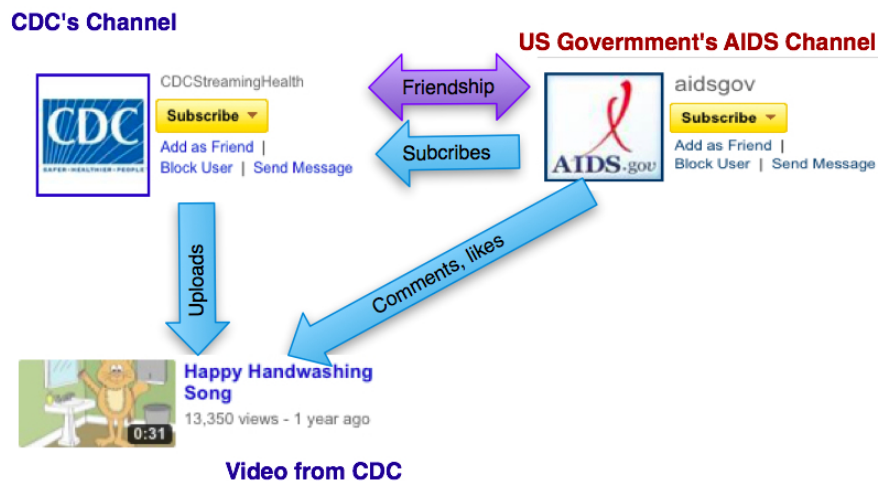
Figure 7: Example of YouTube Social Network (RQ4.Paper2)

Videos on YouTube do not just consist of the video itself, they are enriched by a set of metadata. The metadata includes, but is not limited to, description, subtitles, authorship, duration, views, ratings and comments. In addition, the YouTube API provides access to a wider range of parameters of a video or its creator, such as transcoding information and geolocation. YouTube is gradually adding more technical capabilities to the platform such as the ability to create 3D videos and also interactive videos with elements that can be clickable.

YouTube is gradually becoming a major channel for the dissemination of health education and it is also widely used as a communication channel by patients[25] and healthcare providers[80]. Publishing health videos on YouTube is not an easy task and some institutions such as the Centers for Disease Control and Prevention (CDC) have developed a guideline for publishing on YouTube [21]. One of the main characteristics of YouTube is that anyone can publish regardless of their message or intention. Therefore, concerns about the quality of health information on YouTube are common. YouTube has also become a major channel for the dissemination of pseudo-scientific scams, such as anti-vaccination theories[2], and even harmful content such as the promotion of anorexia (RQ1.Paper5).

Information Retrieval in YouTube

YouTube is a major example of the need for information retrieval tools because finding content in such a vast repository is difficult without automatic tools. The most common ways of finding new videos on YouTube are 1) search engines and 2) recommended videos (aka related videos). Little is known about YouTube search, but in a recent paper the mechanisms used for video recommendation have been explained [33]. As described by Davidson et al.[33], there are basically two types of data to be considered in the information retrieval process: i) content data which includes information about the video itself and its metadata and ii) user interactions which can be divided into explicit (e.g. rating a video) or implicit (e.g. watching a video). In all cases, the quality of the data is very heterogeneous since there may be almost no metadata or user interaction and watching does not always mean one likes the content. In short, in order to give recommendations the system estimates which videos are often co-watched in the same user sessions. The related videos are then ranked based on video quality, user preferences and diversification. The video quality is measured taking into account views, ratings, comments, favoriting, sharing activity and upload time.

The search algorithms implemented on YouTube search have not been published. However, in a panel the search leader at YouTube said that they are mainly based on the relevance algorithms used by Google [160], which are enhancements from the original PageRank. In the PageRank search algorithm (explained in more detail in Section 6.1), the more income links a website has, the higher the endorsement and PageRank of the website. In the case of YouTube, we expect user interaction/popularity to be the main source of data used for calculating trust-based scores.

The design of YouTube search and related videos explains why there is a concern about the quality of retrieved videos in the health domain. A health video that is attractive for an audience will tend to be highly ranked regardless of why it is "attractive". Gruesome health videos have a high ranking because sadly many users find bloody and nasty videos interesting. For example, in a video about a diabetic foot infection one often sees comments such as "I liked the 'crunch' part", meaning the part where the toe bone is amputated. Similar problems occur with health videos with little informative value that are highly popular because they feature a celebrity. The search and recommendation

algorithms designed by YouTube are designed to increase the views and engagement of users, not to retrieve the most valuable videos from a health point of view.

3.2. Study I: Social Motivations of patients video bloggers

Related paper: RQ1.Paper 1

One of the main innovations brought about by social media is the democratization of the generation of content. Nowadays, with a simple click on a mobile phone anybody can upload a video to YouTube or post a photo on a blog. The rise of user generated content is also common in the health domain, both for professionals and patients. To understand the motivations, challenges and social aspects that are faced by those individuals it is crucial to understand the overall picture of the phenomena.

Due to the lack of studies about the motivations of the creators of online health videos I interviewed several patients about their experience of sharing videos on YouTube. We aimed to explore their motivations and challenges in order to gather more knowledge about community insights into patient-led content creation. In a similar study, other researchers have been studying the main topics of videos created by patients with cancer[25].

3.2.1. Methods

The study design was based on qualitative methods. I searched YouTube for patients who were highly visible in the diabetes and multiple sclerosis communities. I then invited them to make a video about their experience of making videos about their disease on YouTube. Finally, four patients shared videos about their experience that were presented in the Medicine 2.0 2010 conference in Maastricht (The Netherlands). In that conference, several Internet psychologist researchers offered to perform a further analysis of the videos. We performed an analysis of the videos about their self-reported motivations and challenges. Two judges led the transcription of the videos which were further analyzed to categorize the content of the videos into major themes.

Table 4: ePatient videos about publishing on YouTube (RQ.Paper1)

YouTube Channel	Age	Sex	Disease	Video link
Vbeachy (VB)	47	Male	Multiple sclerosis	http://www.youtube.com/watch?v=CuOVDAm6tlE
1HappyDiabetic (BW)	31	Male	Diabetes	http://www.youtube.com/watch?v=jsP9nZXIpME
Laurenvparrot (LP)	29	Female	Multiple sclerosis	http://www.youtube.com/watch?v=Atq0LP_mfYE
Sixuntilme (KS)	32	Female	Diabetes	http://www.youtube.com/watch?v=lc9hA4gmlhw

3.2.2. Results

The videos addressed four major topics: the reason for making the first videos, the objectives that they achieved, community aspects and negative consequences.

Overall, the initial motivation was to fill a gap in the content available on YouTube. Patients described the lack of content from a patient's point of view. One of the patients said that he felt a lack of content depicting diabetes from a positive viewpoint. Another patient described the lack of videos about patients with multiple sclerosis suffering a relapse, she decided to share her experience so others could be spared the anxiety of not knowing what to expect. This highlights that the content generated by patients normally complements that generated by professionals or institutions. The next quotation shows that the lack of information was the main reason for starting:

"I was actually not a habitual user of YouTube videos and wasn't making videos on YouTube before I started doing diabetic videos, but I was actually looking for information and diabetes in a video format and there was nothing there.[BW]"

A common reason among all the patients for continuing posting was the support from the community. Some patients even described making friends thanks to YouTube, feeling less alone. They also mentioned that "followers" kept asking about the health of the video makers or even suggested new topics. The creators´ awareness of the community was crucial to sustain quality due to the feeling of being watched by their communities. For that reason, one of the patients had systematically refused to support any commercial offer to endorse a product since that could destroy his reputation. One ePatient even reported that the need for socialization with peers was the main motivation for starting to post videos:

"...And part of why I started my blog in the first place was because, even though I've lived with diabetes for such a long time, I didn't know (sic) anyone else who had it, and I literally felt like the only diabetic on the planet. [KS]"

Among the negative consequences described by the video creators they mentioned the disclosure of private information that could worry relatives among other things. Another negative aspect was the perception of criticism and even insults by some of the viewers.

3.2.3. Discussion

Patients' main reason for making videos was to cover aspects and angles not commonly addressed by other video makers such as health authorities. The role of the community was omnipresent in all the patients who were endorsed by their viewers and their awareness of the need to maintain their online reputation was also clear.

In relation to the questions addressed in this dissertation, the findings in the study, although preliminary due to the small sample size, support the approach of inferring community support as a crucial part of the trust concerning content. That trust is built slowly via multiple interactions with other peers via comments, videos, etc. ePatients are fully aware of the need to maintain and sustain the quality of their videos to avoid damaging trust from the community.

3.3. RQ1.Study 2: Metadata of Health videos

Related papers: RQ1.Paper 2, RQ1.Paper 3

This study focused on the characterization of the features related to health videos. Web information retrieval is normally based on the textual analysis of metadata such as descriptions, tags and comments. The lack of previous studies characterizing metadata in health videos was one of the key challenges of my dissertation. This study was divided into two different experiments: one characterizing the disclosure of private information in comments on health-related videos (RQ1.Paper2), and the use of standardized medical vocabulary in health video tags (RQ1.Paper 3). We selected two case studies of the disclosure of private health information videos by patients with multiple sclerosis, while we opted for surgery videos for the study of the use of medical vocabulary.

3.3.1. Private Health Information Comments on Multiple Sclerosis Videos

The RQ1.Paper2 describes this study on the comments made on social videos about multiple sclerosis. As mentioned earlier, YouTube is a social network where users can enrich somebody else's videos with comments. In these comments, people can reflect on their own experiences and support the videos, etc. The importance of comments for my dissertation lies in the ability to use comments for the modeling of either users or content. For example, I observed that sometimes the topic of videos without a description could be inferred by reading the comments of the viewers who in many cases shared the same health issues as the maker of the video. In order to gather some real data to support this observation we decided to explore the characteristics of comments posted on patients' videos about multiple sclerosis. To my knowledge, this is the sole study on comments posted on health-related videos.

3.3.1.1. Methods

During the first week of December 2008 we searched for YouTube users who had published at least three videos in English about living with multiple sclerosis. In total, we found 27 such users. I developed a web crawler that was used to extract a total of 769 videos, 7,047 comments generated by 2,426 users. Using a random method we selected 25 of those videos with their 557 comments. Comments that were not in English or posted in non health-related videos were excluded. A final set of 320 comments were analyzed and classified as follows:

- Personal health information: comments containing personal health experiences such as diagnosis, symptoms, etc.
- Video discussion: comments with discussions about the videos (e.g. adding information about the videos).
- Appreciation: appreciations from the commenters towards the video author.
- Criticism: complaints about the quality of the video or any other aspect.
- Unrelated: comments not covered in any of the other categories (e.g. comments about the hair-style)

Two of the study's authors performed the data analysis and in the event of discrepancies consensus was reached.

3.3.1.2. Results

The majority of the comments were classified in several categories (see Figure 8). For example, many comments discussing the video also added personal health information. Nearly three quarters of the comments discussed the videos (77%, n=247) and the majority of the comments showed gratitude (55%, n=177). Twenty percent of the comments contained personal health information; of those, the majority disclosed medications and symptoms. Many of the comments contained detailed discussions of the videos enriched by descriptions from personal experiences. The comment in Figure 9 is a good example.
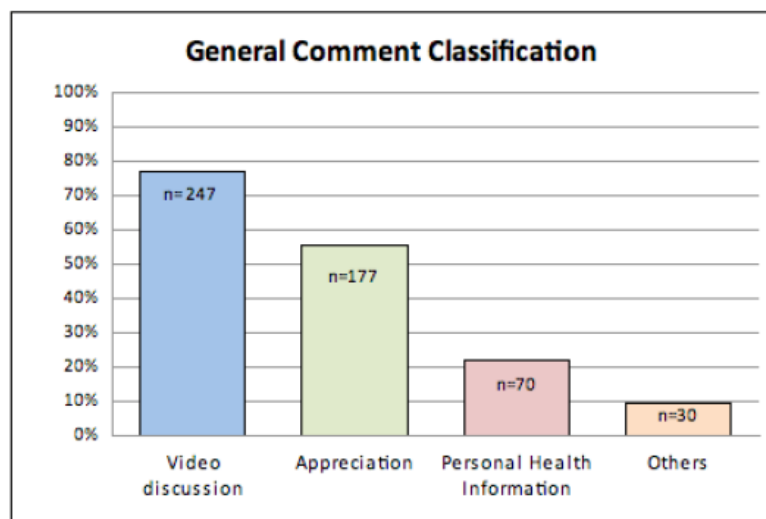
Figure 8: Total numbers of comments in ePatients' videos classified into the main categories (RQ1.Paper2)
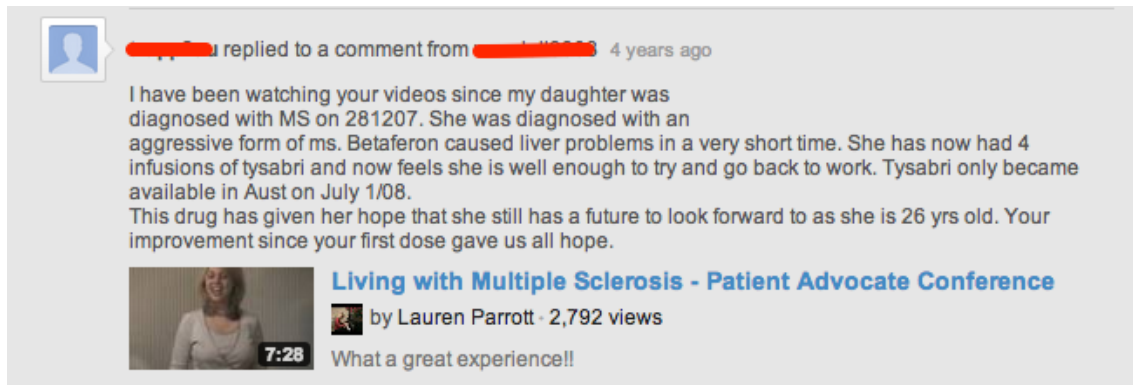
Figure 9: Example of a comment with personal health information on a video (RQ1.Paper2)

3.3.1.3. Discussion

In this preliminary study, we found that around 20% of the comments could eventually be used to model the content of the video or the user because it contained personal health information. In addition, many of the comments discussing the video could also provide clues about the topic of the video and therefore be used for modeling. However, that approach would face several major challenges such as the complexity of dealing with natural language. Comments often use acronyms and more informal language (e..g. Ty for the drug Tysabri). The use of semantic technologies to analyze health consumer text has been researched and it is not a trivial approach[88,144]. Special caution needs to be taken with the extrapolation of the results due to the small sample used: 25 videos (with over 500 comments).

There is one major limitation for the use of user-generated comments in the modeling of users and videos that seems insurmountable. The comments contain highly sensitive and private information. Although there are ways of protecting privacy while modeling users or content [91], there would undoubtedly be user resistance to the possibility of extracting information without their consent.

3.3.2. Semantics and YouTube Medical Surgery Videos

Online videos are enriched by a wide set of metadata as described in Figure 13. One of the most common types of metadata in social media content is the use of tags. YouTube also uses tags to facilitate the indexing of videos. The importance of tags for modeling social media is very important and it has been widely researched [64,131,168], including in the health

domain[66,89,144]. The extracting of semantic information from tags can help to find similar content [64].

In the health domain, semantic technologies can help to translate layperson vocabulary into a medical vocabulary and vice-versa [88]. A wide range of semantic technologies have been used to model clinical documents, facilitate interoperability, etc. In fact, there are well-known standardized medical thesauri, such as SNOMED CT, which provide descriptions of more than 311,000 medical related concepts.

To gain a better understanding of the metadata of medical videos, with some fellow researchers I designed a study about the use of SNOMED CT terminology as tags describing health videos. That study was published and it is RQ1.Paper 4 of this dissertation. My objectives for this study were: 1) to identify the use of standardized medical thesauri in YouTube Health video tags; and 2) to find out how this could be used to facilitate information discovery within online health videos.

3.3.2.1. Methods

We implemented a video portal based on the prototypes developed during my PhD and also the prototypes developed within the European project m-Educator. In the video portal we integrated videos from more than 500 US Hospitals in a list curated by Ed Bennett [13]. We included videos with the word *surgery* in their title, description or tags. In total, we extracted 4,307 videos with 64,367 tags, of which 7,798 were unique tags. The BioPortal REST services were used to match the tags to SNOMED CT terms by both exact match and non-exact match. Figure 10 describes in detail the architecture of the system developed for that process, where I implemented the crawler for the YouTube metadata extraction.

*Linking Open Data cloud diagram, by Richard Cyganiak and Anja Jentzsch. http://lod-cloud.net/ under CC-BY-SA
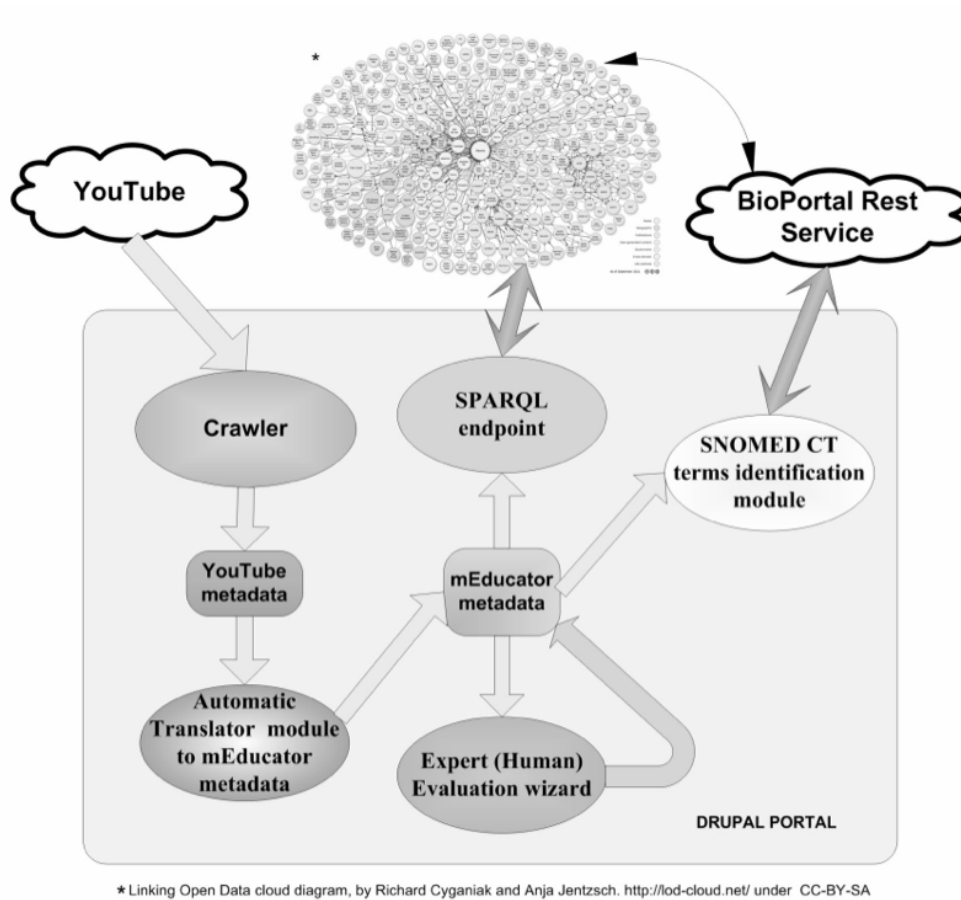
Figure 10: YouTube Health-videos Drupal portal architecture (P9)

3.3.2.2. Results

The videos had on average 16 tags each. As shown in Figure 11, of the 7,798 unique tags found, SNOMED CT terms were present in 20.6% without "exact-match" and 4.7% with "exact-match". Taking into account the total number of tags (n=64,367), 37.6% of them were SNOMED CT non-exact matches, while 7.4% were exact matches. These differences between unique tags and overlapped tags can be explained by the common repetition of SNOMED CT tags across videos. We also found that 17.7% of the videos contained no SNOMED CT tags whatsoever. In those videos with at least one SNOMED CT tag, the average percentage was 41.7%.
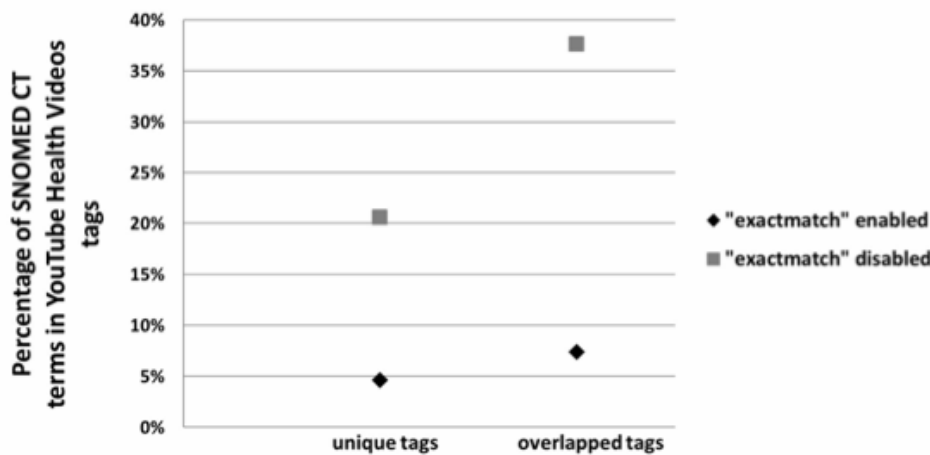
Figure 11. Average percentage of SNOMED CT terms in YouTube health video tags (P9).

3.3.2.3. Discussion

The presence of medical thesauri on the tags describing the videos was modest but significant. The possible combination of tags with other types of videos' textual metadata (e.g. comments, titles, description) could increase the percentage of SNOMED CT terms associated to a video and be used to index or facilitate information retrieval. SNOMED CT terms could be used to enrich videos with additional information such as scientific publications, health dictionaries, the explanation of terms, etc.

However, using semantic information from tags for information retrieval will have to contend with several obstacles. Many videos did not have terms, even in this dataset, which only contained videos from American hospitals, which in general are more likely to use medical vocabulary. The use of additional textual metadata (e.g. descriptions and comments) will also be a challenge due to the language gap between medical vocabulary and layperson language which is more common in social media. This issue has been already addressed by researchers working towards the definition of Consumer Health Vocabulary [172] and it has also been studied in patients' forums such as PatientsLikeMe[144]. We have started to explore the use of semantics derived from tags within health video recommender systems [134].

3.4. RQ1. Study 3: Quality Aspects of Health Videos on YouTube

Related paper: RQ1.Paper 4

YouTube has become the biggest repository of online videos and it is a social network where people can find, view, share and comment videos[22,105,171]. There are over 100 million videos on YouTube and many of them are health related, published both by patients[25]and healthcare organizations[21]. Some recent studies have explored the use of YouTube for health promotion[7,34]and education[50,146]. However, the nature of YouTube also makes it perfect for disseminating low quality health videos and even videos promoting unhealthy behaviors that can cause harm[95].

As explained before, the quality of health-related videos on YouTube can be very heterogeneous. In this project one of the main issues I faced was how to find high quality social media. The first challenge I met in my research was to establish what "good" social media was because different stakeholders have completely different opinions. When I asked some patients, they said that they preferred videos that were informative, funny and above-all sensitive to their situation. Other patients (and many doctors) told me that they preferred videos by authoritative authors (e.g. medical doctors). Video experts worried about the sound, images, while computer scientists always complained about bad metadata. I teamed up with several researchers to address the important issue of identifying the quality parameters that have been considered for health videos on YouTube. In our study presented here, we performed a systematic review of the literature on YouTube health videos aimed at identifying quality parameters reported in previous studies.

We normally consider quality from a more general point of view; however, the quality of health information is highly dependent on the viewer. For example, a video about the nutritional advantage of using whole sugar will be irrelevant and even dangerous for somebody with Diabetes Type I. As Purcell et al mentioned, "the quality of information, like beauty, is in the eye of the beholder, and it is users' views we should be seeking"[129]. The quality of a health video will also depend on its relevance for the viewer.

3.4.1. Methods

A detailed description of this study can be found in RQ1.Paper 4. The methodology employed in the study was an adaptation of the systematic review approach PRISMA[100]. The objective was to gather information about the concept of quality information for patient education on YouTube. The search was performed using the knowledge databases MEDLINE, ISI Web of Knowledge, Scopus, and PsychINFO.

As shown in Figure 12, the study selection and data extraction were performed in several phases. Two authors independently reviewed all the references (n=456) to identify those which were relevant and excluding those: 1) where the scope was not YouTube and/or 2) the concept of quality was absent and/or 3) an abstract was missing and/or 4) duplicated. In the event of discrepancy about an abstract all the authors reviewed the abstract until consensus was reached. In total 13 abstracts were selected, the inter-rater reliability for this first review was found to be Kappa=0.73 (P<.001), 95% CI (0.662-0.792), and considered "moderate" [21]. The 13 abstracts selected in the first round were analyzed by two additional independent reviewers who classified them as included or excluded using the same pre-determined criteria. After this second abstract review, 4 references were considered for inclusion by two reviewers, 5 references by three reviewers, and 4 references by all four reviewers. Overall, the 13 abstracts that were selected by at least two reviewers were incorporated for full text analysis. After that, the full-texts of the studies with at least two reviewers were retrieved for data extraction of the concepts and definitions regarding the quality of YouTube videos for patient education.
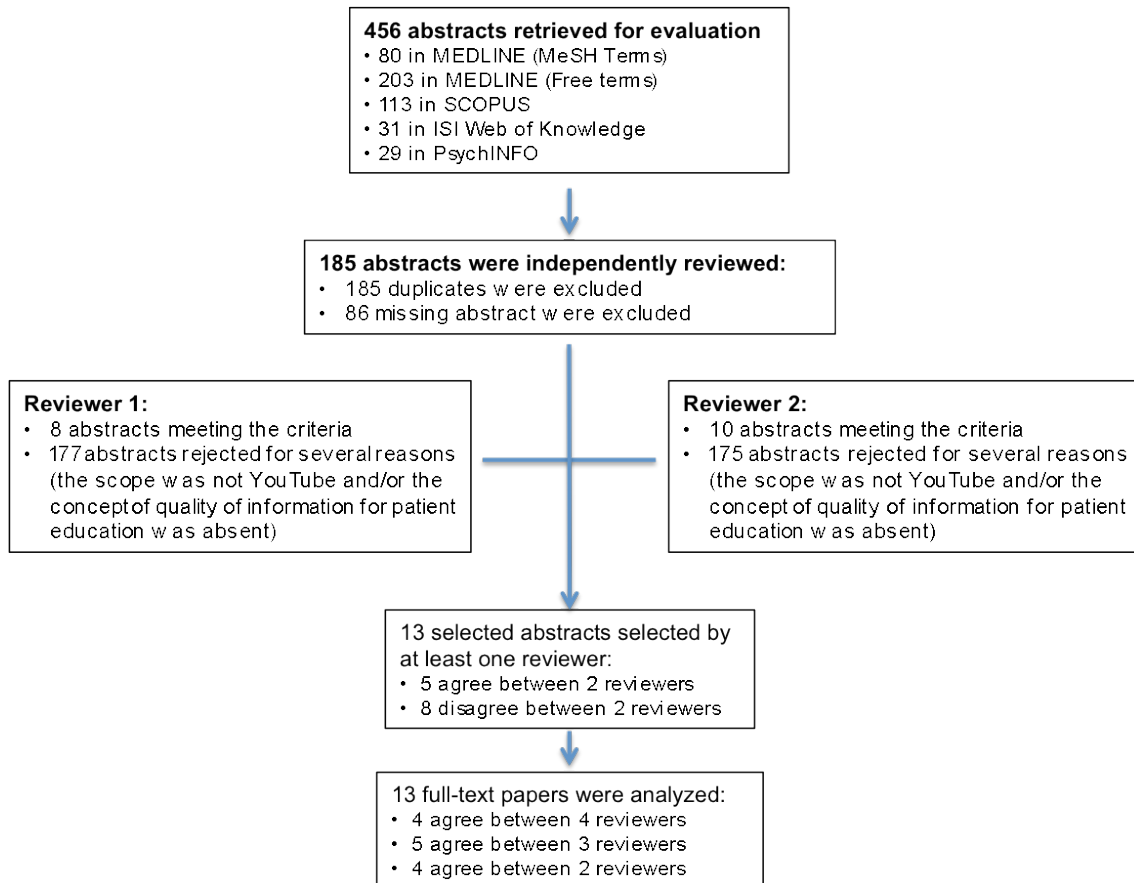
Figure 12: Study design of "Identifying Measures Used for Assessing Quality of YouTube Videos with Patient Health Information: A Review of Current Literature" (RQ1.Paper 4)

3.4.2. Results

In our review of the selected articles we identified three main categories for quality criteria: expert-driven; popularity driven; or heuristic-driven. Expert driven measures were those related to the perspective of health professionals (e.g., validity of the content from a clinical point of view). Measures related to the popularity of the videos (e.g. view count) were also commonly mentioned and they were included in the popularity-driven category. Less common was the use of heuristic-driven measures based on the metadata of the video (e.g. proper description of the videos).

Table 5 summarizes the different quality measures found in our review.

Table 5: Quality metrics associated to health videos on YouTube (RQ1.Paper 4)

Quality metrics	Frequency N=13 (%)
Quality content (includes accuracy-credibility of content; scientifically correct information and/or evidence-based practices)	10 (77%)
View Count / Popularity	9 (69%)
Rated by expert (medical staff)	8 (62%)
Adequate length / Duration	6 (46%)
Public ratings	5 (39%)
Good description of the video/ comprehensive narrative provided	4 (31%)
Technical quality (light, sound, angle, resolution)	4 (31%)
Further contact info provided / Credentials	4 (31%)
Suitability as a teaching tool	4 (31%)
Comments (by viewers)	4 (31%)
Title and Tags	3 (23%)
Amount of content / enough information to identify its objective	3 (23%)
Viewership share (number of links to the video and/or number of shares in other social media)	2 (15%)
Description of video	2 (15%)
Health professional and patient/s seen in video	1 (8%)
Mention intended target audience	1 (8%)
Judgment includes patients/parents/users	1 (8%)

Expert-driven measures

The most common approach to evaluate the quality of a YouTube health video was the manual assessment of one or more experts (8 of 13 publications, 62%). The experts were mainly health professionals and other researchers [7,34,50,57,67,119,122,140,146,148,155], but surprisingly the assessment by patients was only found in one study (8%). In this category, the quality concept referred to the 1) accuracy-credibility of the content, 2) scientific correctness of the information, and 3) evidence-based practice. These aspects

are normally analyzed by the active watching of the whole content of the video and the review of the metadata (e.g. description). This is a highly complex task because medical information does not just have to be accurate but also up to date. That is why in many cases the evaluation was not made by one single expert but by a panel of experts.

Popularity-driven measures

The popularity of the videos was also considered related to its quality in most of the papers (9 papers, 69%)[4,7,34,50,67,119,122,146,155]. The rationale for linking popularity with quality was that popularity might be the result of the attractiveness of the video (e.g. enjoyable, funny, lay person language). Popularity was often measured by: 1) the mean daily views since the video was posted, 2) public ratings provided by YouTube and 3) the number of shares of the videos (e.g. sharing in FaceBook).
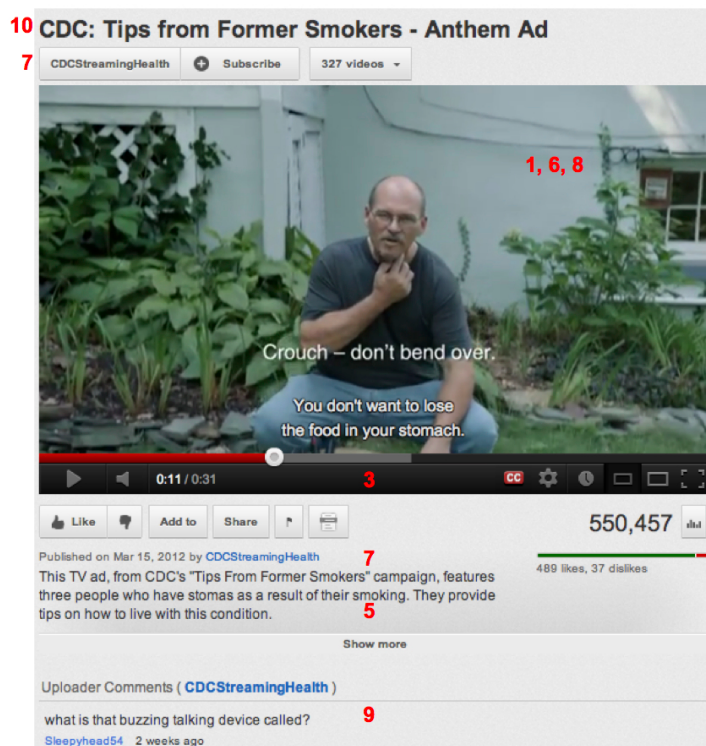
Heuristic-Driven Measures

Figure 13: Examples of Metadata used for quality evaluation (RQ1.Paper 4)

As Figure 13 shows, the metadata and other attributes of the video were also considered to assess quality. For example, technical parameters such as the length of the videos were

considered relevant to assess quality in many papers [67,119,122,146,148,155]. Other technical parameters included appropriate light, sounds, angle and resolution[50,67,148,155]. The quality of the metadata (e.g. title, tags and description) was also a common quality parameter[4,57,67]. The metadata was assessed by: (1) good description and comprehensive narrative [4,57,67], (2)evidence-based[34,122,140,146], (3) suitability as a teaching tool[50,122,140,155], (4) credentials or contact information provided in the video[67,119,140], (5) amount of content or the presence of enough information and ability to identify its objective [4,57].

3.4.3. Discussion

In this discussion, I shall focus on how those quality aspects can be taken into consideration to build new technical solutions for the retrieval of high quality health videos from social networks such as YouTube.

Quality is no simple aspect and it involves multiple factors. In fact, there is a lack of consensus in the scientific literature over which should be the main criteria to assess the quality of an online health video. The majority of papers focus on the validity of the content since it is well known that misleading health information can be harmful. A second level included all the quality measures related to the technical quality of the video and its popularity among the public.

The assessment of quality based on **expert-driven measurements** presents a major problem for application on a wide scale. As the volume of online videos grows exponentially (72 hours of video uploaded every minute[171]), using only expert evaluation to assess the quality of all videos posted on YouTube is not a realistic or sustainable long-term solution. An alternative way to evaluate the credibility of content is to infer it from the creator of the video, since an expert-based review of the content's creators will be less time consuming.

The use of **popularity-driven measurements** was described as highly relevant and it is very easy to integrate in information retrieval tools. However, great care must be taken because the study was limited to 13 papers. For example, we did not find any paper addressing the quality of videos about sexual health or other topics (e.g. eating disorders) where it might be plausible for popularity to be related more to polemical content or even graphical content. Popularity is also prone to manipulation (e.g. advertising campaigns) and it is not well understood why some videos become viral. Popularity could be a

parameter to be taken into consideration as a measurement of the likability of a video by the audience. The problem lies in determining whether the audience stems from its quality from a health point of view or something else. It has also been demonstrated that online crowd influence can potentially lead consumers to make unsafe health decisions[95,96].

The use of **heuristic-driven measurements**, mainly based on the quality of the metadata, is an important aspect in many of the reviewed papers. It has been shown that poor quality in those aspects will result in fewer views, either by making the video less attractive or harder to find. However, once again, there is room for manipulation of those aspects. In many cases, I have found videos promoting misleading health information (RQ1.Paper 5) which were of very good quality in terms of metadata and vice-versa.

3.5. RQ1.Study 4: Popularity and quality of anorexia-related videos

Related paper: RQ1.Paper 5

The number of videos disseminated over the Internet has reached unprecedented levels due in particular to platforms such as YouTube and Vimeo. The native digital generation (i.e. teenagers and young adults) is more accustomed to YouTube than traditional TV channels. That population faces an increased risk of finding unreliable health information when searching for health content which is more relevant to their age (e.g. sexual health and body image). A very interesting case is the rise of social networks promoting anorexia as a lifestyle choice; this is a subject which has been documented for nearly a decade[117]. This is a big societal problem, since a study in 2009 found that ca. 10% of teenage girls in Belgium had searched for pro-anorexia content[31].

I collaborated in a study (RQ1.Paper5) with colleagues from Taiwan and the USA where we studied the presence of pro-anorexia videos on YouTube; see Figure 14 for an example. My interest in the study was to explore the prevalence of such types of videos and the similarities between informative content and misleading videos (e.g. tags, descriptions, popularity).

Figure 14: Example of a pro-anorexia video (RQ1.Paper 5)

3.5.1. Methods

As shown in Figure 15, we retrieved anorexia-related videos from YouTube using a set of keywords that included terms related to anorexia, including those traditionally used by the pro-anorexia community. The videos were extracted during October 2011 using a Java crawler I developed using the YouTube API. I retrieved 4,000 videos for each query (thinspo, proana, anorexia and anorexia nervosa) and sorting criteria (relevance, view count, rating and date of publishing). In total, I retrieved 16,000 search results that contained 7,583 videos from 3,968 users (aka channels).
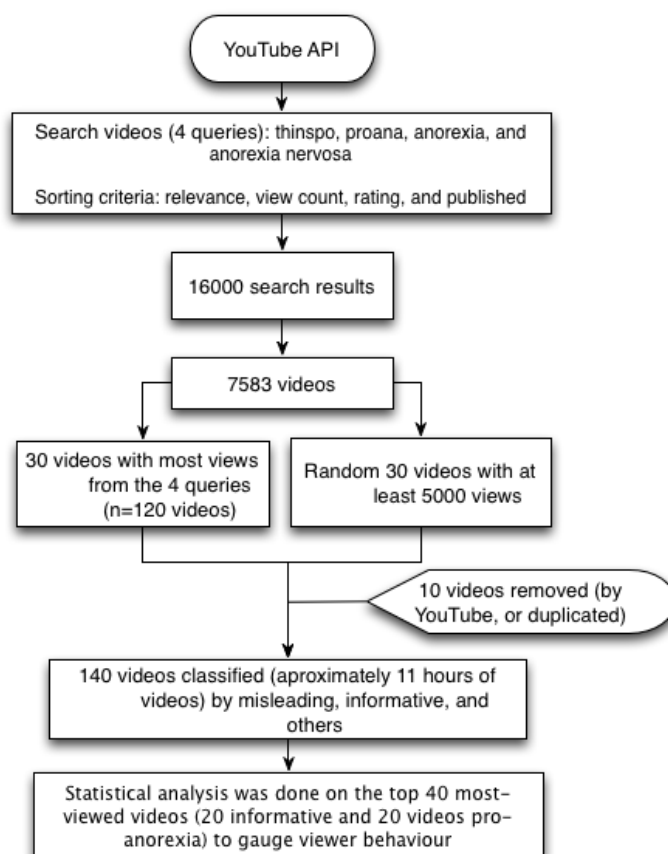
Figure 15. Study design of RQ1.Study 4 (RQ1.Paper 5)

I further selected the 30 most-viewed videos from each query (n=120) and a subset of 30 random videos with at least 5,000 views. The random selection explored the overall prevalence of pro-anorexia videos. Of the final sample of 150 videos, 8 were duplicates and 2 were removed from YouTube during the analysis (i.e. the video creator may have decided to remove it). The final set of 140 videos contained around 11 hours in total.

Three medical doctors independently reviewed those videos in order to classify them as pro-anorexia, informative or unrelated. Informative videos were those describing the health consequences of anorexia and how to cope with it, while pro-anorexia videos were those promoting anorexia as a fashion, a source of beauty and providing tips and methods for extreme weight loss. As explained below, the inter-rater agreement was measured using Fleiss Kappa [58] and, in the event of disagreement, consensus was used to classify those videos.

In addition, we analyzed the 40 most viewed videos (20 pro-anorexia and 20 informative) to compare their features. These top videos were not necessarily the ones shown first when searching for videos, since many users sort search results by relevance and not views. These videos had a total of 61.13 million views and good examples to determine the differences between the categories. The statistics software SPSS v17 was used to analyze the different features.

3.5.2. Results

The inter-rater agreement was considered moderate (Fleiss' kappa=0.5) [58]. Of the 140 videos, 41 (29.3 %) were classified as pro-anorexia, 78 (55.7%) as informative and 21 (15%) as unrelated (see Table 6). These percentages were similar in the random selection of videos. If those percentages could be extrapolated to the entire sample of 7,583 videos we could expect YouTube to contain more than 2,000 videos promoting anorexia.

Table 6: Results of the classification of anorexia-related videos on YouTube (RQ1.Paper 5)

Video selection	Total, n	Pro-anorexia, n (%)	Informative	Others
Top 30 videos with most views for each query	110	32 (29%)	61 (55.5%)	17 (15.5%)
Random selection with more than 5,000 views	30	10 (33%)	17 (57%)	3 (10%)
Total reviewed videos	140	41 (29.3%)	78 (56.7%)	21 (15%)

In addition, we analyzed the 40-most viewed videos (20 pro-anorexia and 20 informative). These videos are of particular importance since most users only browse the first pages of results. In absolute terms the informative videos had more views, nearly 5 times more. However, in terms of ratings there was a similar ratio between likes and dislikes. Surprisingly, the pro-anorexia videos were favored 3 times more than the informative videos (odds ratio [OR] 3.3, 95% CI 3.3-3.4, P<.001) after adjusting for number of views. Also, there was twice as much viewer engagement in pro-anorexia videos, calculated by

taking into account how many like-dislikes a video had per view. Table 7 summarizes those results.

Table 7: Assessment of the 20 most-viewed anorexia-related videos on YouTube (RQ1.Paper 5)

Features		Pro-anorexia	%	Informative	%	OR (95% CI)	P value
Total views		9,51 million	100	51,62 million	100		
Favorite		24,462	0,26	39,424	0,08	3,37 (3.32-3.43)	<.001
Total Likes/Dislikes		15,209	0,16	45,486	0,09	1,82(1,78-1,85)	<.001
	Likes	12,506	82,58	40,332	88,67	0,61 (0,58-0,64)	<.001
	Dislikes	2,649	17,42	5,154	11,33	1,65 (1,57-1,74)	<.001

Pro-anorexia videos commonly featured extremely thin female models, with the few exceptions of a video featuring boys and another with overweight models (also referred to as reverse thinspo). In many cases those videos included "inspirational" quotes and tips for losing weight. For example, Figure 16 shows a video with a pro-anorexia nutritional pyramid where smoking and taking diuretic drugs are featured as recommendations for losing weight. Although not studied in detail in this paper, we found that pro-anorexia videos contained tags related to anorexia nervosa, weight loss, diet, fashion, etc.

Figure 16 Pro-anorexia video with misleading diet advice (RQ1.Paper5)

Figure 17. Example of Pro-anorexia video with self-portrait (RQ1.Paper 5)

The informative videos came from a wide range of angles, from TV reports about the disease from major news agencies (such as CBS), to health organizations and individuals recovering from the disease. The vocabulary changed a lot depending on the video creator, however terms related to weight loss, diet and fashion were not common. The videos found that were not related to anorexia were mainly from a heavy metal music band called Anorexia.

Figure 18. Example of informative anorexia video (RQ1.Paper5)

Some of the videos included demographic information. We analyzed the demographic information available in 15 pro-anorexia videos. Of those videos, 80% (n=12) had minors (13-17 years) in the top-3 age group as viewers, even though one-third of the videos had age restrictions. The only reasonable explanation for that is that those videos were highly popular among younger viewers before they were flagged as inappropriate for minors.

3.5.3. Discussion

In this study we found that some characteristics often considered relevant for inferring the quality of health videos (e.g. ratings, popularity) are in fact misleading. Therefore, relying solely on metrics related to popularity to find trustworthy health videos could be very challenging in scenarios where there is "popular" demand for those misleading videos. The problem of misleading information is not insignificant in anorexia since it represents nearly one third of the videos.

4. Modeling Health Social Media

> **Related paper:** RQ2.Paper 1
>
> **RQ1**: **Are there any techniques for extracting and modeling health social media?**

4.1. RQ2.Study I: Literature Review of Modeling Health Social Media

In parallel with the study of the context of health social videos, described in the previous section, I decided to perform a wide review of the different techniques that might be used to extract and model health social media content and users. Many of those techniques could eventually be used to facilitate the retrieval of high-quality health social media content. However, I faced the challenge of the lack of reviews about the use of those techniques in the context of health social media.

I decided to explore 1) user and content modeling for health social media, 2) personalization and tailored health education and 3) community analysis of health social networks. These topics were selected for the following reasons: 1) User and document modeling is a crucial part of any information retrieval system, including future systems or what is already implemented in YouTube or Google. 2) Since health information seeking is rather personal, the study of tailored health education could provide insights into possible methods to support personalized information seeking. 3) In addition, the study of social networks is central to understanding the dynamics of health social media.

It is necessary to model users and social media for the retrieval of social media adapted to users with particular information needs. Information retrieval techniques rely on the need to extract knowledge from different data sources for the indexing of the social media and also about the information needs of users. The extraction of knowledge about online documents and users is defined as document and user modeling. Modeling is not only used for information retrieval, but also for building systems that are personalized.

As we describe below in Table 8, there is a wide range of data sources and techniques to model users and social media. These are relevant research areas to take into consideration for the retrieval of health social media. The different research areas can be divided into more general computer science and health informatics areas.

Table 8: Relevant research areas for the extraction of information from health social networks and personalization (RQ2.Paper 1)

Research Area	Importance
Health Informatics: Tailored Health education [102]	Modeling of users and educational materials to provide personalized education.
Health Informatics: Personal Health Records [83,153]	Modeling of user needs based on health information about the users.
Health Informatics: Biomedical text mining [78]	Extraction of information about users and documents based on textual analysis.
Health Informatics: Consumer Health Vocabulary[172]	Study of the vocabulary used by laypersons when searching for or commenting on health information
Health Informatics: Computer-aided diagnostics[125]	Analysis of text, audio and video to diagnose health problems, which can be used for user and document modeling.
Computer Science: User modeling and personalization[5]	Adaptation of web systems to users and document modeling.
Computer Science: Computer vision[94]	Extraction of information for modeling from videos and photos.
Computer Science: Affective Computing[20,161]	Extraction of information about user emotions and social clues.
Computer Science: Collaborative computing[141]	To extract information from the community of users (e.g. tagging, ratings)
Computer Science: Web data mining[65]	Techniques to analyze web content (e.g. link analysis).

4.1.1. Methods

We designed this review to identify relevant research literature that addressed aspects related to health personalization in the social web, mainly related to user and document modeling.

The search for relevant literature took place in major scientific databases in computer science (e.g. ACM digital Library) and biomedicine (e.g., PubMed). In addition, we searched through the references of the selected papers, conferences (e.g. UMAP) and grey literature. The search was performed in the different relevant research areas described in Table 8.

The team of authors, under my leadership, performed the selection of relevant articles. The different backgrounds (i.e. computer science, medical informatics and public health) of the authors were crucial to ensure the implications of the research were fully understood.

4.1.2. Results

RQ2.Paper 2 provides a detailed description of all the approaches for the extraction of information from the health social web in a review paper containing more than 100 references. In this section, I will only focus on the main findings, which are summarized in Table 8.

Modeling users is crucial in order to personalize health education and also retrieve social media. In fact, search engines also model users when they enter a search query that represents a user information need. In the context of health social media, there is a wide range of possible data sources where information about content and users can be extracted (see

Table 9). The Social Web has facilitated the creation of a wide range of content (e.g. blogs, videos, comments, tags, user profiles). That content can be analyzed to extract clues about the personal traits of the users who posted the content. For example, it is common for users to disclose personal health information in user profiles and also comments (RQ1.Paper2). Furthermore, many other sources are available such as ratings, links and web-usage data. In addition, new types of web applications are starting to use information integrated in Personal Health Records (PHRs) to model users[53].

Table 9: Main sources of information from health social web

Source	Examples of relevant information for user and document modeling
Personal Health Records	Personal health information and demographics.[51,53,83,153]
Textual content	Textual analysis for modeling users (creators of content) and content (metadata).
User profiles	Demographic information, disclosed health risks and preferences.[15,101,111,116]
Forum posts and comments	Personal health information, emotional status, and categorization of content. [52,74,126]
Search queries	User interests.[130]
Tags	User interests and topics of the tagged content. [64]
Audio	Users' emotional and health status (e.g., depression). [98,118,150]
Facial photos	Emotions, gender and age.[30,94]
Videos	Diagnosis (e.g. Parkinson's), topics of the videos.[106,162]
Ratings	User preferences and similarity of content.[141]
Social Network and Links	Community discovery, reputation.[3,23,65,92]
Web usage data	Classification of users based on navigation patterns.[113]

4.1.3. Discussion

On the basis of the reviewed literature, there is no doubt about the potential of natural language processing for the modeling of health social media. However, advanced techniques for extracting relevant health information can be challenged by ethical issues, in particular privacy. Textual analysis will also face the challenge of the heterogonous character of medical vocabulary, including the use of acronyms. Furthermore, textual content can be prone to spamming issues. A less intrusive approach could be based on web-usage data mining (e.g. query analysis), but, even so, if somebody searches for sensitive information (e.g., abortion, erectile dysfunction) they might strongly oppose the use of such information to create a user model to improve information retrieval.

In the health domain, users can be modeled on the basis of Personal Health Records, which contain rich user profiles. However, their use is limited due to the low adoption and also concerns about the quality and quantity of information stored in such profiles. Audio and video analysis can even be used to diagnose diseases, but those techniques require complex processing which make them unsuitable for most information retrieval solutions.

Collaborative filtering represents an interesting option; the problem being the high ratings of bad content. Collaborative filtering can only be applied if we are certain that "harmful" sub-communities are not relevant. Furthermore, Social Network Analysis can be used to identify "harmful" sub-communities and avoid them when computing the trust within health communities.

5. Trust and Social Network Analysis of Health Communities

Related papers: RQ3.Paper1 and RQ3. Paper2

RQ3: How can Social Network Analysis be used to extract information about the characteristics of health social media?

- **RQ3.1**: Can social network analysis be used to infer the misleading nature of social photos in the case of anorexia?
- **RQ3.2**: Do the central users' online health communities have different features to the rest of users in the case of diabetes?

As the title of the dissertation emphasizes, one core aspect of my research is the study of how trust is built within health social networks. Specifically **RQ3**, which is covered in this section, addresses the study on **how social network analysis can be used to extract information about the characteristics of health social media**. In order to study that broad question I decided to focus on/analyze several case-studies. First of all, the case of online communities about anorexia was of special interest due to the presence of sub-communities promoting misinformation. Secondly, the case of diabetes was selected as a prime example of a chronic disease where online communities are thriving and very popular.

The reliability of health information is a crucial aspect to take into consideration when retrieving online content[42]. Misleading and harmful content coexists with educational content that can help people manage their health better. Trust in the creator and the content is one of the most complex and challenging quality features to take into account. In the off-line world, the health sector has implemented many safeguards, such as medical licenses, to ensure that only reliable people can provide health information. In the online world, the situation is far more complex since there are no boundaries and health quality codes are not well known by the public and they do not ensure quality [14].

As described in the previous chapter, a possible mechanism for inferring the trustworthiness of health social media is the use of social network and link analysis of online communities. This is a very promising approach because findings suggest that it enables online communities of patients to filter misleading content[45]. In addition, link and

social network analysis are successfully applied by most web information retrieval systems.

However, there are several areas where more research was needed in relation to RQ3. In health social media, online health communities are no longer isolated because they share the same general-purpose social networks (e.g. YouTube, Facebook, Flickr). In this context, we do not know how antagonistic communities try to influence each other. In Section 5.2, I present a study on the interactions between pro-anorexia and pro-recovery peers on a photo-sharing site. There was also a lack of knowledge about the features of the most reputable members within online diabetes communities; this is addressed in the study presented in Section 5.3.

5.1. Trust and Health Social Networks

Trust is a crucial part of online health information delivery since information that is untrusted by the receiver will scarcely be taken into consideration. There have been many studies looking at what makes online information trustworthy [59,107,121,138]. Most of them have looked into the validity of the information provider and also the information itself. Websites can be designed to enhance credibility, by careful design of both graphical and textual elements[59]. However, trust is always subjective and rather personal, thus some other studies have explored personalizing recommendations based on trust [121]. In the health domain, trust becomes even more complex since the information seeker may trust information but still be misleading and harmful; and vice versa. Studies show that psychological, social (e.g. educational level) and many other factors affect our trust in online health information[143,154,163]. Information about a specific treatment in the USA may not be valid in Europe, where the same treatment might even be illegal. I believe that very few domains have more complex trust models than the health domain.

As explained in the studies of this chapter, trust in online health information is a complex and multidimensional aspect. In the context of social media, trust is also related to the confidence of the messenger within the online community. Many studies have looked at the influences in health social networks, both online[28] and real-life social networks[26]. For example, in a study of a large smoking cessation online community the position of the members within the network was found to correlate with their health behavior[28]. These findings are not directly related to trustworthiness but we can conjecture that health

behavior was more influenced by the top-influencers of the networks, those with more social trust. There are many psychological theories, such as the Social Cognitive Theory[159], explaining why people tend to put more trust in leaders and also similar peers. These social influences can be studied using social network analysis.

In the context of health social networks, several aspects have not received sufficient attention from the research community. For example, most studies have focused on social networks that address one common health interest, such as communities promoting anorexia as a lifestyle. However, there are no studies looking into the network and social influence dynamics between antagonistic communities, an aspect I address in RQ3.Study I.

Trust is a very complex concept that has been used across disciplines. According to the Encyclopedia Britannica's dictionary, trust is the "belief that someone or something is reliable, good, honest, effective, etc"[37]. Overall, trust models are present in many computer science domains[138], but in this dissertation I focus on information retrieval and social trust. Most current algorithms for web information retrieval combine both relevance and trustworthiness, since untrusted but relevant content is clearly useless (e.g. spam). Trustworthiness in web information retrieval is in most cases calculated using link analysis metrics which resemble social network metrics predating the Internet. A link to a website is inferred as endorsement towards it. As explained by the authors of the TrustRank[72] algorithm, untrusted webpages are those that maliciously try to influence search engines. These websites are defined as web spam[73]. In recommender systems, trust often refers to users' belief in the quality/usefulness of the recommendations[121].

The dichotomy between usefulness and trustworthiness can also be applied to health information retrieval[42,48,143,154,163]. Health information is only useful if trusted by the receiver. From a public health point of view, the issue of trust is often described as a communication problem since information needs to be valid but also believed by the public [43,87]. In fact, health information is expected to be used to improve the health in our societies, by acquiring healthy habits, for example. As demonstrated by the following case related to anorexia, the trustworthiness of health information is in many cases subjective. For example, a patient with an eating disorder might believe that extreme dieting is part of a lifestyle choice and therefore the patient will mistrust health authorities (and websites) saying the opposite.

5.2.RQ3.Study I: Anorexia Communities in Flickr

Anorexia nervosa is highly prevalent in developed countries and on the rise in the rest of the world [79]. This eating disorder has a huge impact in the quality of life and health of the affected people and their relatives[6].

The case of anorexia nervosa and social media is paradigmatic due to the presence of an online movement of people affected by anorexia nervosa who advocate that anorexia is not a disease but rather a lifestyle choice [63,76,103,117,120,135]. These communities have been studied for many years and are commonly defined as "pro-ana" or lately the term "thinspo" (inspiration to be thin) has become popular. Members of pro-ana communities share tips, emotional support, pictures, etc. aimed at helping each other achieve unrealistic and dangerous behaviors to lose weight. This pro-anorexia content is a cause of major public concern since it can affect the health of the younger and more vulnerable population[8]. A study found that more than 10% of female high-school students in Belgium have consumed online pro-anorexia content [31].

In relation to this dissertation, the case of anorexia health social media is of special interest due to the co-existence of a community promoting harmful content with others providing informative content. The use of social network analysis to identify trustworthy content providers faces a serious challenge in this situation if one takes into consideration that pro-anorexia members can have high reputations within the community. In this study, I personally wanted to explore the differences in network dynamics across the different types of user and also establish which metrics are best at predicting whether a user is pro-anorexia or pro-recovery.

5.2.1. Methods

The data was extracted from the Yahoo photo-sharing site Flickr. As Figure 19 shows, there are four kinds of links between content and users: contacts, favorites, comments and tags. Users can comment on profiles and pictures of other users, and favorite and tag pictures of other users. In addition, users can join groups.

Figure 19: Flickr example of a pro-anorexia picture (RQ3.Paper 1)

In our study we used data that was public, in order to avoid privacy problems. The data was collected during February 2012 using the Flickr API and crawling of actual pages. The selection of seed users was performed using four methods; in total we identified 753 users. The first method consisted of searching for users who uploaded at least two photos using pro-anorexia keywords (e.g., thinspo, thinspiration, proana). The second method was to select users who posted at least two photos in anorexia-related groups (e.g. Anorexia Help). We also identified commenters on at least two anorexia-related photos. Finally, we selected users based on favoring anorexia-related photos. Of those seed users (n=753), we extracted all the publicly available information about their activity on Flickr.

We obtained metadata information of 543,891 photos; 2,229,489 comments; 642,317 favorite links; 237,165 outgoing links of contacts between users.

The collected users were labeled according to their degree of pro-anorexia or pro-recovery by researchers using a Likert-based scale. A good agreement was achieved according to Kappa inter-rater agreement (0.51, P<.001)[29]. There were 172 pro-recovery users and 319 pro-anorexia users. Further tags related to anorexia content were identified selecting the tags that were at least 10 times more likely to appear on non-anorexia related photos. A total of 25,689 photos contained at least one of these tags.

5.2.2. Results

Posting activity: the two sub-communities presented a correlation between the posting activities (0.82, P<.001). In general, pro-recovery users are more active, posting a median of 196 photos, compared to 105 photographs by pro-anorexia users (statistically significant, ranksum, P< 10-5).

Tagging: The pro-anorexia community used fewer tags related to self-portraits. There was a set of tags more commonly found in pro-anorexia videos such as "thispo", "doll" and "skinny". The pro-recovery users had a wider variety in tags including unrelated tags such as "garden". Overall, the most commons tags of pro-anorexia users refer to body image and interestingly also to "cigarette".

The tags used by each sub-community were modeled using a vector-space model weighted by Inverse Document Frequency (TF-IDF)[167] with the distance between photos being measured using cosine similarity.

The similarity of tags between users within the pro-anorexia group was 0.259, while in the pro-recovery group it was 0.202. That means pro-anorexia users share more vocabulary than pro-recovery. In fact, the average similarity between the two groups was 0.225 (differences significant at P< 10-5, ranksum test). That means that the similarity among tags was more common between the pro-anorexia groups than within the pro-recovery group.

In addition to the wider range of topics addressed by pro-recovery users, another reason was that pro-recovery users often used tags associated with the pro-anorexia camp such as "pro-anorexia". A plausible explanation for this is that pro-recovery users are trying to

get the attention of the other sub-community. Overall, the Spearman correlation between tag frequencies in both communities was 0.67 (P< 10-5).

Inter and intra community connectivity: contacts were more likely in the same group: 72% of pro-recovery contacts by pro-recovery users were to users in the same group. In the case of pro-anorexia users the percentage was 59%. Similar results were found regarding the commenting activity. Pro-recovery users were almost as likely to favorite a photo regardless of the group, but pro-anorexia users were 8.4 times more likely to favorite a photo from a fellow pro-anorexia user than from a pro-recovery user (89% vs. 11%, statistically significant at P< 10-5, chi2test).

Figure 20 shows the networks according to the different types of connections. The graph in the top left, which is based on contact links between users, shows a clear grouping between the different pro-anorexia and pro-recovery groups. The bottom-left graph, which is based on tags, shows a much more intermingling network due to the use of shared vocabulary across the groups.
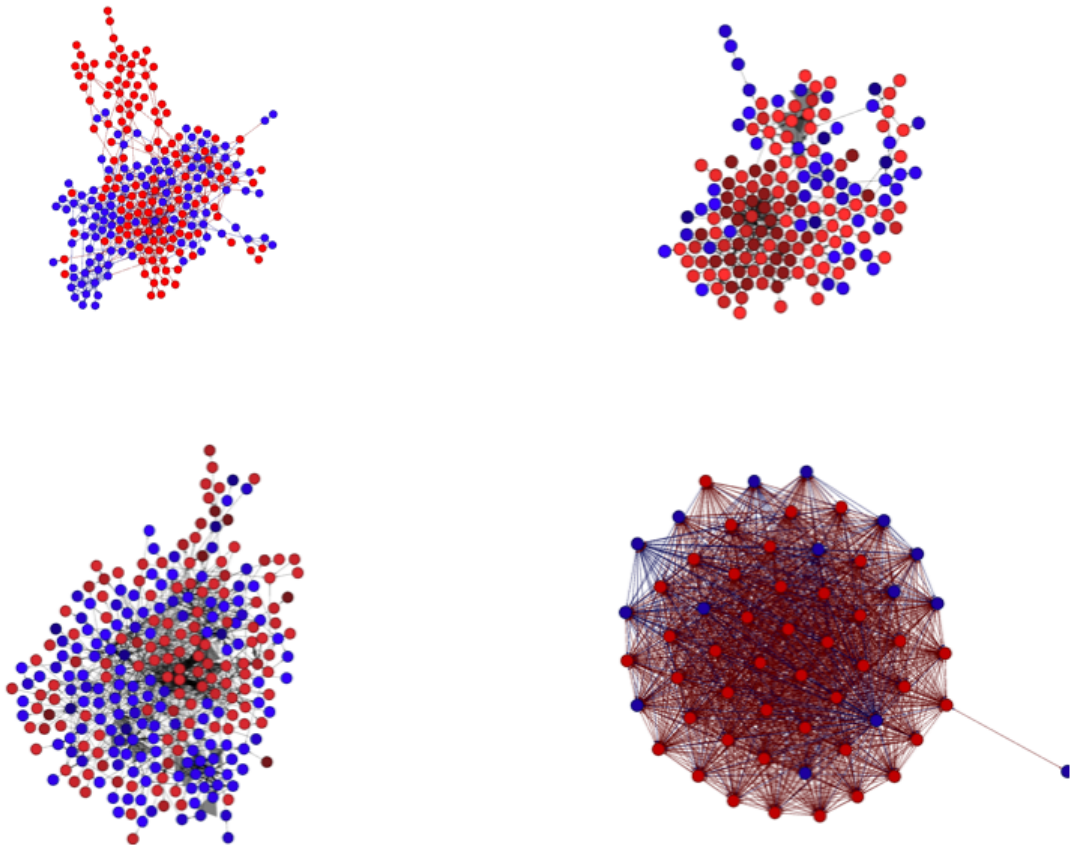
Figure 20: Flickr Pro-anorexia and Recovery Network graphs according to four connection types (from top left, clockwise): Contacts, Favorites, Tags, Comments. Blue represents pro-recovery and red pro-anorexia (RQ3.Paper 1)

Automatic classification of the groups: using the ROC curve[40] metric we explored the predictive value of the different connectivity data. The ROC curve is used to visualize the performance of a classifier, its area can range from 0 to 1 (perfect classification) and 0.5 equals random classification. The ROC using the comments or contacts was 0.74, using favorites was 0.53 and 0.52 using the tags network. Two users were considered similar in tags if their cosine distance between their tags was greater than 99%. Thus, the social links between the users were the best predictive values.

Inter-community posting as an intervention: some effects on user behavior were found due to inter-community posting. That shows that the sub-communities were intermingling. For a more detail description of this phenomenon see the study paper .

5.2.3. **Discussion**

There are two sub-communities related to anorexia nervosa in Flickr, one supporting people trying to recover (pro-recovery) and the other supporting eating disorder behavior (pro-anorexia). This study represents novel work looking into those two sub-communities since previous research has focused solely on pro-anorexia communities.

An important aspect of this study for the retrieval of trustworthy information is the finding that sub-communities heavily altered popularity measurements. For example, pro-anorexia users were 8.4 times more likely to favorite pro-anorexia photos and this was mirrored in the pro-recovery group. This user behavior contributed to the creation of two clearly distinct sub-communities taking into consideration social network dynamics.

This study showed that social links (i.e., favorite, friends) were the best metrics to automatically classify users within the two different groups. That finding is of paramount importance in this dissertation since it reinforces the idea that social network metrics can be used to identify misleading content. Misleading content in this case are all the photos created within the pro-anorexia sub-community, as our research shows that the pro-recovery community does not intermingle much with the pro-anorexia community. In contrast, text based classification was not very useful for classification since many words were often used by both sub-communities.

This study has also been very important for defining future areas of research. A better understanding about harmful information (e.g., pro-anorexia) is of great importance for public health authorities.

Limitations

This study had several limitations. First of all, it considered a very specific type of social network (photo sharing). We can expect similar results to be found in other content-sharing social networks (e.g. video sharing); nevertheless, these findings cannot be generalized without further testing in other types of networks.

Secondly, our classification method was based on a manual review of user profiles and we are unable to guarantee that the classification was correct without asking the users. Thirdly, our data collection only took into consideration publicly available information and

therefore it is likely that information of members with more restrictive privacy settings was missing.

5.3. RQ3. Study 2: Diabetes Online Communities

Diabetes is a good case-study for health social media since people who are affected by that set of diseases have a great deal to learn about the management of their condition and they also need a lot of emotional support. That is one of the reasons why diabetes online communities have been very successful. For example, the online community TuDiabetes already has more than 50,000 members. That diabetes community is led by patients and has already developed videogames, viral video campaigns, applications for personal health records, etc. The use of social network analysis to study online communities has attracted a lot of attention in the health domain as reviewed by Dunn and Westbrook[41].

The goal of the study presented here was to gather more information about diabetes social networks, including their structure and also the parameters related to the user profiles. As explained below, observing several social networks, the study looked into the application of several algorithms for community detection and the characteristics of the most relevant members.

5.3.1. Methods

Data acquisition and network modeling: in this study, public forum activity from five diabetes forums was extracted using a web crawler. These forums contained 140,000 registered users and 1.6M posts. That user interaction data was used to model social networks where the user was the node and the edges between users were based on comments about the original topic. In addition, publicly available information from the users' profiles was extracted which included health and demographic information. The names of the diabetes communities were not disclosed in the paper due to certain legal restrictions.

Studying community structure: the network structure was analyzed using PageRank and Authority-Hub algorithms[16,90] to measure the users' characteristics. The authority score related to the authoritativeness of the node within the network, this in turn was co-related to the number of incoming links for a given node. On the other hand, the hubs score was

high in those nodes that were acting as connectors between multiple nodes. To perform this type of analysis, we had to build a directed social network. Other parameters from the social networks, such as density and centrality[164], were also analyzed. The community detection was based on algorithms such as Affinity Propagation and Greedy Optimization algorithms[27,149].

5.3.2. Results

The results of the study showed that it was possible to model diabetes social networks and sub-communities using the forum data. The network structure, as Figure 21 shows, presented a distinctly dominant star topology[41]. Low user participation was also found in the forums, suggesting a core of users who were heavily active compared to the rest.
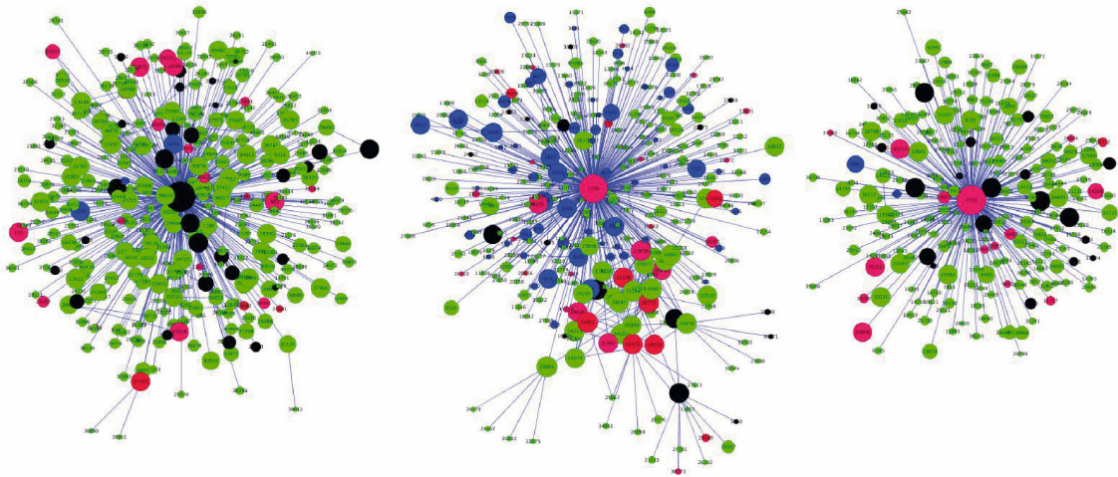
Figure 21: Communities' structures in a diabetes online community (F1) based on Greedy Optimization Algorithm, black nodes represent users who were diagnosed more than 10 years ago, red between 2-9 years, green less than one and blue without data. (RQ3.Paper2)

To further explore the characteristics of the active and centric users we looked at the demographic information available in the F1 forum. In F1, 82% of the users provided health information. The majority of the users (78%) who provided data had been diagnosed less than two years earlier (green nodes in Figure 21). The more central nodes had at least two years' experience with diabetes. It appears that patients with between two and ten years of experience living with diabetes are highly active supporting less experienced diabetes users. We also found different patterns of activity that may represent help seeking. For example, some "experienced" nodes had many incoming links

from less experienced nodes, which might indicate a newly diagnosed patient struggling with the disease seeking support from an experienced diabetes patient.

5.3.3. Discussion

This study provided very interesting results as far as this dissertation is concerned. It showed that the most reputable users (i.e. central nodes) were in fact those with more experience managing the disease (i.e. a longer post-diagnosis period). That reinforces the assumption that network metrics to find the more reputable and trustworthy users may be correlated with more knowledge about the topic based on greater experience with the disease. In the case of diabetes, a patient with greater experience handling the disease will be better prepared to discern misleading information than a newly diagnosed user.

The main limitation of this study concerned the nature of the data since it did not contain all the social networking activity (e.g. no private messages between users, incomplete demographic information) and information could have been forged (e.g. minors forging their age to join networks for adults). A second limitation was the fact that the networks were moderated. It is therefore possible thatusers promoting misleading information had been removed by administrators.

6. The HealthTrust Metric

> **Related papers: RQ4.Paper 1 and RQ4.Paper 2**
>
> **RQ4: Can trust-based metrics improve the retrieval of social videos about diabetes?**
>
> - **RQ4.1**: Can a metric of trustworthiness within a health community be used to retrieve relevant trustworthy providers of diabetes videos?
>
> - **RQ4.2**: Can a metric of trustworthiness within a health community be used to search for relevant trustworthy diabetes videos?

Despite the increased availability of social media published by health authorities and medical associations, finding online health information is no easy task[46,47,70,108].There is a great deal of misinformation, including content promoting anorexia or discouraging vaccinations[2,31,152,169]. That type of content can become extremely popular and viral (e.g., conspiracy theories about vaccination). What is more, quality labels and certificates for online health information are not always effective[14,81]. Therefore, sifting through this to find trustworthy health information remains one of the main challenges faced by health consumers.

Participation and socialization are intrinsic parts of social media. As Figure 22 illustrates, when a patient publishes a video about their diabetes this could potentially be part of a social conversation where people endorse the video (ratings, favorite) and also comment. In addition, users on social platforms such as Flickr or YouTube form groups and communities. In RQ1.Paper1, patients publishing videos about their health described themselves as members of an online community.
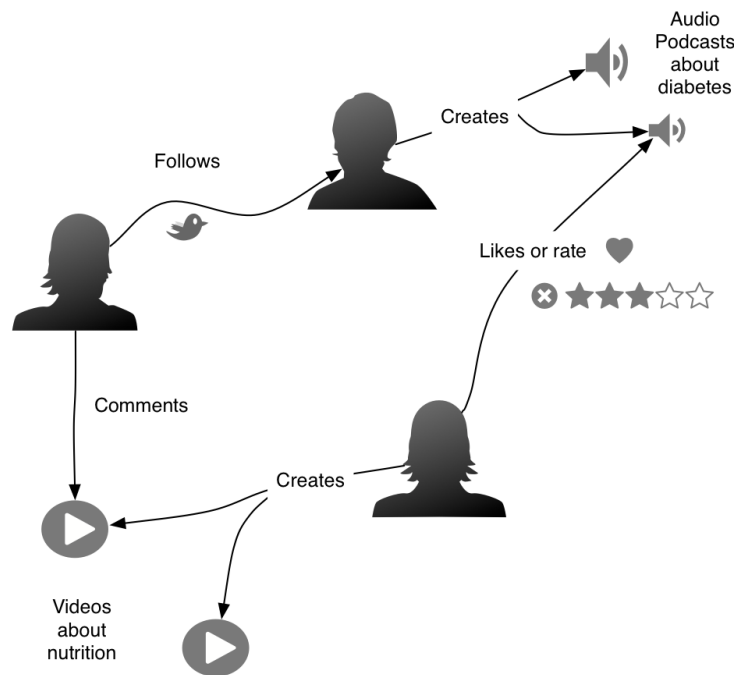
Figure 22: Example of Health Social Network (RQ4.Paper2)

These health social networks in many cases also become hubs of health information. As described in several of the previous studies, it is common for consumers of health social media to ask the creators for content addressing specific topics or enrich the content with comments. Esquivel A et al. reported that misleading information about cancer in an online forum was detected and deleted by community members within a few hours[45]. Thus, health communities can become powerful information retrieval tools. In addition, these networks are very robust to spammers since reputation in an online health community is not easy to create.

The objective of this study, which is the core part of the dissertation, was to explore the feasibility of extracting metrics about the trustworthiness of content and providers within online health communities. The assumption is that trustworthiness within a health community can be used to predict the quality of the health content. Traditionally, relying heavily on a particular community to model trust has been considered pernicious, as those communities might not represent the interests of the general public (see Tightly Knit Community description in Sub-section 0). An authoritative member of the community, such as the Joslin Diabetes Centre, tends to publish or endorse content of better quality than non-authoritative members of the community, such as somebody selling herbal supplements for diabetes.

The trustworthiness metric used in this dissertation was built based on the integration of different link-based analyses as described in the following sections. As a case-study, I decided to choose online videos about diabetes. Diabetes is one of the most prevalent diseases worldwide and this prevalence is increasing[166]. In addition, people with diabetes need to acquire a thorough understanding of their disease that can affect many aspects of their daily life[44]. Therefore, people affected by diabetes are a prime example of users who need health-related information retrieval tools and are actively engaged in social media[24,71]. The selection of online videos was based on the fact that online videos are one of the most popular and fastest-growing types of social media, and YouTube is the leader worldwide[18,105,147]. I also had some strong indicators of the presence of a diabetes community on YouTube based on the interviews with patients described in RQ.Paper1.

Relation to other Research Questions: the studies performed to address these research questions are central to my dissertation and they rely on the knowledge acquired in the previous studies. First of all, the design of HealthTrust is based on the knowledge acquired about the unique characteristics of health social media (RQ1) and videos in particular. In addition, the review about techniques for modeling health social media (RQ2) highlighted significant advantages of using social network analysis for modeling content and users of health social media. Finally, the studies performed in RQ3 provided insights into how social network dynamics could affect the trust of health social media content.

6.1. Web Search Engines

Nowadays, modern information retrieval is heavily knitted to the Internet, as described in the book Modern Information Retrieval [114]. Web Search Engines are of particular interest in the health domain since they are used by a significant number of health consumers. In addition, most web search engines rely on link analysis that is also a technique used in HealthTrust.

Searching for information was one of the first applications of computers. In fact, Mooers coined Information Retrieval as a computing term in 1951 [115]. In those early days of Information Retrieval, the main task was to search for information within big databases, especially for corporate use (e.g. searching medical records of a given patient).Information

Retrieval is a crucial part in the field of Medical Informatics as reviewed by W Hersh in his iconic book "Information Retrieval – a health and biomedical perspective"[77]. An Information Retrieval system, such as a search engine, consists of content, which is stored in a computer either locally or in a network such as Internet. The content is grouped in a collection, which can be defined as database, site, repository, records, and documents. In search engines, users can actively seek content by providing an input of queries that the system will use to retrieve results. Traditionally, user input is based on a text query.

The system matches the query with the metadata describing the information stored in the content collection. The metadata is a set of attributes about the content, normally including a set of terms used for indexing and other attributes such as the authoritativeness of the content. Indexing is a crucial part in the process and it also facilitates fast information retrieval. Retrieval is the process of interaction between the user seeking information and the system. The information requirement of a user is modeled in the search query that is used to match the content.

Finding high quality and relevant information is not a new problem in the online world. At the beginning of the World Wide Web there were some approaches towards the creation of a global repository classifying all online content. However, that approach was rapidly overcome by the development of web search engines. Figure 23 below, shows the typical structure of a web search engine as explained by A. Sonawane[145]. Several aspects of web searches are unique in the information retrieval context. Web space is immense and in nearly all cases the number of relevant webpages is far too large for a human to classify (e.g. there are millions of websites about diabetes). In addition, there is no central file or system with a list of all the available web resources. For that reason, web search engines have "crawlers" or "spiders" that systematically scrutinize the entire web in search of new content. The crawling of websites is generally based on the analysis of the links between websites. These crawlers extract all the relevant metadata for any web content. That metadata is the integrated in index files. A user's search query is processed and used to find relevant webpages from the index files. However, the process does not finish here, since, in many cases, a huge quantity of webpages will match the search query. That is why ranking is one of the most crucial aspects. Ranking is normally based on several metrics about the importance of the webpage such as PageRank, as explained below.
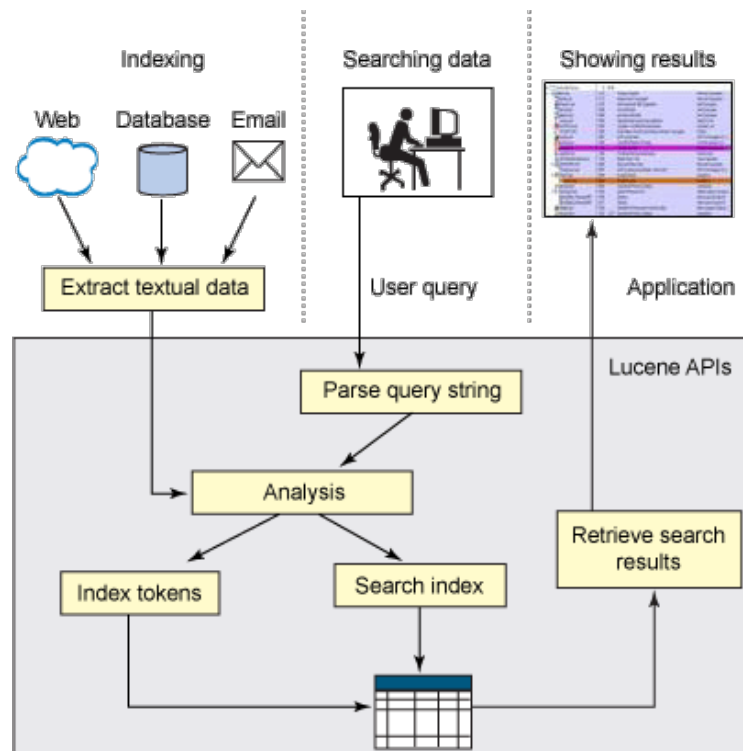
Figure 23: Architecture of Lucene (Open Source Web Search Engine). Image from A. Sonawane[145]

Link analysis for rating websites

The most relevant aspect of web searching for this dissertation is the process of determining the importance of a website. Web search engines have gradually had to adopt new techniques for finding relevant and high quality content while avoiding low quality content such as spam[73]. The most common approach has been the use of link analysis[16,69,72,90,99]. The links between websites are used to model a graph where nodes represent webpages and the edges are the links between them. That web graph can be interpreted as a social network where the nodes (aka webpages) with more incoming links can be presumed to be more reputable, because linking normally denotes endorsement. Metrics about the importance of a webpage are calculated based on link-analysis. These metrics are then used to put the search results into order for a given search query. The most famous metric is PageRank, which was the original metric used in the algorithm of Google[16]. In Google's PageRank, links from one site to another can be modeled as an endorsement, and they are used to calculate a global rank of all the websites. Another example is the hyperlink-induced topic search (HITS) algorithm[90].

Other algorithms, such as TrustRank, take into account trustworthiness in online communities, aimed at making the search more robust to Web spam[72]. TrustRank is similar to PageRank but takes into special consideration webpages of well-known trustworthiness. The social network of the users has also been taken into account to personalize the search results. For example, Gou et al explored how to use social network analysis for the personalized ranking of online videos[69]. Mislove et al., studied the integration of social networks with online searches[112].

Overall, these link-based metrics are designed to find the more reputable webpages and also to fight "spammers". In addition, they are designed to counter the effect of webpages which, deliberately or not, have artificially high scores in those metrics.

Hyperlink-induced Topic Search (HITS): Hubs and Authorities scores

For illustrative reasons, I will explain the basic aspects of the HITS algorithm that was used in the study described here[90]. The objective of the HITS algorithm is to identify the most authoritative webpages to rank the search results. HITS is based on two scores: authoritativeness and hubs. The hubs represent webpages that link to many authoritative pages and consequently act as directories. The authorities are those webpages that are highly linked by many of the most representative webpages; they can be seen as authorities within the web space. The idea behind that approach is that linking to a particular webpage can be seen as conferring authority.
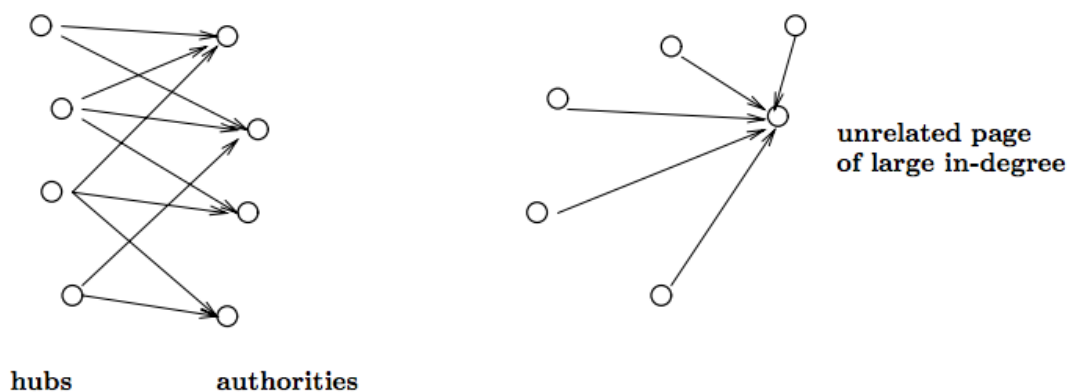
hubs   authorities

Figure 24. Diagram showing the structure of linked pages representing hubs and authorities.
Images from Kleinberg 1999[90]

The calculation of Hubs and Authorities is based entirely on the analysis of the link structure of a set of webpages. Each website has two scores for its authority (x = authority

weight) and hubs value (*y*=hub weight). These values are normalized across the whole dataset so their squares sum to 1. There is a relationship between hubs and authorities in the sense that if a *p* page points to many pages with a large *x-value*, then it should receive a large *y-value*. Similarly, if many pages with a large y-value point to a *p* page, then *p* should receive a larger x-value. In other words, the hub score is the sum of the authority scores of all its linking pages. Based on that relationship, there are two operations to update the *x-values* and *y-values*. Very briefly, the algorithm starts with a collection of linked pages where all their x-values and y-values start with 1, then several iterations are made in which the x-values and y-values are updated and normalized.

Figure 25 shows the iterative operation that is calculated for each page to update the weight values for both authoritative and hub scores. If a p page points to many pages with a large x value (aka, good authorities) then p should have a larger y weight, which means it is a good hub. The left of **Figure 25** shows operation *I* to calculate authority weights, while the figure on the right shows operation *O* to calculate hub values. As the operations show, both authoritativeness and hubs reinforce each other as follows:

- x[p] = sum of y [q], where q represents all the pages pointing at p.
- y [p] = sum of x [q], where q represents all the pages pointed to by p.

$$x^{\langle p \rangle} \leftarrow \sum_{q:(q,p) \in E} y^{\langle q \rangle}. \quad \bigg| \quad y^{\langle p \rangle} \leftarrow \sum_{q:(p,q) \in E} x^{\langle q \rangle}.$$

Figure 25: Operations to update authority (x) and hub weights (y). Images from Kleinberg 1999[90]

As explained in detail by Kleinberg[90], that operation is performed several times for all the pages in the collection in an iterative manner, including normalization. This is to calculate all the authoritative and hub weights which are later used to retrieve the pages.

PageRank relation to HITS

The PageRank[16] algorithm determines the probability of a person randomly clicking on links to arrive at a particular page. Thus the probabilistic distribution of PageRank will be affected by the links between the pages. The more incoming links to a website the higher the chances of being clicked. Recursive operations to determine the PageRank of a given page can be implemented in an iterative algorithm.

Comparing HITS with the well-known PageRank algorithm used by Google, the main difference in HITS is that the scores are calculated at query time over a root set of webpages (typically 1000-5000 pages) which are related to the search query. The scores are calculated over that set. The PageRank algorithm is designed to be applied to the entire WWW without any query dependency.

6.1.1. Health Web Search

One of the main challenges in the health domain is that misleading health information can be very popular (e.g., anti-vaccination videos) and therefore may be paradoxically highly rated and not considered spam by general information retrieval algorithms.

Information retrieval has been widely used in the health domain as described by William Hersh in his book "Information Retrieval: a health and biomedical perspective"[77]. Search engines have been widely used to retrieve medical records (e.g. searching for a certain condition in a patient´s record) and also biomedical information (e.g. searching genotype information). Web search has been studied mainly for the retrieval of research papers from online knowledge databases such as PubMed. For example, under the paradigm of evidence-based medicine, a doctor with a patient with a rare disease may be interested in finding out the latest guidelines about treating that particular disease. Web health search is emerging as a new area of research with workshops and special issues in leading journals and conferences [38,170].

The problem of finding trustworthy health websites is not new. For over a decade, organizations such as the Health on the Net Foundation have been working on the creation of quality labels and guidelines for providers of health information[32,47,48,62,81,108]. The use of quality labels and seals to identify websites that adhere to health quality information guidelines is now common. That metadata information about adherence to quality guidelines has been automatized in some research projects by applying different natural language and semantic technologies[48,108]. That metadata can then be used to rank search results of websites with quality seals [47,108].

There are some major limitations to the previously described approaches. First of all, the evaluation of the quality of those adhering websites is revised manually, normally on a yearly basis. Therefore the quality certificates are not about a particular content (e.g. video, blog post) but rather a general aspect for the entire web portal where the content is

posted. Some studies have pointed out cases where those guidelines were not that effective for finding good health information[14,62,81]. Overall, these quality seals face major problems due to the lack of awareness among information providers and health consumers which is partially the result of the emergence of dozens of different quality seals.

To our knowledge, none of these projects have focused on link-based analysis and trust metrics of health websites as generic search engines do. In addition, despite the popularity of health videos, we have not come across any project specifically aimed at developing tools to help find relevant health videos.

6.2. RQ4. Study 1: Design and evaluation of the HealthTrust Metric

In this section, I will describe the HealthTrust metric followed by the explanation about how it was used to search for diabetes videos posted in YouTube. The methodology for the evaluation consisted of two studies for comparing the correlation between the HealthTrust scores and the perceived quality by end-users of diabetes videos and channels (aka users). We designed these experiments to evaluate our hypothesis that HealthTrust's metric can be used to improve the retrieval of health social media. The first study focused on diabetes channels HealthTrust scores while the second study was on videos about two specific topics: diabetes A1c testing and diabetic foot.

The Merriam Webster Dictionary defines trust as an "assured reliance on the character, ability, strength, or truth of someone or something"[37]. *Authoritativeness* ("clearly accurate or knowledgeable"[36]) and *reputation* ("overall quality or character as seen or judged by people in general"[132]) are also often used as synonyms of trust. In the Web Search domain, the reputation of a website is generally inferred by the analysis of the web structure. Basically, a link from one website to another infers an endorsement of the linked website. This approach is very similar to the impact factor used to rank scientific journals based on the citations of the journals' papers. In the health domain, trust is related more to the authoritativeness of the content provider, the reliability of the creator of the information, such as whether the messenger is qualified (i.e. medical doctor). There are, however, many additional aspects related to trust such as appearance and impartiality[75] since some actors may be trying to manipulate information. Within the scope of this study, we define trust as

the "assured reliance on the quality of users and content within an online health community."

An important motivation of the HealthTrust approach is that in some cases online health communities have been found to be very effective filtering out misleading health information[45]. Gaining trust within an online community requires a great deal of interaction with fellow users, thus users are very cautious about sharing misinformation that could damage their reputation. This is something reported by the patients interviewed in RQ1.Paper1. For example, a user sharing videos about herbal cures for diabetes would receive less endorsement from the diabetes community than a video created by the American Diabetes Association.

The HealthTrust metric is based on the assumption that misleading information will be endorsed less within a health community[45]. Therefore, trustworthiness within the health community will correlate with higher quality. The algorithm for calculating the HealthTrust metric is designed to estimate the trustworthiness of social media content within the health community to which it belongs.

HealthTrust and the TKC effect (Tightly Knit Community)

The HealthTrust approach is related to the social network dynamics of tightly knit communities. SALSA (Stochastic Approach for Link-structure Analysis) is another web ranking algorithm designed by Lempel and Moran[99]. SALSA is based on HITS but includes some improvements to make it more computationally efficient. More importantly, SALSA also has a different behavior with the TKC effect (Tightly Knit Community) described by Lempel and Moran: "*A tightly knit community is a small but highly interconnected set of sites*"[99]. The TKC effect occurs when pages from such tightly knit communities score high on link-based analysis due to the highly connected nature of the community. As a result of that, "problem" pages of such communities will rank in unjustifiedly high positions. As an example, Lempel and Moran ran some experiments searching webpages about abortion. They found that pro-life websites form a TKC and their websites therefore tended to be highly ranked. Their SALSA algorithm gave a more balanced list of authorities between pro-choice and pro-life webpages, avoiding the over representation of pro-life websites.

The consequence of the TKC effect in this dissertation is important. I conjecture that the TKC effect is positive since those health-related tightly connected sub-communities have

the potential to filter out low-quality content. Therefore, the approach in some web search engines of undermining the TKC effect could lead to lower quality when searching for health content. The HealthTrust algorithm described in this study is designed to increase the influence of health-related tightly knit communities to improve the quality of search results.

Health Trust Implementation and Methodology

A set of steps must be followed to calculate HealthTrust : (1) identification of the health community and extraction of its social network, (2) calculation of the authoritativeness scores of content and members, and (3) calculation of HealthTrust metric for content and members. HealthTrust can then be used to rank search results as explained in the subsection "HealthTrust for Search."

As shown earlier in Figure 22, users and content are heavily intermingled in social media via favorite links, ratings, friendship, and groups. These links form a massive social network that can be modeled as a directed graph. Using social network analysis, one can derive metrics about the trustworthiness (or social relevance) of any give node. Often those health-related social networks are formed due to a common interest (e.g. diabetes).

$$HealthTrust\ (content, community)$$
$$= Authoritativenes\ (content, community)\ x\ (1 - InheritanceFactor)$$
$$+ Authoritativeness\ (author\ of\ content, community)\ x\ InheritanceFactor$$

Figure 26: Calculation of HealthTrust (RQ4.Paper2)

In order to calculate the HealthTrust (Figure 26) score of given content we need to take into account the following parameters:

- ***content***: the content can be any type of social media, such as videos or photos.
- ***author of content***: the author of the content is the user who uploaded the content to his profile.
- ***community***: the community is a group of interlinked users with a common interest, such as diabetes. The community is modeled as a graph where nodes are the users and edges the relationship between them (e.g., friendship, favoring content).

- **InheritanceFactor**: The *InheritanceFactor* models this process of inheriting reputation from an author to his content. A higher *InheritanceFactor* will result in a more "transference" of the reputation from the content creator towards the content itself. The *InheritanceFactor* is a value from 0 to 1, a value of 0 will mean that the reputation of the author of content does not affect the HealthTrust value of his content. In some cases, it is common that by default we trust content provided by a very reputable content creator. For example, when the very reputable American Diabetes Association publishes a video, it gains strong reputation which in inherited from its creator.

- **Authoritativeness**: the authoritativeness is a function that calculates reputation of a given content (or author) based on link analysis algorithm. In the case of the calculation of the authoritativeness of content, the content is modeled as a page that is part of a graph of interlinked pages. These interlinked pages are in fact the community. As explained later, authoritativeness can be implemented using algorithms such as HITS or PageRank.

1.-Community Extraction

The first step to calculate HealthTrust is to extract a health community from which its social network graph can then be extracted. As explained in the previous chapter, there are many ways of extracting online communities[19,27]. Extracting the community is a core aspect in HealthTrust since it relates to the specific trust for a content or user within the health community. For example, an MTV video about a rock star with diabetes may be seen as more authoritative than videos from health agencies about diabetes, because MTV is trusted by a broader viewership on YouTube. In contrast, in HealthTrust the focus is exclusively on intra-community authoritativeness. In the case of diabetes, health authorities will be more reputable than MTV.

In our case study, we tested HealthTrust within YouTube´s diabetes community. As Figure 27 shows, YouTube is a social network where users share videos. Users who subscribe to others, favorite or comment videos shape the social network graph of YouTube[19,128]. To keep things simple, we only considered subscriptions and favorite links.

The community extraction varied between the two studies reported in this chapter. In the first, we searched and extracted all the channels (aka users) that had the word diabetes using YouTube´s API (Application Programming Interface). In the second, we extracted

channels based on a search of videos retrieved by searching for a set of diabetes-related queries. In both studies, we extracted all the accessible information about the channels (e.g., uploads, subscriptions, and favorites).
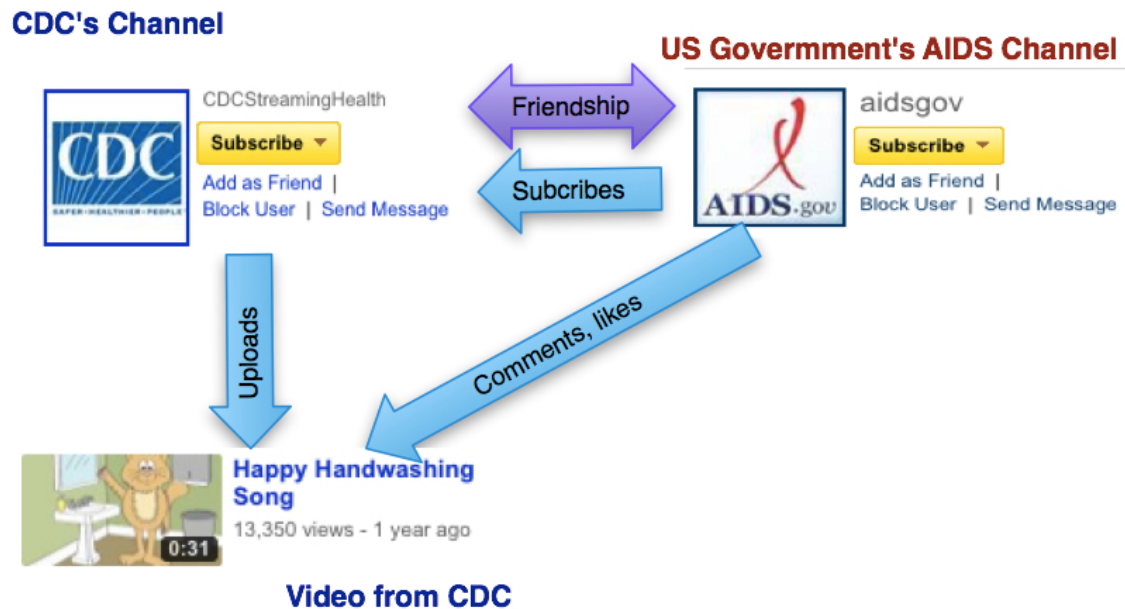
Figure 27: YouTube´s social network. CDC = Center for Disease Control and Prevention (RQ4.Paper2)

2.-Authoritativeness Scores

Having extracted the social graph, the next step is to calculate the authoritativeness scores. In HealthTrust these scores are meant to be calculated using well-known algorithms such as PageRank scores[16] or HITS authoritativeness[90]. HITS authoritativeness was used in the studies described in this chapter.

In order to apply HITS a directional graph needs to be created where nodes can either be videos or users, and their links (e.g. favorite, subscription) are the edges. As explained in the figure below, the authoritativeness of content and users are calculated as follows. First, the authoritativeness of content is calculated based on the links between users and content, considering both content and users as nodes. Second, the authoritativeness of users is calculated based on the links between all users excluding the videos as nodes.
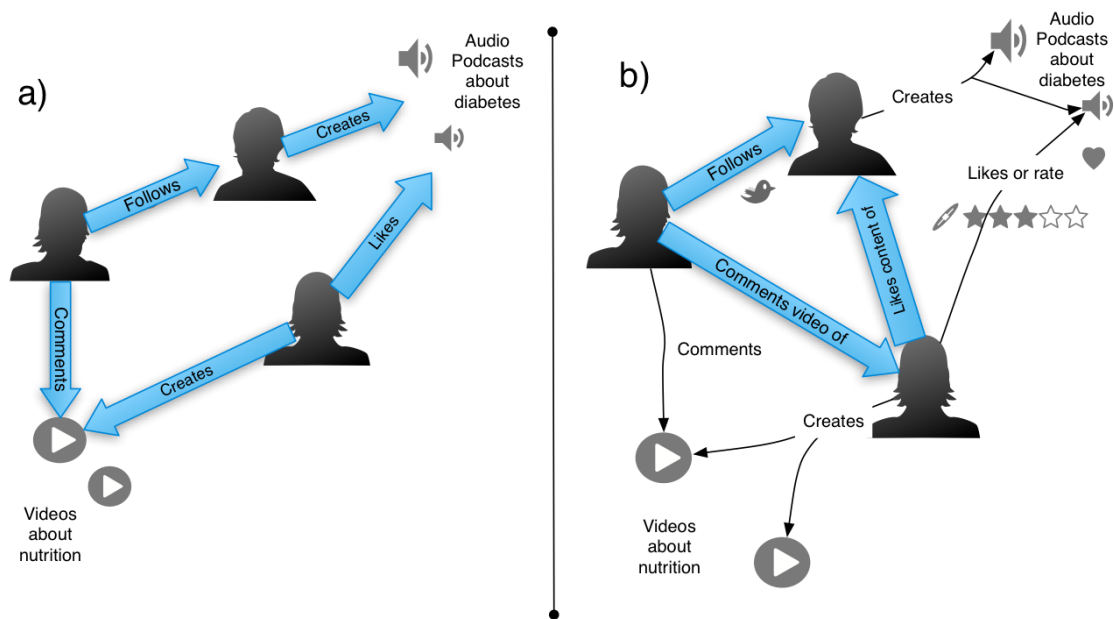
Figure 28. Calculation of HealthTrust: Links (in blue) used to Calculate authoritativeness of users (left) and content (right) (RQ4.Paper 2)

In order to calculate the authoritativeness values used in HealthTrust, I implemented the following algorithm:

1. Extraction of data from YouTube about videos and channels related to diabetes
2. Calculation of authoritativeness of channels:
   a. Create a directed graph where nodes are the channels.
   b. Create directed edges between nodes representing subscriptions) (channel X subscribed to channel Y) and favorites (channel X subscribed to video of channel Y).
   c. Calculate HITS authoritativeness values of the channels using JUNG (Java Universal Network/Graph, http://jung.sourceforge.net/)
3. Calculation of authoritativeness scores of the videos:
   a. Create a directed graph where nodes are both channels and videos.
   b. Edges are favorites (channel X subscribed to video Z) and subscriptions (channel X subscribed to video of channel Y).
   c. Calculate the HITS authoritativeness values of the videos using JUNG.
4. Normalize between 0 and 1 the authoritativeness scores of videos and users that were calculated independently.

3.-Calculation of HealthTrust

The next step is to calculate the HealthTrust score that is used to rank search results. The goal of HealthTrust is to give a metric about the trustworthiness of content from health social networks. As shown in Figure 29, the HealthTrust score of content is the weighted combination of the normalized authoritativeness scores.

$$
\begin{aligned}
HealthTrust\ &(content, community) \\
&= Authoritativenes\ (content, community)\ x\ (1 - InheritanceFactor) \\
&+ Authoritativeness\ (author\ of\ content, community)\ x\ InheritanceFactor
\end{aligned}
$$

Figure 29: Calculation of the HealthTrust metric (RQ4.Paper 2)

The combination of the authoritativeness scores of content and users depends on the *Inheritance Factor*. That variable will regulate how much "trust" is inherited by the content from the user. Thus, new content from a trusted author will have a higher HealthTrust score than new content from an untrustworthy author. A high Inheritance Factor will imply that the authoritativeness of the author is very important, which is the most common case in the health domain. An Inheritance Factor of 0 implies that there is no inheritance transfer of trust from the authors to their content; so all the authoritativeness

is based on the video's score. As you can see in Figure 30, in our case-study the diabetes videos we used had an *Inheritance Factor* of 0.7. Therefore, the HITS authoritative value of videos weighed 30% and the author's authoritativeness 70%. That high *Inheritance Factor* can be explained by the fact that we observed many videos from trusted sources that had low authoritativeness scores due to the novelty of the content and the subsequent lack of links towards them.

$$HealthTrust\ (video\ v, diabetes\ comunity)$$
$$= Authoritativenes\ (video, community)\ x\ (1-0.7)$$
$$+ Authoritativeness\ (author\ of\ video, community)\ x\ 0.7$$

Figure 30. HealthTrust calculation for diabetes videos from YouTube (RQ4. Paper 2)

4.-HealthTrust for Search

The search in HealthTrust is based on the combination of separate scores from trust and relevance. These two parameters can be combined in different ways providing finer granularity to adjust to the particular requirements of the study case. The combination (see Figure 31) of syntactical relevance and HealthTrust scores provides a new trust-based relevance score, which is used to rank search results.
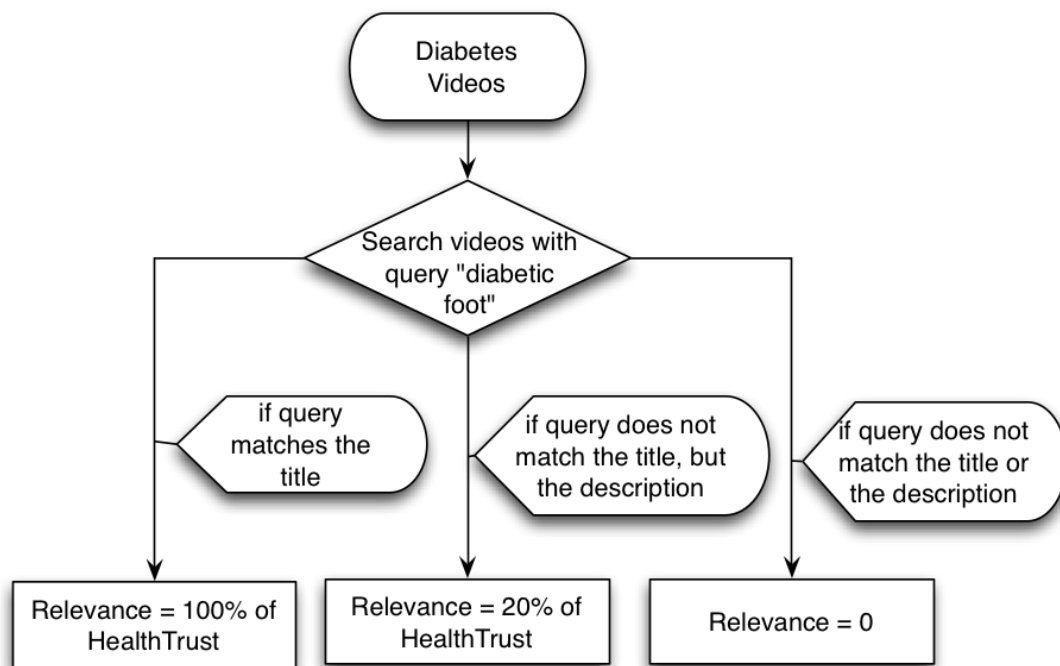
Figure 31: Relevance calculation for HealthTrust-based search (RQ4.Paper 2)

The search algorithm is based on the combination of two scores: (1) the relevance of the content to the search query, and (2) the HealthTrust score of the video. The relevance score was implemented with a simple query matching (eg, the query is contained in the title or description of the video). The algorithm (see Figure 31) to calculate the trust-based relevance was implemented as follows:

1) *videos* = all videos with query *X* in title or description.

2) For all the *videos* ($v_1,...v_k$):

      2.1) If (v_k has the query in the title): Trust-based relevance (v_k) = HealthTrust Score (v_k).

      2.2) Else If (v_k has the query in the description): Trust-based relevance (v_k) = HealthTrust Score (v_k) *0.2.

      2.3) Else: Trust-based relevance (v_k) = 0.

The reason for choosing such weighting was our observation of many lengthy descriptions of videos in many cases as part of an SEO (Search Engine Optimization) strategy. We also observed cases of videos with many tags, and in some cases irrelevant.

6.2.1. Evaluation of Diabetes Channels Search with HealthTrust

As described in RQ4.Paper1, we performed one first study on the use of HealthTrust to find trustworthy diabetes channels on YouTube. The main objective was to establish the feasibility of using authoritativeness scores to identify trustworthy channels (aka users). The **Data Collection** of that sub-study involved several steps:

1.-Data Collection (see Figure 32): using the YouTube API we extracted channels with the query diabetes. From those channels we extracted data such as links, videos and descriptions.

2.-As described in the previous subsection, we calculated the authoritativeness score of each of the YouTube channels.

3.- We created a list with the top-20 channels according to YouTube search (by YouTube´s relevance with the query diabetes) and the top-20 with higher trust-based relevance (as illustrated in Figure 1) using the algorithm of HealthTrust.
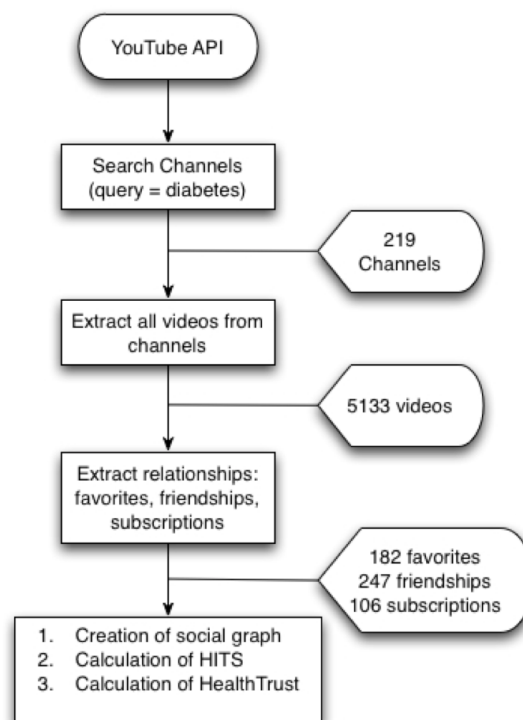
Figure 32. Data extraction in the study of diabetes channels and HealthTrust. API = application programming interface; HITS = hyperlink-induced topic search (RQ4.Paper 2)

The **evaluation** of the perceived quality of retrieved channels was performed by two healthcare professionals who classified the channels independently. They had to answer whether they would recommend the channel to a patient with diabetes. The quality of the classification was measured with Cohen Kappa[29] using the statistical framework R for psychology research [133] and it resulted in good agreement (.61). As explained in the results subsection, the results were analyzed using the metrics Precision at K and Discounted Cumulative Gain (DCG)[17].

Results for Diabetes Channels Search

Our approach to retrieve diabetes channels was evaluated using the well-known metric Precision at K (k=5, k=10 and k=20). We compared precision at K for the top results retrieved by HealthTrust and YouTube Search. The relevance of a channel was determined when at least one reviewer (out of 2) recommended the channel (with the second reviewer being neutral). As illustrated in the table bellow, Precision at K (K=20, K=10, K=5) was better for HealthTrust in 11 the 12 of the evaluated metrics. Evaluation based on DCG also resulted in better score for HealthTrust (see R4.Paper 1).

Table 10: Evaluation of the top20 channels in HealthTrust Study

Recommended by / Precision	Both reviewers		At least one reviewer	
	YouTube	HealthTrust	YouTube	HealthTrust
Precision at 5	80% (4)	80% (4)	80% (4)	100% (5)
Precision at 10	60% (6)	70% (7)	70% (7)	90% (9)
Precision at 20	50% (10)	65% (13)	60% (12)	85% (17)

Filtering out misleading content is particularly important in the health domain. We analyzed how well the algorithms performed taking into consideration channels that none of the reviewers recommended. For P@K20, HealthTrust's list had only 3 bad channels (15%) versus 8 (40%) on the YouTube list. Similar results were found for P@K10 and P@K5 as described in R4.Paper 1.

Some of the "bad" channels, according to the reviewers, retrieved by YouTube search featured famous singers with diabetes or were commercials about diabetes products, etc.

The HealthTrust list did not contain any channels with advertising, but it did have some channels from e-patients with very heterogeneous quality. Surprisingly, some well-known diabetes institutions (e.g. Diabetes Research Foundation) were not represented in the top-20 of HealthTrust, but top-ranked in YouTube. The over representation of patients in HealthTrust might be explained by their high level of social interactions within the diabetes community on YouTube.

6.2.2. Evaluation of Diabetes Videos Search with HealthTrust

In RQ4.Paper2, we describe the main sub-study to evaluate HealthTrust. In this case the objective was to evaluate how HealthTrust can be used to retrieve trustworthy videos (and not users as in RQ4.Paper 1). As explained below, the approach of this study involved several steps.

The Data Collection (see figure below) involved the use of the YouTube API to extract relevant videos using several diabetes-related queries such as diabetic foot, diabetes and diabetes ketoacidosis (complete list is shown in **Figure 33**). From those videos we extracted information about channels, subscriptions, etc.
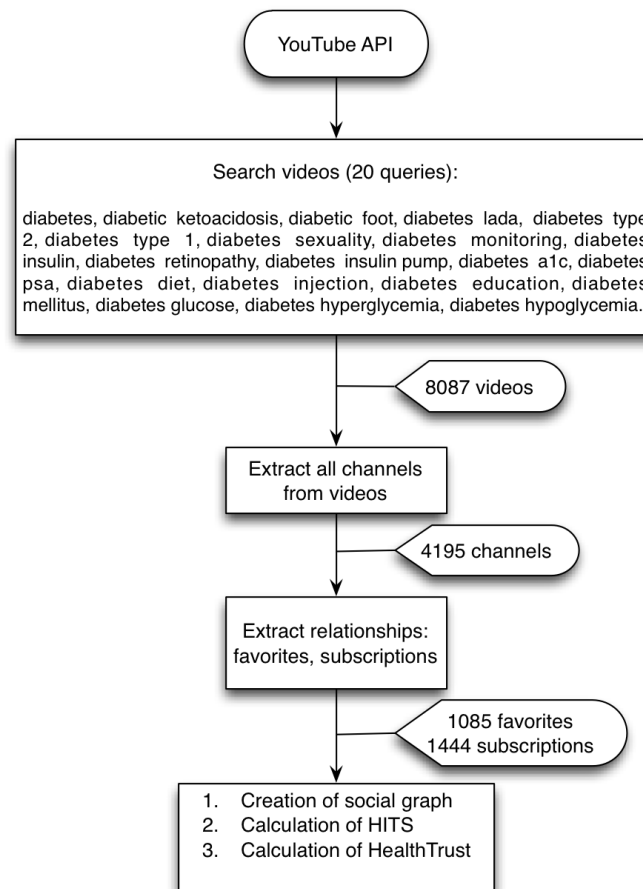
Figure 33: Data extraction on the search study on diabetes videos and HealthTrust (from RQ4.Paper 2)

The evaluation was limited to two search queries representing common information needs in patients with diabetes (diabetic foot and hemoglobin A1c -glycated hemoglobin testing). Lists with the top-7 ranked videos for each query were created using YouTube search and HealthTrust search. Patients and professionals evaluated these lists. In the case of patients, only videos from trusted channels were used to avoid showing potentially harmful videos to them.

Professionals were recruited using a Snow Ball [68] sampling mainly via mailing lists. In total, professionals provided 162 ratings of 23 videos. Patients were recruited via a mailing list managed by the diabetes online community TuDiabetes.org. In total, consumers provided 427 ratings of 17 videos.

Video Surveys

Using a web survey, we evaluated the top 7 video search results for the queries "diabetic foot" and "diabetes A1c" using 1) HealthTrust and 2) YouTube search ordered by relevance. As explained in Figure 34: Rating process for list retrieved by HealthTrust (RQ4.Paper 2) was as follows: after the acceptance of the informed consent, the patients and professionals received an email with a random link to different web surveys with the video lists. Patients were assigned randomly to lists for the same queries. On the web form, they had to respond to questions about the videos with Likert scale questions such as "Would you recommend this video to a patient with diabetes and questions about diabetic foot?".
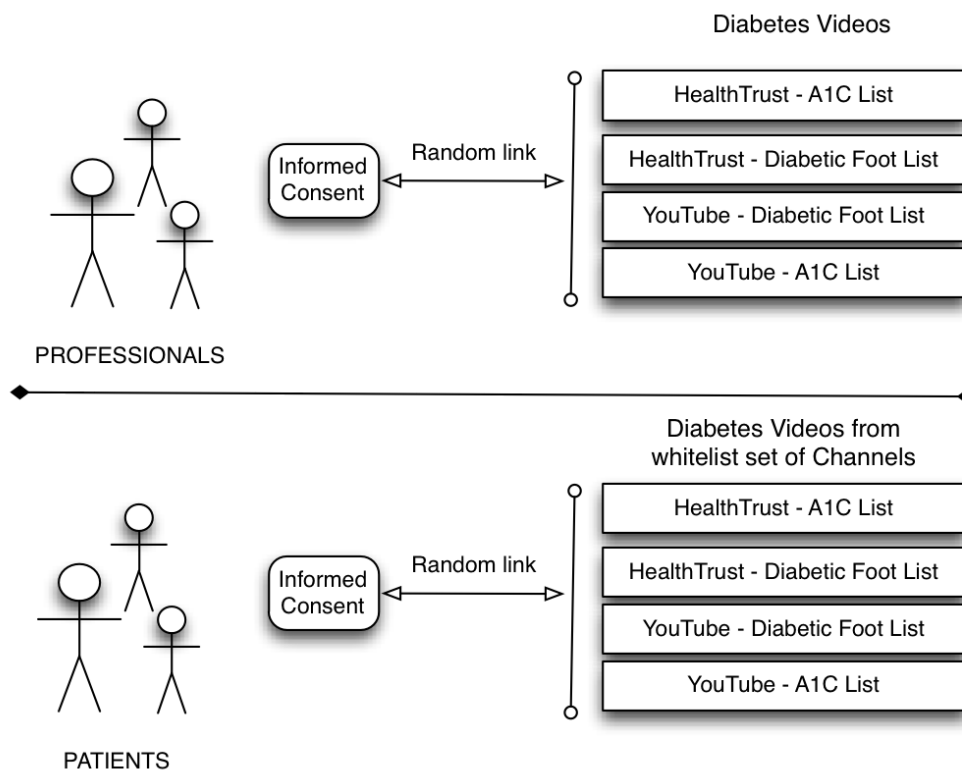
Figure 34: Rating process for list retrieved by HealthTrust (RQ4.Paper 2)

The results from the list of professionals were evaluated using the precision at K and DCG metrics, as in the previous study. We did not calculate either of these metrics for the health consumers, as they had a pre-filtered dataset to avoid showing them misleading videos. In addition, Pearson correlation[137]was used to study the correlation between average ratings and HealthTrust scores.

Results of HealthTrust for searching for diabetes videos

The evaluation of the search results retrieved using the HealthTrust approach was done with the Precision at K metric, and these were compared with the results returned by YouTube for the same queries. The relevance of a video was determined by an average video rating equal to or greater than 3.5 (value ranges from 1 to 5). As shown inTable 11, Precision at K (k = 3, k=7) was better in all the cases using HealthTrust vs YouTube search. As explained in detail in the paper, similar results were found by using the DCG metric.

Table 11: Precision at K for videos evaluated by professionals retrieved by HealthTrust and YouTube

Precision at K	A1C		Diabetes Feet	
	YouTube	HealthTrust	YouTube	HealthTrust
K = 3	66% (2)	100% (3)	33% (1)	66% (2)
K = 7	57% (4)	70% (5)	43% (3)	57% (4)

Another aspect evaluated in this study was the correlation between the average ratings for a given video and its HealthTrust score. We compared the normalized scores of HealthTrust and ratings for the videos using the Pearson correlation. Table 12 shows that there was a positive and statistically significant correlation in the case of videos about hemoglobin A_{1c} in the case of ratings by professionals (Pearson r_{10} = .646, P = .02). For the same query and ratings by patients, the results were not statistically significant, although they were close to significance (r_7 = .649, P = .06). In the case of diabetes foot videos, the results were inconclusive.

Table 12: Pearson correlation between ratings and HealthTrust scores

	A1C		Diabetes Feet	
	Pearson	P-Value	Pearson	P-Value
Professionals	r (10) =.646	.023 *	r(9) = .275	.413
Health consumers	r (7) = .648	.058	r(6) =-.019	.964

6.2.3. Discussion of HealthTrust Evaluation

<u>HealthTrust Metric Performance</u>

Our results indicate that social network analysis of health communities can be used to infer the quality of the content disseminated within the social network, in our case-study this was the diabetes community. The performance of that approach appeared to be equal or better than the results retrieved using a standard YouTube search.

Looking into the specific videos, some patterns emerged that may help explain why HealthTrust performance was good. Some disturbing videos, such as graphic surgery for diabetes foot (see Figure 35), were highly popular on YouTube. Watching those videos would be a traumatic experience for patients affected by diabetes, and that would explain why they had a very low HealthTrust score (nobody from the diabetes community would link to them) and also why the professionals would not recommend them. Looking at the comments of these disturbing videos, we see how popular they are among watchers who like gruesome videos. If the misleading-harmful videos have a community of viewers which is more active or larger than the diabetes community's, it is easy to understand why they are more highly ranked by YouTube. That would also explain why videos with little informative value, featuring singers with diabetes are highly represented in a YouTube search.

Figure 35. Highly ranked YouTube video about a diabetic foot infection (RQ4. Paper 2)

It is very difficult to compare the HealthTrust algorithm versus YouTube since the YouTube algorithm is not public, but the recommender algorithm has been published[33]. However, we assume that YouTube search is based on link analysis as are most search engines. General purpose search engines have been trying to cut down the effect of tightly knit communities that link heavily between them and thereby increase their reputation and ranking. This phenomenon is called the "tightly knit community effect" [99] and most search engines have been attempting to minimize this effect to reduce the influence of these communities and raise general public satisfaction. In contrast, HealthTrust reinforces the influence of the diabetes tightly knit community. Therefore, the HealthTrust approach serves the diabetes community better than it does general public satisfaction. The HealthTrust algorithm assesses health-related trust but not general trust within YouTube.

HealthTrust Weaknesses

The diabetes-centered approach also resulted in lower HealthTrust scores for video providers with a solid reputation within the health domain, such as governmental public health channels (eg, CDC Streaming Health's Channel). The reason is that these channels have weaker ties with the diabetes community since they do not engage in much social interaction with the diabetes community on YouTube. One possible improvement in the HealthTrust algorithm would be to develop a new HealthTrust algorithm adapted to compute trust within different health-related communities. For example, an algorithm that would make it possible to combine HealthTrust scores for both the health and diabetes communities.

In the study looking into the correlation between HealthTrust scores and average ratings, we did not find statistically significant results for all the cases. That is not surprising if we consider that the perceived quality of a video is not just related to its trust (RQ1.Paper3). Personal taste and other quality parameters play a major role. For example, the video *O is for outrage – Type 1 diabetes* (Figure 36) was given a higher rating by health consumers (average of 4.2) than by professionals (average of 2.75). One plausible explanation is that the video appeals to emotional aspects and that may be more attractive to patients than professionals. Professionals may prefer to recommend the most informative videos.

Another problem faced by HealthTrust concerned those videos with no public links (e.g. favorites, playlist). These videos might be liked by users who preferred not to make public

which videos they liked. The scarce availability of social data meant that some HealthTrust scores of videos were based only on the scores of the providers. An example of this problem is shown in the following two videos from the diabetic foot list for consumers: (1) *Baseball great Ron Santo & Diabetes--INCREDIBLE Story*, and (2) *Miami Ink's Darren Brass: Tattoos and Diabetes*. Both videos have the same HealthTrust score, as both are from the same diabetes channel, dLifedotcom. However, the Miami Ink video was less appealing to health consumers. The lack of enough social links is mainly due to the limitations for crawling social private data (unavailable with the YouTube API), such as which videos are being watched by users or who rated a particular video. The only way to overcome that barrier is to create our own video portal where we store that information.

Figure 36. O is for outrage - Type 1 diabetes (RQ4. Paper 2)

6.3. Limitations

In our evaluation, we created a simulated context for retrieving diabetes-related videos or channels. We need to be careful when generalizing about the findings, since, ideally, information retrieval evaluation should be performed in a real search engine with users

having personal information needs. Currently, I am working on the implementation of several video portals to capture more data for evaluation within a real user context.

Another limitation of our study is that it was based only on the case of diabetes. We cannot assume that the performance will be similar in other health topics for several reasons. Not all health-topics are likely to have an appropriate health community for the calculation of HealthTrust. For example, in the case of videos about erectile dysfunction, it would be surprising to find a community of users on YouTube similar to the one for diabetes. Another potential limitation is the calculation of HealthTrust in cases where there is a large sub-community promoting harmful content, such as pro-anorexia (RQ1.Paper5, RQ3.Paper1). In such cases, the misleading sub-community would have to be subtracted before the calculation of HealthTrust.

As mentioned earlier, HealthTrust only considers trust within the online community. That approach has the limitation that some trusted sources, which are not as active online, will have low HealthTrust scores. One possible way of overcoming this problem would be to obtain more clues about the trustworthiness of a channel by analyzing their metadata.

An additional limitation of this study was the lack of experiments using other algorithms to calculate authoritativeness scores. The use of the HITS algorithm is comparable with PageRank or more recently developed algorithms. Another technical limitation was the simplistic approach towards the community extraction instead of using more advanced and well-known algorithms, as used in RQ3.Paper1 and RQ3.Paper2. These advanced algorithms will be necessary in more complex cases, such as anorexia-related videos.

Finally, our study is limited to the particular context of online health videos on YouTube. We will have to evaluate HealthTrust in other video portals and also with other types of social media content (e.g. blogs and photos). The HealthTrust algorithm will require some adaptation in implementation since the data sources will be different. The type of content might affect the structure and dynamics of the health communities.

7. Discussion

7.1. Research Questions and Key Findings

7.1.1. RQ1: What are the characteristics of health social videos?

The research problem addressed in this research question (RQ1) was to acquire new knowledge about the context of health social videos. Due to the lack of background research in the areas covered in this dissertation, I had to participate in several research studies to grasp a better insight of the problems. The broad scope of the research question had to be narrowed down to a pair of addressable sub-questions.

RQ1.1: What are the motivations and challenges of lay people creating health videos?
Trust within social networks cannot be fully understood without looking at the individual. This question was chosen to gain a better understanding of the motivational aspects behind those laypersons sharing their health stories in social media. As explained in **RQ1. Paper 1**, one of the best ways to acquire that knowledge is to ask individuals about their experience in semi-structured interviews.

All the ePatients interviewed mentioned the importance of their online community for emotional support and also as a key factor for them continuing to share their experience. They also commonly mentioned how important quality was for them, since their online reputation was hard to gain and easy to lose.

K1.1: For experienced users, trust within the community is the key to their motivation to publish high quality content. (RQ1.Paper1)

RQ1.2 Do health videos contain relevant medical vocabulary in their textual metadata?

The presence of medical vocabulary describing health videos is of particular interest since it can be used to model videos and ultimately facilitate content retrieval. However, very little was known about the characteristics and nature of medical vocabulary within health video metadata.

In **RQ1.Paper 2**, we studied the presence of personal health information in videos' comments. We found that around 20% of the comments of videos from patients with multiple sclerosis contained personal private health information.

In **RQ1.Paper 3**, the focus was not on the disclosure of health information but rather on the use of medical thesauri in videos' metadata. Only a modest percentage of tags describing the videos were in fact found to be standard medical thesauri.

The answer to this research question is positive, at least for the two studies in which I participated. However, based on the results from those two studies, I concluded that modeling content based on textual language could be a feasible approach, but many problems would have to be overcome such as privacy issues, tag spamming and the use of jargon.

K1.2: Textual metadata may be of very heterogeneous quality, but it still contains a lot of relevant health information for modeling (RQ1.Paper 2, RQ1.Paper 3)

RQ1.3: What are the quality features of online health videos?

A major challenge in this dissertation was the lack of clear guidelines about what is meant by a high quality, and trustworthy video. Each time I asked this question to a different person I got a different answer. In addition, the research literature was equally inconclusive.

The literature study presented in **RQ1.Paper 4** was meant to answer this question. A systematic search of the literature was performed to create a matrix with the most common cited quality feature for health videos. The findings suggest that quality is a highly complex aspect for health videos, encompassing a wide range of aspects such as audio quality. The most common feature was related to the trustworthiness of the content provider.

Although there can be no categorical answer, the study I participated in has been highly cited since many authors have used it as a guideline for the quality evaluation of online health videos.

K1.3: Quality of health videos is a multidimensional concept, reliability of the content and provider is very important (
RQ1.Paper 3)

RQ1.4: Do misleading and informative online videos about anorexia have different characteristics?

General purpose search engines rely heavily on popularity metrics, such as favorites, number of views, etc. The study presented in RQ1.Paper5 was designed to explore differences in popularity metrics between misleading and informative videos about anorexia. The number of views was found to be higher for informative videos, but favoring and commenting ratios (divided by total views) was higher for misleading videos.

K1.4: Common popularity metrics, such as favoring ratio, correlate negatively with the trustworthiness of anorexia-related videos. (RQ1.Paper5).

7.1.2. RQ2: What are the technical solutions for extracting and modeling health social media?

Ultimately, this research question (RQ2) cannot be answered satisfactorily since it is too broad. The objective of this question was to create new knowledge by means of systematically reviewing different approaches for extracting and modeling social media within the health domain.

In the review paper (RQ2.Paper 1), we identified the problems and advantages of using several methods to extract knowledge from health social networks. Social network analysis was identified as one of the most promising approaches. The literature review of RQ2.Paper 1 provided the following findings:

- **K2.1**: Most technical solutions for modeling social media have shortcomings in the health domain due to textual analysis complexities, privacy issues and popular but harmful content. Link and Social Network Analysis is promising but has not been studied in sufficient detail in the health domain. These are some of the sub-findings:
  - Textual content poses a major challenge due to the complexity of dealing with jargon, vocabulary gaps, spamming, etc. In addition, textual analysis can raise privacy concerns.
  - Metadata is heterogeneous, especially user profiles (including Personal health records), descriptions, etc.

- o Images, Audio and video analysis requires heavy processing power due to their complexity.
- o Collaborative filtering is an interesting option; the problem is the high rating of misleading (but popular) content.
- o Links and social network analysis have been very successfully applied in web information retrieval, but little is known about its application in the health domain

7.1.3. RQ3: How Social Network Analysis is used to extract information about the characteristics of health social media?

In RQ,3 I studied whether social network analysis could be used to infer the trustworthiness of health social media. The findings seemed to support the idea that it could because the most experienced users in diabetes were the central nodes. In addition, social network metrics were the best parameters for predicting the misleading nature of content providers. However, a major problem is how to discern whether a particular sub-community is providing misleading, rather than informative, information.

RQ3.1: Can social network analysis be used to infer the misleading nature of social photos in the case of anorexia?

Some people might argue that social network analysis is of little use in scenarios where there is a misleading and highly active sub-community, such as with pro-anorexia communities. The study presented in RQ3.Paper 1 actually demonstrates the opposite. Social network metrics performed better than textual analysis for the classification of members disseminating harmful information about anorexia. In addition, we found that social network analysis could be used to separate misleading from informative social media providers. The main problem found in textual analysis for modeling anorexia-related content was that misleading information providers used similar vocabulary to trusted sources as a technique for gathering more views.

- **KF3.1**: the best predictors of users belonging to a misleading or informative sub-community are social network metrics. Tag-based classification performed worse due to the use of common vocabulary by both sub-communities. (RQ3. Paper 1)

RQ3.2: Do the central users' online health communities have different features to the rest of users in the case of diabetes?

The research presented in this dissertation has found that social network analysis can be used to identify users who are more likely to have longer life experience with the disease (and supposedly knowledge). In RQ3.Paper 2, we found that in the case of diabetes communities those more centric nodes according to social network analysis belonged to patients with at least a decade's experience of living with diabetes.

- **K3.2**: Most experienced users in the health area occupy central nodes in the social network of diabetes communities. (RQ3. Paper 2)

7.1.4. RQ4: Can trust-based metrics improve the retrieval of social videos about diabetes?

In RQ4, I explored whether it is possible to improve the retrieval of social videos about diabetes using metrics that capture trust within a health social network. I designed a metric called HealthTrust, which is calculated using social network metrics. The metric was used in two studies aimed at improving the retrieval of diabetes videos and also relevant publishers of videos.

RQ4.1: Can a metric of trustworthiness within a health community be used to retrieve relevant trustworthy providers of diabetes videos?

In the first study of HealthTrust, we found that the metric HealthTrust was better than YouTube for avoiding unrelated and misleading video providers. Therefore, we can conclude that within the context of diabetes videos on YouTube the answer to this question is affirmative. I need to highlight that performance was not that good in the sense that some good video providers were not highly ranked by HealthTrust mainly due to their low social interaction.

- **K4.1**: The search for diabetes videos and channels using a social network metric, such as HealthTrust, had positive results compared to a standard YouTube search. Therefore, we can assume that social network features may be used to infer a degree of content quality within health communities.

RQ4.2: Can a metric of trustworthiness within a health community be used to search for relevant trustworthy diabetes videos?

HealthTrust performed better than YouTube in retrieving diabetes videos. In line with the study of diabetes providers, the main improvements were due to the capability of HealthTrust to avoid misleading and gruesome videos. Similar issues were found regarding the penalization in the HealthTrust algorithm of relevant videos for health authorities that were less active in the social network.

- **K4.2**: The HealthTrust approach was especially effective in filtering out the more misleading and potentially harmful content.

7.2. Main Limitations

The main limitation of this study lies in the impossibility of generalizing the findings. As I have stated throughout the dissertation, each health community has been found to have a different structure and network dynamics. It might have been possible to address a specific health problem (e.g. diabetes) and therefore possibly have a narrower answer to all the research questions. Instead, I decided to perform my research using case studies that have the advantage of producing more generalizable results.

The main contribution of the dissertation, the HealthTrust Study, has only been tested in a simulated environment within a particular health community and with one type of content. This is definitely only an approximation: a real setting evaluation will be required to conclude the evaluation of HealthTrust. Furthermore, the computing performance of the algorithms involved in HealthTrust has not been tested. Thus, as a conclusion we can only say that HealthTrust represents a feasible approach for the retrieval of health social media.

Another aspect of vital importance is that my dissertation demonstrates just how complex the process of health information seeking can be, because each community and case-study had unique characteristics. Thus the application of an approach such as HealthTrust will require adaptation to each particular health problem in order to maximize its performance. A possible solution would be to combine HealthTrust with personalization techniques.

7.3.Implications

Public Health: it is well known that access to health information is crucial for professionals and patients who need to take informed decisions on a daily basis. In this dissertation, I have collaborated with other researchers to gain more insight into how health information is disseminated and trusted, and the way it evolves in social networks. That knowledge should be taken into consideration by health authorities before they invest resources in social media. For example, we found that pro-anorexia online communities are very effective in reaching those searching for health information about eating disorders. The understanding of the online social dynamics of pro-anorexia communities may help to increase the chances of success for online interventions regarding eating disorders.

In this dissertation, we have seen how good content from a health point of view guarantees neither popularity nor visibility in social media. The outreach of health social media depends on multiple factors (e.g. social dynamics, metadata quality). Therefore, online public health authorities should take into account all those factors (or the ones that have been shown to be more relevant) when designing online public health campaigns. The common approach of simply relying on the generation of more online content is less efficient and thus a waste of resources.

Policy Making: the promising results of HealthTrust shows that low investment, like this PhD, in improving online health information retrieval can yield good results. Therefore, I recommend that more resources be directed to research of Online Health Information Retrieval. Another aspect to consider is the feasibility of social network analysis for identifying providers of harmful information (e.g. promotion of anorexia). Policy makers can enforce the adoption of specific health-related filtering tools in web browsers to protect the most vulnerable population (e.g. minors).

ICT Research: this dissertation shows the potential of algorithms specifically designed for the online health sphere to improve results versus more general approaches. This dissertation is one of the few looking from a computing perspective into the specific challenges of finding relevant and trustworthy online health information. While modest, these very promising results, and the enormous potential societal impact, should

encourage others in ICT research to address the challenge of finding trustworthy online health information.

When I started the PhD, I believed that the online health domain was very complex. I was aware that "good" information from a technical point of view could be misleading and harmful. After seven years of research, I have discovered that the scenario is far more complex than I had imagined and the problem needs to be approached from multiple perspectives to understand it. In my dissertation, I had to collaborate with anthropologists, physicians, patients, caregivers, nurses, etc. ICT research in this domain requires a fairly complex multidisciplinary approach. That extra complexity for ICT research in this domain is compensated by multiple and societally important research challenges.

7.4. Recommendations for Future Research

This PhD work has increased our knowledge of how health information flows in social media, with a special focus on the role of the different online communities. The research presented here has identified the mismatch between the quality of online health information and results retrieved by general web information retrieval solutions. I identified that social network dynamics play a major role in the visibility of health information. However, our comprehension of the dynamics and processes involved is very limited. For example, we do not really understand why misleading information such as anti-vaccination or pro-anorexia are so highly visible on social media platforms. A better understanding of misinformation content providers in social media would help us to improve the definition of algorithms for information retrieval or define better strategies for different health organizations when it comes to disseminating content in social media.

Another area of special interest in relation to my dissertation concerns the process of information seeking in the health domain. Very few studies have looked into how health consumers create their search queries as part of their information seeking process. Without that knowledge it is impossible to provide effective tools to help consumers (and professionals) find trustworthy and relevant content. I am collaborating with Dr. EldYov-Tom from Microsoft Research in a new study addressing those issues in the context of online vaccination information.

In my dissertation I decided to focus on online health videos due to their popularity among patients and professionals. However, that is a simplification of the current use of social media. Users watch videos on YouTube that are shared in Facebook and YouTube. In many cases, those videos might link to presentations in Slideshare, blogs, etc. A future research line should look into how to apply HealthTrust in different types of social media, or with a combination of different types of content.

8. Conclusions

RQ1. What are the characteristics of health social videos?

"I promise nothing complete; because any human thing supposed to be complete, must not for that very reason infallibly be faulty." - 1851, Herman Melville, *Moby-Dick, Chapter 32*

After many years studying health social media I find many analogies with evolutionary biology. In this dissertation, I have been looking at making "life" easier for patients and health professionals living in the health social media ecosystem. Therefore, I needed to study the ecosystem of health social media in great depth.

As in nature, each time a new innovation emerges in social media some users adapt these innovations to achieve their unique goals. Every single user of social media lives as an individual organism; their similar peers form a population, and they coexist with others in an ecosystem that is part of a more general biosphere. Many patients embrace social media technologies as a way of communicating, creating content and socializing with fellow patients and healthcare professionals. Users adapt these technologies to their unique needs and context, thus making it very hard to draw conclusions about what the main characteristics are of health social media content and videos in particular.

In nature, a desert can be a perfect living place for some organisms or a death-zone for others. Thus the desirable characteristics of a living place will vary across organisms. My research has made me realize that in the case of health social media and videos in particular the situation is very similar. Health professionals tend to consider as very important the characteristic of how credible the author of the content is and the clinical validity of the content. Many patients consider emotional and affective characteristics (e.g. how positive, friendly, enjoyable) as important. Information retrieval experts are concerned about the quality of the data describing the content (aka metadata) to facilitate the search process. In addition, as new features are brought into social media (e.g. wearable technologies with Google Glass) new characteristics will emerge with different levels of importance for each individual in the health social media context. My research provides a snapshot of many important characteristics that need to be considered in health videos. In a nutshell, these social videos need to be from a trusted source and they have to be enjoyable for the intended audience. Many characteristics can be used to measure those aspects (e.g. who made the video, when was it made, quality of the content, popularity among the intended audience).

RQ2. Are there technical solutions for modeling health social media?

Taking into consideration all the important characteristics identified in **RQ1**, we needed to get an overview of all the different technical solutions to model health social media. In biology there are many techniques for measuring the characteristics of a living organism. In **RQ2.Paper 1**, I identified dozens of types of techniques that have been used to model social media content and users, some of which have already been applied in the health domain. Of those, the most promising were those involving textual processing (Natural Language Processing) and social network analysis. Natural Language Processing faced serious ethical issues, the risk of being used against the privacy of users. Social Network Analysis was found to be less challenging in that sense, because the scope of social network analysis is mainly limited to characterizing communities and the reputation of its members and content. Also Social Network Analysis is one of the best techniques for measuring reputation and trust, which in turn is one of the most important characteristics in health social media content.

RQ3. How can Social Network Analysis be used to extract information about the characteristics of health social media?

Social Network Analysis has been used for many decades to explore socialization processes. For example, Social Network Analysis can provide clues about how diseases spread in the case of an epidemic outbreak. Also, since we influence each other's behavior it is possible to study the dissemination of health lifestyles. In addition, by studying the social network we can identify the leader, the facilitator, or the outsiders.

In my dissertation, I decided to get involved in two studies looking into the social network characteristics in diabetes and eating disorders (i.e. anorexia). Diabetes communities are representative in social media, since they are homogenous and successful social networks. In **RQ3.Paper 2**, I learned that the most reputable members of diabetes social networks are those with more experience living with the disease. Thus, there was a link between knowledge and reputation in the social network. In contrast, online communities about anorexia are well known for promoting dangerous information (i.e. pro-anorexia). In **RQ3.Paper 1**, we found that the anorexia social network in Flickr was in fact an ecosystem of intermingled and antagonistic sub-communities promoting either trustworthy information about the diseases or promoting anorexia as a lifestyle. In **RQ3.Paper 1**, we also found that social network data was better than textual data for classifying members of those sub-communities.

As in nature, the study of the interactions between individuals, communities and the ecosystem turned out to be of crucial importance to understand the dynamics and trust within the entire ecosystem. It is also clear that each community and ecosystem is unique in its structure and dynamics, so my research gives only a small glimpse of the health social media biosphere.

RQ4. Can trust-based metrics improve the retrieval of social videos about diabetes?

"Beauty is in the eye of the beholders" – English Expression

I decided to explore the feasibility of using trust-metrics within health communities to search for health social media content. The rationale for that approach is that trust is in many cases nurtured within a community, thus *trust in the eyes of the health community* might be very significant for finding trustworthy health content for that particular community.

The HealthTrust study (**RQ4.Paper 1** and **RQ.Paper 2**) was designed to test whether a trust-based metrics could be used to identify diabetes videos that were highly recommended by both patients and professionals. The HealthTrust approach turned out to perform better than simply searching for diabetes videos on YouTube. Although that study is very limited in scope it does support the feasibility of the HealthTrust approach.

YouTube Search, as an example of any general web search, aims to help the entire Internet biosphere. A general web search is intended to optimize information retrieval for the largest number of users. In contrast, the HealthTrust approach focuses on serving the particular ecosystem of a health community. In the case of diabetes, YouTube retrieved frightening videos about toe amputations that are of interest to many viewers on YouTube. These videos had no connections within the diabetes community and they were therefore not listed using the HealthTrust search.

As for future work, there is a very clear limitation in the HealthTrust approach that needs to be overcome. HealthTrust is designed for a particular community. As in nature, different communities are interlinked and co-exist with other communities of living organisms who belong to that same ecosystem. For example, videos from other communities (e.g. general health users) were not represented in the HealthTrust results although some of them were clearly good. Another potential problem is the application of HealthTrust in communities that are sub-divided in different antagonistic sub-groups (e.g. pro-anorexia vs pro-recovery, RQ3.Paper 1). A possible way of overcoming this

limitation would be to design a general HealthTrust for the entire health ecosystem that combined trust metrics for each community (e.g. DiabetesTrust, SclerosisTrust).

References

1. Abdul SS, Lin C-W, Scholl J, et al. Facebook use leads to health-care reform in Taiwan. *Lancet*. 2011;377(9783):2083-4. doi:10.1016/S0140-6736(11)60919-7.

2. Ache KA, Wallace LS. Human papillomavirus vaccination coverage on YouTube. *Am J Prev Med*. 2008;35(4):389-92. doi:10.1016/j.amepre.2008.06.029.

3. Agichtein E, Castillo C, Donato D, Gionis A, Mishne G. Finding high-quality content in social media. In: *Proceedings of the international conference on Web search and web data mining - WSDM '08*. New York, New York, USA: ACM Press; 2008:183. doi:10.1145/1341531.1341557.

4. Almeida J, Goncalves MA, Figueiredo F, Pinto H, Belem F. On the Quality of Information for Web 2.0 Services. *IEEE Internet Comput*. 2010;14(6):47-55. doi:10.1109/MIC.2010.102.

5. Anand SS, Mobasher B. Introduction to intelligent techniques for Web personalization. *ACM Trans Internet Technol*. 2007;7(4):18-es. doi:10.1145/1278366.1278367.

6. Arcelus J, Mitchell AJ, Wales J, Nielsen S. Mortality Rates in Patients With Anorexia Nervosa and Other Eating Disorders: A Meta-analysis of 36 Studies. *Arch Gen Psychiatry*. 2011;68(7):724-731. Available at: http://archpsyc.ama-assn.org/cgi/content/abstract/68/7/724.

7. Backinger CL, Pilsner AM, Augustson EM, Frydl A, Phillips T, Rowden J. YouTube as a source of quitting smoking information. *Tob Control*. 2011;20(2):119-22. doi:10.1136/tc.2009.035550.

8. Bardone-Cone AM, Cass KM. What does viewing a pro-anorexia website do? an experimental examination of website exposure and moderating effects. *Int J Eat Disord*. 2007;40(6):537-548. Available at: http://dx.doi.org/10.1002/eat.20396.

9. Basagoiti I, ed. *Alfabetización en Salud: de la información a la acción*. Valencia, Spain: Instituto Itaca, Universidad Politecnica de Valencia; 2014:539. Available at: http://www.salupedia.org/alfabetizacion/. Accessed May 6, 2014.

10. Van de Belt TH, Berben SAA, Samsom M, Engelen LJLPG, Schoonhoven L. Use of social media by Western European hospitals: longitudinal study. *J Med Internet Res*. 2012;14(3):e61. doi:10.2196/jmir.1992.

11. Van De Belt TH, Engelen LJLPG, Berben SAA, Schoonhoven L. Definition of Health 2.0 and Medicine 2.0: a systematic review. *J Med Internet Res*. 2010;12(2):e18. doi:10.2196/jmir.1350.

12. Benito L. "Apps" dirigidas a prevenir patologías en los jóvenes - DiarioMedico.com. *D Med*. 2013. Available at: http://www.diariomedico.com/2013/01/31/area-profesional/entorno/apps-dirigidas-prevenir-patologias-en-jovenes. Accessed May 6, 2014.

13.   Bennett E. 2010 O 19. HSNL. Hospital Social Network List. *Found In Cache*. 2010. Available at: http://ebennett.org/hsnl/. Accessed November 9, 2011.

14.   Bernstam E V, Walji MF, Sagaram S, Sagaram D, Johnson CW, Meric-Bernstam F. Commonly cited website quality criteria are not effective at identifying inaccurate online information about breast cancer. *Cancer*. 2008;112(6):1206-13. doi:10.1002/cncr.23308.

15.   Bonander J. Personally tailored health information: a health 2.0 approach. *Med 20*. 2008.

16.   Brin S, Page L. The anatomy of a large-scale hypertextual Web search engine. *Comput Networks ISDN Syst*. 1998;30(1-7):107-117. doi:10.1016/S0169-7552(98)00110-X.

17.   Buettcher S, Clarke CLA, Cormack G V. *Information Retrieval: Implementing and Evaluating Search Engines*. The MIT Press; 2010:632. Available at: http://www.amazon.com/Information-Retrieval-Implementing-Evaluating-Engines/dp/0262026511. Accessed February 10, 2014.

18.   Burgess J, Green J. *YouTube: Online Video and Participatory Culture (DMS - Digital Media and Society)*. Polity; 2013. Available at: http://www.amazon.com/YouTube-Participatory-Culture-Digital-Society-ebook/dp/B00CFHJAS4. Accessed February 7, 2014.

19.   Burton S, Morris R, Dimond M, et al. Public health community mining in YouTube. In: *Proceedings of the 2nd ACM SIGHIT symposium on International health informatics - IHI '12*. New York, New York, USA: ACM Press; 2012:81. doi:10.1145/2110363.2110376.

20.   Carberry S, de Rosis F. Introduction to special Issue on "Affective modeling and adaptation." *User Model User-Adapt Interact*. 2008;18(1-2):1-9.

21.   CDC. Centers for Disease Control and Prevention. Social media guidelines and best practices. YouTube and other online video. 2012. Available at: http://www.cdc.gov/SocialMedia/Tools/guidelines/pdf/onlinevideo.pdf. Accessed April 18, 2013.

22.   Chau C. YouTube as a participatory culture. *New Dir Youth Dev*. 2010;2010(128):65-74. doi:10.1002/yd.376.

23.   Chin A, Chignell M. Automatic detection of cohesive subgroups within social hypertext: A heuristic approach. *New Rev Hypermedia Multimed*. 2008;14(1):121-143. Available at: http://dx.doi.org/10.1080/13614560802357180.

24.   Chomutare T, Arsand E, Fernandez-Luque L, Lauritzen J, Hartvigsen G. Inferring Community Structure in Healthcare Forums. An Empirical Study. *Methods Inf Med*. 2013;52(2). doi:10.3414/ME12-02-0003.

25.   Chou W-YS, Hunt Y, Folkers A, Augustson E. Cancer survivorship in the age of YouTube and social media: a narrative analysis. *J Med Internet Res*. 2011;13(1):e7. doi:10.2196/jmir.1569.

26.   Christakis NA, Fowler JH. The spread of obesity in a large social network over 32 years. *N Engl J Med*. 2007;357(4):370-9. doi:10.1056/NEJMsa066082.

27. Clauset A, Newman M, Moore C. Finding community structure in very large networks. *Phys Rev E*. 2004;70(6):066111. doi:10.1103/PhysRevE.70.066111.

28. Cobb NK, Graham AL, Abrams DB. Social network structure of a large online community for smoking cessation. *Am J Public Health*. 2010;100(7):1282-9. doi:10.2105/AJPH.2009.165449.

29. Cohen J. A Coefficient of Agreement for Nominal Scales . *Educ Psychol Meas* . 1960;20 (1):37-46. Available at: http://epm.sagepub.com/content/20/1/37.short.

30. Cottrell GW, Metcalfe J. EMPATH: face, emotion, and gender recognition using holons. 1990:564-571. Available at: http://dl.acm.org/citation.cfm?id=118850.105194. Accessed February 21, 2012.

31. Custers K, Van den Bulck J. Viewership of pro-anorexia websites in seventh, ninth and eleventh graders. *Eur Eat Disord Rev*. 2009;17(3):214-219. Available at: http://dx.doi.org/10.1002/erv.910.

32. Darmoni SJ, Dahamna B, Roth-Berghofer TR. Seal of transparency heritage in the CISMeF quality-controlled health gateway. *BMC Med Inform Decis Mak*. 2004;4:15. doi:10.1186/1472-6947-4-15.

33. Davidson J, Livingston B, Sampath D, et al. The YouTube video recommendation system. In: *Proceedings of the fourth ACM conference on Recommender systems - RecSys '10*. New York, New York, USA: ACM Press; 2010:293. doi:10.1145/1864708.1864770.

34. Dawson AL, Hamstra AA, Huff LS, et al. Online videos to promote sun safety: results of a contest. *Dermatology Reports*. 2011;3(1):e9. doi:10.4081/dr.2011.e9.

35. DeBronkart D. *¡Dejad que los pacientes ayuden! (Spanish Edition)*. (DeBronkart D, Fernandez-Luque L, eds.). CreateSpace Publishing; 2014:122. Available at: http://www.amazon.com/¡Dejad-pacientes-ayuden-Spanish-Edition/dp/1495345416. Accessed May 6, 2014.

36. Definition of Authoritative. *Merriam-Webster Dict*. 2014. Available at: http://www.merriam-webster.com/dictionary/authoritative. Accessed February 7, 2014.

37. Definition of Trust. *Merriam-Webster Dict*. 2014. Available at: http://www.merriam-webster.com/dictionary/trust. Accessed February 1, 2014.

38. Denecke K, Brooks E. Web science in medicine and healthcare. *Methods Inf Med*. 2013;52(2):148-51. Available at: http://www.ncbi.nlm.nih.gov/pubmed/23508344. Accessed February 7, 2014.

39. Diaz A. YouTube, una herramienta para la salud... o para el desastre | Tecnología Médica | elmundo.es. *El Mundo*. 2011. Available at: http://www.elmundo.es/elmundosalud/2011/12/26/tecnologiamedica/1324926130.html. Accessed May 6, 2014.

40.   Duda RO, Hart PE, Stork DG. *Pattern Classification (Pt.1)*. Wiley-Interscience; 2000:680. Available at: http://www.amazon.com/Pattern-Classification-Pt-1-Richard-Duda/dp/0471056693. Accessed February 3, 2014.

41.   Dunn AG, Westbrook JI. Interpreting social network metrics in healthcare organisations: a review and guide to validating small networks. *Soc Sci Med*. 2011;72(7):1064-8. doi:10.1016/j.socscimed.2011.01.029.

42.   Dutta-Bergman M. Trusted online sources of health information: differences in demographics, health beliefs, and health-information orientation. *J Med Internet Res*. 5(3):e21. doi:10.2196/jmir.5.3.e21.

43.   ECDC. *Communication on immunisation - building trust*. Stockholm: European Centre for Disease Prevention and Control; 2012. Available at: http://www.ecdc.europa.eu/en/publications/publications/ter-immunisation-and-trust.pdf. Accessed February 1, 2014.

44.   Ellis SE, Speroff T, Dittus RS, Brown A, Pichert JW, Elasy TA. Diabetes patient education: a meta-analysis and meta-regression. *Patient Educ Couns*. 2004;52(1):97-105. doi:10.1016/S0738-3991(03)00016-8.

45.   Esquivel A, Meric-Bernstam F, Bernstam E V. Accuracy and self correction of information received from an internet breast cancer list: content analysis. *BMJ*. 2006;332(7547):939-42. doi:10.1136/bmj.38753.524201.7C.

46.   Eysenbach G, Powell J, Kuss O, Sa E-R. Empirical studies assessing the quality of health information for consumers on the world wide web: a systematic review. *JAMA*. 287(20):2691-700. Available at: http://www.ncbi.nlm.nih.gov/pubmed/12020305. Accessed February 23, 2012.

47.   Eysenbach G, Yihune G, Lampe K, Cross P, Brickley D. Quality management, certification and rating of health information on the Net with MedCERTAIN: using a medPICS/RDF/XML metadata structure for implementing eHealth ethics and creating trust globally. *J Med Internet Res*. 2000;2(2 Suppl):2E1.

48.   Eysenbach G. An ontology of quality initiatives and a model for decentralized, collaborative quality management on the (semantic) World-Wide-Web. *J Med Internet Res*. 2001;3(4). Available at: 10.2196/jmir.3.4.e34.

49.   Eysenbach G. Medicine 2.0: Social Networking, Collaboration, Participation, Apomediation, and Openness. *J Med Internet Res*. 2008;10(3):e22+. Available at: 10.2196/jmir.1030.

50.   Fat MJL, Doja A, Barrowman N, Sell E. YouTube videos as a teaching tool and patient resource for infantile spasms. *J Child Neurol*. 2011;26(7):804-9. doi:10.1177/0883073811402345.

51.   Fernández-Alemán JL, Seva-Llor CL, Toval A, Ouhbi S, Fernández-Luque L. Free Web-based Personal Health Records: An Analysis of Functionality. *J Med Syst*. 2013;37(6):9990. doi:10.1007/s10916-013-9990-z.

52.   Fernandez-Luque L, Elahi N, Grajales FJ. An analysis of personal medical information disclosed in YouTube videos created by patients with multiple sclerosis. *Stud Health Technol Inform*. 2009;150:292-6. Available at: http://www.ncbi.nlm.nih.gov/pubmed/19745316. Accessed December 13, 2011.

53.   Fernandez-Luque L, Karlsen R, Krogstad T, Burkow TM, Vognild LK. Personalized health applications in the Web 2.0: the emergence of a new approach. *Conf Proc IEEE Eng Med Biol Soc*. 2010;2010:1053-6. doi:10.1109/IEMBS.2010.5628087.

54.   Fernandez-Luque L, Karlsen R, Melton GB. HealthTrust: A Social Network Approach for Retrieving Online Health Videos. *J Med Internet Res*. 2012;14(1). doi:10.2196/jmir.1985.

55.   Fernandez-Luque L, Karlsen R, Vognild LK. Challenges and opportunities of using recommender systems for personalized health education. *Stud Health Technol Inform*. 2009;150:903-7. Available at: http://www.ncbi.nlm.nih.gov/pubmed/19745443. Accessed October 19, 2011.

56.   Fernández-Luque L, Traver V, Ochoa A, Grajales FJ. *La nueva sanidad y los nuevos pacientes en la web 2.0 [New health and new patients in the Web 2.0]*.; 2010:53. Available at: http://static.diariomedico.com/docs/2010/07/12/sanidad_pacientes_web_2_0.pdf. Accessed May 6, 2014.

57.   Figueiredo F, Belém F, Pinto H, et al. Evidence of quality of textual features on the web 2.0. In: *Proceeding of the 18th ACM conference on Information and knowledge management - CIKM '09*. New York, New York, USA: ACM Press; 2009:909. doi:10.1145/1645953.1646070.

58.   Fleiss JL, Cohen J. The equivalence of weighted kappa and the intraclass correlation coefficient as measures of reliability.

59.   Fogg BJ, Swani P, Treinen M, et al. *What makes Web sites credible?* New York, New York, USA: ACM Press; 2001:61-68. doi:10.1145/365024.365037.

60.   Gabarron E, Fernandez-Luque L, Armayones M, Lau AY. Identifying Measures Used for Assessing Quality of YouTube Videos with Patient Health Information: A Review of Current Literature. *Interact J Med Res*. 2013;2(1):e6. doi:10.2196/ijmr.2465.

61.   Gabarrón E, Fernández-Luque L. *Libro Blanco ePaciente (translation of ePatient White Paper)*. Boston, MA, USA; 2011:126. Available at: http://e-patients.net/u/2011/11/Libro-blanco-de-los-e-Pacientes.pdf. Accessed May 6, 2014.

62.   Gagliardi A, Jadad AR. Examination of instruments used to rate quality of health information on the internet: chronicle of a voyage with an unclear destination. *BMJ*. 2002;324(7337):569-73. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=78993&tool=pmcentrez&rendertype=abstract. Accessed February 23, 2012.

63.   Gavin J, Rodham K, Poyer H. The Presentation of "Pro-Anorexia" in Online Group Interactions . *Qual Heal Res* . 2008;18 (3):325-333. Available at: http://qhr.sagepub.com/content/18/3/325.abstract.

64.   De Gemmis M, Lops P, Semeraro G, Basile P. Integrating tags in a semantic content-based recommender. In: *Proceedings of the 2008 ACM conference on Recommender systems - RecSys '08*. New York, New York, USA: ACM Press; 2008:163. doi:10.1145/1454008.1454036.

65.   Getoor L, Diehl CP. Link mining. *ACM SIGKDD Explor Newsl*. 2005;7(2):3-12. doi:10.1145/1117454.1117456.

66.   Good BM, Tennis JT, Wilkinson MD. Social tagging in the life sciences: characterizing a new metadata resource for bioinformatics. *BMC Bioinformatics*. 2009;10:313. Available at: 10.1186/1471-2105-10-313.

67.   Gooding LF, Gregory D. Descriptive analysis of YouTube music therapy videos. *J Music Ther*. 2011;48(3):357-69. Available at: http://www.ncbi.nlm.nih.gov/pubmed/22097103. Accessed April 18, 2013.

68.   Goodman LA. Snowball Sampling. *Ann Math Stat*. 1961;32(1):148-170. Available at: http://projecteuclid.org/euclid.aoms/1177705148. Accessed May 6, 2014.

69.   Gou L, Chen H-H, Kim J-H, Zhang X (Luke), Giles CL. SNDocRank. In: *Proceedings of the 19th international conference on World wide web - WWW '10*. New York, New York, USA: ACM Press; 2010:1103. doi:10.1145/1772690.1772825.

70.   Greenberg L, D'Andrea G, Lorence D. Setting the public agenda for online health search: a white paper and action agenda. *J Med Internet Res*. 2004;6(2):e18. doi:10.2196/jmir.6.2.e18.

71.   Greene JA, Choudhry NK, Kilabuk E, Shrank WH. Online social networking by patients with diabetes: a qualitative evaluation of communication with Facebook. *J Gen Intern Med*. 2011;26(3):287-92. doi:10.1007/s11606-010-1526-3.

72.   Gyöngyi Z, Garcia-Molina H, Pedersen J. Combating web spam with trustrank. 2004:576-587. Available at: http://dl.acm.org/citation.cfm?id=1316689.1316740. Accessed February 23, 2012.

73.   Gyongyi Z, Garcia-Molina H. Web Spam Taxonomy. *First Int Work Advers Inf Retr Web (AIRWeb 2005)*. 2005. Available at: http://ilpubs.stanford.edu:8090/771/1/2005-9.pdf. Accessed February 1, 2014.

74.   Harper FM, Moy D, Konstan JA. Facts or friends?: distinguishing informational and conversational questions in social Q&\#38;A sites. *CHI '09 Proc 27th Int Conf Hum factors Comput Syst*. 2009:759-768. Available at: http://doi.acm.org/10.1145/1518701.1518819.

75.   Harris PR, Sillence E, Briggs P. Perceived threat and corroboration: key factors that improve a predictive model of trust in internet-based health information and advice. *J Med Internet Res*. 2011;13(3):e51. doi:10.2196/jmir.1821.

76.   Harshbarger JL, Ahlers-Schmidt CR, Mayans L, Mayans D, Hawkins JH. Pro-anorexia websites: what a clinician should know. *Int J Eat Disord*. 2009;42(4):367-70. doi:10.1002/eat.20608.

77. Hersh W. *Information Retrieval: A Health and Biomedical Perspective (Health Informatics)*. Springer; 2010:486. Available at: http://www.amazon.com/Information-Retrieval-Biomedical-Perspective-Informatics/dp/1441926976. Accessed February 7, 2014.

78. Himmel W, Reincke U, Michelmann HW. Text mining and natural language processing approaches for automatic categorization of lay requests to web-based expert forums. *J Med Internet Res*. 2009;11(3):e25. doi:10.2196/jmir.1123.

79. Hoek HW, van Hoeken D. Review of the prevalence and incidence of eating disorders. *Int J Eat Disord*. 2003;34(4):383-396. Available at: http://dx.doi.org/10.1002/eat.10222.

80. Hughes JP, Quraishi MS. YouTube resources for the otolaryngology trainee. *J Laryngol Otol*. 2012;126(1):61-2. doi:10.1017/S0022215111002337.

81. Jadad AR, Gagliardi A. Rating health information on the Internet: navigating to knowledge or to Babel? *JAMA*. 1998;279(8):611-4. Available at: http://www.ncbi.nlm.nih.gov/pubmed/9486757. Accessed February 23, 2012.

82. Johnsen E, Fernandez-Luque L, Hagen R. Det Sosiale Nettet og helse [Health and the Social Web]. In: Tjora A, Sandaunet A, eds. *Digitale pasienter*. Oslo, Norway: Gyldendal akademisk; 2010. Available at: http://www.gyldendal.no/Faglitteratur/Medisin/Profesjon/Digitale-pasienter. Accessed May 6, 2014.

83. Kaelber DC, Jha AK, Johnston D, Middleton B, Bates DW. A research agenda for personal health records (PHRs). *J Am Med Inform Assoc*. 2008;15(6):729-736. Available at: 10.1197/jamia.M2547.

84. Kannan R. Experts warn against trusting health information on social media - The Hindu. *The Hindu Times*. 2013. Available at: http://www.thehindu.com/news/cities/chennai/experts-warn-against-trusting-health-information-on-social-media/article4434572.ece. Accessed May 6, 2014.

85. Kaplan AM, Haenlein M. Users of the world, unite! The challenges and opportunities of Social Media. *Bus Horiz*. 2010;53(1):59-68. doi:10.1016/j.bushor.2009.09.003.

86. Keelan J, Pavri-Garcia V, Tomlinson G, Wilson K. YouTube as a source of information on immunization: a content analysis. *JAMA*. 2007;298(21):2482-4. doi:10.1001/jama.298.21.2482.

87. Keselman A, Logan R, Smith CA, Leroy G, Zeng-Treitler Q. Developing informatics tools and strategies for consumer-centered health communication. *J Am Med Inf Assoc*. 2008;15(4):473-483. Available at: 10.1197/jamia.M2744.

88. Keselman A, Smith CA, Divita G, et al. Consumer health concepts that do not map to the UMLS: where do they fit? *J Am Med Inf Assoc*. 2008;15(4):496-505. Available at: 10.1197/jamia.M2599.

89. Khan SA, Cohall A, Kukafka R. A tag based recommendation engine to suggest information resources in an online community for health promotion. *AMIA Annu Symp Proc*.

2008:1002. Available at: http://www.ncbi.nlm.nih.gov/pubmed/18999253. Accessed December 13, 2011.

90. Kleinberg JM. Authoritative sources in a hyperlinked environment. *J ACM*. 1999;46(5):604-632. doi:10.1145/324133.324140.

91. Kobsa A. Privacy-enhanced personalization. *Commun ACM*. 2007;50(8):24-33.

92. Kumar R, Raghavan P, Rajagopalan S, Tomkins A. Trawling the Web for emerging cyber-communities. *Comput Networks*. 1999;31(11-16):1481-1493. doi:10.1016/S1389-1286(99)00040-7.

93. Kummervold PE, Chronaki CE, Lausen B, et al. eHealth trends in Europe 2005-2007: a population-based survey. *J Med Internet Res*. 2008;10(4):e42. doi:10.2196/jmir.1023.

94. Kwon YH, Lobo N da V. Age Classification from Facial Images. *Comput Vis Image Underst*. 1999;74(1):1-21. doi:10.1006/cviu.1997.0549.

95. Lau AY, Gabarron E, Fernandez-Luque L, Armayones M. Social media in health – what are the safety concerns for health consumers? Authors. *Heal Inf Manag J*. 2012:1-11.

96. Lau AYS, Coiera EW. Impact of web searching and social feedback on consumer decision making: a prospective online experiment. *J Med Internet Res*. 2008;10(1):e2. doi:10.2196/jmir.963.

97. Lau AYS, Siek KA, Fernandez-Luque L, et al. The Role of Social Media for Patients and Consumer Health. Contribution of the IMIA Consumer Health Informatics Working Group. *Yearb Med Inform*. 2011;6(1):131-8. Available at: http://www.ncbi.nlm.nih.gov/pubmed/21938338. Accessed October 3, 2011.

98. Lee CM, Narayanan SS. Toward detecting emotions in spoken dialogs. *Speech Audio Process IEEE Trans*. 2005;13(2):293-303. Available at: 10.1109/TSA.2004.838534.

99. Lempel R, Moran S. The stochastic approach for link-structure analysis (SALSA) and the TKC effect. *Comput Networks*. 2000;33(1-6):387-401. doi:10.1016/S1389-1286(00)00034-7.

100. Liberati A, Altman DG, Tetzlaff J, et al. The PRISMA statement for reporting systematic reviews and meta-analyses of studies that evaluate health care interventions: explanation and elaboration. *PLoS Med*. 2009;6(7):e1000100. doi:10.1371/journal.pmed.1000100.

101. Liu H, Maes P. InterestMap: Harvesting Social Network Profiles for Recommendations. *Beyond Pers*. 2004.

102. Lustria MLA, Cortese J, Noar SM, Glueckauf RL. Computer-tailored health interventions delivered over the Web: review and analysis of key components. *Patient Educ Couns*. 2009;74(2):156-73. doi:10.1016/j.pec.2008.08.023.

103.   Lyons EJ, Mehl MR, Pennebaker JW. Pro-anorexics and recovering anorexics differ in their linguistic Internet self-presentation. *J Psychosom Res*. 2006;60(3):253-6. doi:10.1016/j.jpsychores.2005.07.017.

104.   Ma X, Chen G, Xiao J. Analysis of an online health social network. In: *Proceedings of the ACM international conference on Health informatics - IHI '10*. New York, New York, USA: ACM Press; 2010:297. doi:10.1145/1882992.1883035.

105.   Madden M. *Online Video*.; 2007. Available at: http://www.pewinternet.org/~/media/Files/Reports/2007/PIP_Online_Video_2007.pdf. pdf. Accessed January 29, 2014.

106.   Marino S, Sessa E, Di Lorenzo G, et al. Quantitative analysis of pursuit ocular movements in Parkinson's disease by using a video-based eye tracking system. *Eur Neurol*. 2007;58(4):193-197. Available at: 10.1159/000107939.

107.   Massa P, Avesani P. Trust-aware recommender systems. *RecSys '07 Proc 2007 ACM Conf Recomm Syst*. 2007:17-24. Available at: http://doi.acm.org/10.1145/1297231.1297235.

108.   Mayer MA, Karkaletsis V, Stamatakis K, et al. MedIEQ-Quality labelling of medical web content using multilingual information extraction. *Stud Health Technol Inform*. 2006;121:183-90. Available at: http://www.ncbi.nlm.nih.gov/pubmed/17095816. Accessed February 23, 2012.

109.   McDaid D, Park A. *Online Health: Untangling the Web*.; 2011:30. Available at: http://www.bupa.com/media/44806/online_20health_20-_20untangling_20the_20web.pdf. Accessed March 8, 2013.

110.   McLuhan M, Lapham LH. *Understanding Media: The Extensions of Man*. The MIT Press; 1994:392. Available at: http://www.amazon.com/Understanding-Media-The-Extensions-Man/dp/0262631598. Accessed March 13, 2013.

111.   Mishra N. Online Social Networking: Can it Be Used for Risk Behavior Surveillance. *AMIA Spring Congr*. 2008.

112.   Mislove A, Marcon M, Gummadi KP, Druschel P, Bhattacharjee B. Measurement and analysis of online social networks. In: *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement - IMC '07*. New York, New York, USA: ACM Press; 2007:29. doi:10.1145/1298306.1298311.

113.   Mobasher B, Cooley R, Srivastava J. Automatic personalization based on Web usage mining. *Commun ACM*. 2000;43(8):142-151. doi:10.1145/345124.345169.

114.   *Modern Information Retrieval [Paperback]*. Addison Wesley; 1st edition; 1999:544. Available at: http://www.amazon.com/Modern-Information-Retrieval-Ricardo-Baeza-Yates/dp/020139829X. Accessed June 11, 2014.

115.   Mooers CN. The next twenty years in information retrieval. In: *Papers presented at the the March 3-5, 1959, western joint computer conference on XX - IRE-AIEE-ACM '59 (Western)*. New York, New York, USA: ACM Press; 1959:81-86. doi:10.1145/1457838.1457853.

116.   Moreno MA, Parks M, Richardson LP. What are adolescents showing the world about their health risk behaviors on MySpace? *MedGenMed*. 2007;9(4).

117.   Mulveen R, Hepworth J. An interpretative phenomenological analysis of participation in a pro-anorexia internet site and its relationship with disordered eating. *J Health Psychol*. 2006;11(2):283-296. Available at: 10.1177/1359105306061187.

118.   Mundt JC, Snyder PJ, Cannizzaro MS, Chappie K, Geralts DS. Voice acoustic measures of depression severity and treatment response collected via interactive voice response (IVR) technology. *J Neurolinguistics*. 2007;20(1):50-64. Available at: DOI: 10.1016/j.jneuroling.2006.04.001.

119.   Murugiah K, Vallakati A, Rajput K, Sood A, Challa NR. YouTube as a source of information on cardiopulmonary resuscitation. *Resuscitation*. 2011;82(3):332-4. doi:10.1016/j.resuscitation.2010.11.015.

120.   Norris ML, Boydell KM, Pinhas L, Katzman DK. Ana and the Internet: a review of pro-anorexia websites. *Int J Eat Disord*. 2006;39(6):443-447. Available at: 10.1002/eat.20305.

121.   O'Donovan J, Smyth B. Trust in recommender systems. *IUI*. 2005:167-174.

122.   Pandey A, Patni N, Singh M, Sood A, Singh G. YouTube as a source of information on the H1N1 influenza pandemic. *Am J Prev Med*. 2010;38(3):e1-3. doi:10.1016/j.amepre.2009.11.007.

123.   Paton C, Bamidis PD, Eysenbach G, Hansen M, Cabrer M. Experience in the Use of Social Media in Medical and Health Education. Contribution of the IMIA Social Media Working Group. *Yearb Med Inform*. 2011;6(1):21-9. Available at: http://www.ncbi.nlm.nih.gov/pubmed/21938320. Accessed April 8, 2013.

124.   Paton C, Hansen M, Fernandez-Luque L, Lau AYS. Self-Tracking, Social Media and Personal Health Records for Patient Empowered Self-Care. Contribution of the IMIA Social Media Working Group. *Yearb Med Inform*. 2012;7(1):16-24. Available at: http://www.ncbi.nlm.nih.gov/pubmed/22890336. Accessed February 27, 2013.

125.   Peintner B, Jarrold W, Vergyriy D, Richey C, Tempini MLG, Ogar J. Learning diagnostic models using speech and language measures. *Conf Proc IEEE Eng Med Biol Soc*. 2008;2008:4648-4651. Available at: 10.1109/IEMBS.2008.4650249.

126.   Pestian JP, Matykiewicz P, Grupp-Phelan J, Lavanier SA, Combs J, Kowatch R. Using natural language processing to classify suicide notes. *AMIA Annu Symp Proc*. 2008:1091. Available at: http://www.ncbi.nlm.nih.gov/pubmed/19006447. Accessed February 21, 2012.

127.   Plaza J. Las TIC deben quitarse el "sambenito" del alto coste - DiarioMedico.com. *D Med*. 2014. Available at: http://www.diariomedico.com/2014/03/13/area-profesional/gestion/tic-deben-quitarse-sambenito-alto-coste. Accessed May 6, 2014.

128.   Publicly Private and Privately Public: Social Networking on YouTube. Available at: http://jcmc.indiana.edu/vol13/issue1/lange.html. Accessed August 6, 2012.

129. Purcell GP, Wilson P, Delamothe T. The quality of health information on the internet. *BMJ*. 2002;324(7337):557-8. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1122494&tool=pmcentrez &rendertype=abstract. Accessed January 31, 2014.

130. Qiu F, Cho J. Automatic identification of user interest for personalized search. In: *Proceedings of the 15th international conference on World Wide Web - WWW '06*. New York, New York, USA: ACM Press; 2006:727. doi:10.1145/1135777.1135883.

131. Rainie L. *Tagging*. Washington, DC, USA: Pew Internet. American Life Project; 2007:9. Available at: http://www.pewinternet.org/files/old-media//Files/Reports/2007/PIP_Tagging.pdf.pdf.

132. Reputation - Definition and More from the Free Merriam-Webster Dictionary. Available at: http://www.merriam-webster.com/dictionary/reputation. Accessed February 7, 2014.

133. Revelle W. *Procedures for Psychological, Psychometric and Personality Research*.; 2014. Available at: http://cran.r-project.org/web/packages/psych/psych.pdf. Accessed February 10, 2014.

134. Rivero-Rodriguez A, Konstantinidis ST, Sanchez-Bocanegra CL, Fernandez-Luque L. A health information recommender system: Enriching YouTube health videos with Medline Plus information by the use of SnomedCT terms. In: *Proceedings of the 26th IEEE International Symposium on Computer-Based Medical Systems*. IEEE; 2013:257-261. doi:10.1109/CBMS.2013.6627798.

135. Roberts Strife S, Rickard K. The Conceptualization of Anorexia: The Pro-Ana Perspective. *Affilia*. 2011;26(2):213-217. doi:10.1177/0886109911405592.

136. Robichaud P, Hawken S, Beard L, et al. Vaccine-critical videos on YouTube and their impact on medical students' attitudes about seasonal influenza immunization: A pre and post study. *Vaccine*. 2012;30(25):3770-3763. doi:10.1016/j.vaccine.2012.03.074.

137. Rodgers JL, Nicewander WA. Thirteen Ways to Look at the Correlation Coefficient. *Am Stat*. 1988;42(1):59. doi:10.2307/2685263.

138. Sabater J, Sierra C. Review on Computational Trust and Reputation Models. *Artif Intell Rev*. 2005;24(1):33-60.

139. Saigí Rubió F, Sanchez-Bocanegra CL, Fernández-Luque L, Sanchez-Laguna F. *Salud móvil y software libre*. (Saigí Rubió F, ed.). Barcelona: Editorial UOC; 2013:118. Available at: http://www.editorialuoc.cat/saludmvilysoftwarelibre-p-1185.html?cPath=1. Accessed May 6, 2014.

140. Sajadi KP, Goldman HB. Social networks lack useful content for incontinence. *Urology*. 2011;78(4):764-7. doi:10.1016/j.urology.2011.04.074.

141. Schafer J Ben, Frankowski D, Herlocker JL, Sen S. Collaborative Filtering Recommender Systems. *Adapt Web*. 2007;4321:291-324.

142. Shabbir SA, Fernandez-Luque L, Li Y-C, Hsu M-H, Lee P-S, Chiu W-T. Leaks and threats to patient data. *BMJ*. 2011;342(feb11 1):d871-d871. doi:10.1136/bmj.d871.

143. Sillence E, Briggs P, Fishwick L, Harris P. Trust and mistrust of online health sites. In: *Proceedings of the 2004 conference on Human factors in computing systems - CHI '04*. New York, New York, USA: ACM Press; 2004:663-670. doi:10.1145/985692.985776.

144. Smith CAA, Wicks PJ. PatientsLikeMe: Consumer Health Vocabulary as a Folksonomy. *AMIA Annu Symp Proc*. 2008:682-686.

145. Sonawane A. Using Apache Lucene to search text. *IBM Dev*. 2009. Available at: http://www.ibm.com/developerworks/library/os-apache-lucenesearch/. Accessed February 7, 2014.

146. Sood A, Sarangi S, Pandey A, Murugiah K. YouTube as a source of information on kidney stone disease. *Urology*. 2011;77(3):558-62. doi:10.1016/j.urology.2010.07.536.

147. Statistics - YouTube. Available at: http://www.youtube.com/yt/press/statistics.html. Accessed April 18, 2013.

148. Steinberg PL, Wason S, Stern JM, Deters L, Kowal B, Seigne J. YouTube as source of prostate cancer information. *Urology*. 2010;75(3):619-22. doi:10.1016/j.urology.2008.07.059.

149. Steinhaeuser K, Chawla N V. Identifying and evaluating community structure in complex networks. *Pattern Recognit Lett*. 2010;31(5):413-421. doi:10.1016/j.patrec.2009.11.001.

150. Stoltzman W. Toward a Social Signaling Framework: Activity and Emphasis in Speech. 2006.

151. Susannah Fox SJ. The Social Life of Health Information. 2009.

152. Syed-Abdul S, Fernandez-Luque L, Jian W-S, et al. Misleading Health-Related Information Promoted Through Video-Based Social Media: Anorexia on YouTube. *J Med Internet Res*. 2013;15(2):e30. doi:10.2196/jmir.2237.

153. Tang PC, Ash JS, Bates DW, Overhage JM, Sands DZ. Personal health records: definitions, benefits, and strategies for overcoming barriers to adoption. *J Am Med Inform Assoc*. 2006;13(2):121-126. Available at: 10.1197/jamia.M2025.

154. Theng Y-L, Goh LYQ, Tin MT, Sopra R, Kumar SKP. Trust cues fostering initial consumers' trust. In: *Proceedings of the 2nd ACM SIGHIT symposium on International health informatics - IHI '12*. New York, New York, USA: ACM Press; 2012:807. doi:10.1145/2110363.2110462.

155. Tian Y. Organ donation on Web 2.0: content and audience analysis of organ donation videos on YouTube. *Health Commun*. 2010;25(3):238-46. doi:10.1080/10410231003698911.

156. Traver V, Fernández-Luque L, Basagoiti I. *Salud 2.0 entre profesionales [Health 2.0 between professionals]*. (Traver V, Fernández-Luque L, Basagoiti I, eds.). Valencia, Spain: ITACA-Universidad Politecnica de Valencia; 2011:70.

157. Traver V, Fernandez-Luque L. *El ePaciente y las redes sociales (Spanish Edition)*. El epaciente y las redes sociales; 2011:320. Available at: http://www.amazon.com/ePaciente-las-redes-sociales-Spanish/dp/8469441108. Accessed February 16, 2012.

158. Traver V, Fernandez-Luque L. *El ePaciente y las redes sociales [ePatients and Social Networks]*. (Traver V, Fernandez-Luque L, eds.). Valencia, Spain: Fundacion Vodafone; 2011:365. Available at: http://www.amazon.com/ePaciente-las-redes-sociales-Spanish/dp/8469441108/ref=sr_1_1?ie=UTF8&qid=1328430949&sr=8-1. Accessed May 6, 2014.

159. Turner JC. Towards a cognitive redefinition of the social group. *Curr Psychol Cogn*. 1981;1(2):93-118.

160. Ubois J, Davidson J, Grobelnik M, Over P, Westerhof H. Video search. In: *Proceedings of the 19th international conference on World wide web - WWW '10*. New York, New York, USA: ACM Press; 2010:1339. doi:10.1145/1772690.1772918.

161. Vinciarelli A, Pantic M, Bourlard H. Social signal processing: Survey of an emerging domain. *Image Vis Comput*. 2009;27(12):1743-1759. Available at: DOI 10.1016/j.imavis.2008.11.007.

162. Wang J, Chng E, Xu C, Lu H, Tong X. Identify Sports Video Shots with "Happy" or "Sad" Emotions. *Multimed Expo, IEEE Int Conf*. 2006;0:877-880. Available at: http://doi.ieeecomputersociety.org/10.1109/ICME.2006.262641.

163. Wang Z, Walther JB, Pingree S, Hawkins RP. Health information, credibility, homophily, and influence via the Internet: Web sites versus discussion groups. *Health Commun*. 2008;23(4):358-68. doi:10.1080/10410230802229738.

164. Wasserman S, Faust K. *Social Network Analysis: Methods and Applications (Structural Analysis in the Social Sciences)*. Cambridge University Press; 1994:857. Available at: http://www.amazon.com/Social-Network-Analysis-Applications-Structural/dp/0521387078. Accessed February 4, 2014.

165. Whitfield L. Patient Opinion mashes up NHS Choices. *eHealthInsider*. 2008.

166. Wild S, Roglic G, Green A, Sicree R, King H. Global Prevalence of Diabetes: Estimates for the year 2000 and projections for 2030. *Diabetes Care*. 2004;27(5):1047-1053. doi:10.2337/diacare.27.5.1047.

167. Wu HC, Luk RWP, Wong KF, Kwok KL. Interpreting TF-IDF term weights as making relevance decisions. *ACM Trans Inf Syst*. 2008;26(3):1-37. doi:10.1145/1361684.1361686.

168. Xu Z, Fu Y, Mao J, Su D. Towards the Semantic Web: Collaborative Tag Suggestions. *WWW2006 Proc Collab Web Tagging Work*. 2006.

169. Yom-Tov E, Fernandez-Luque L, Weber I, Crain SP. Pro-anorexia and pro-recovery photo sharing: a tale of two warring tribes. *J Med Internet Res*. 2012;14(6):e151. doi:10.2196/jmir.2239.

170. Yom-Tov E. *HSD 2013 - ACM SIGIR Workshop on Health Search & Discovery*. (White R, Yom-Tov E, Horvitz E, Agichten E, Hersh B, eds.). Dublin, Ireland: Microsoft Research; 2013. Available at: http://research.microsoft.com/en-us/events/hsd2013/hsd2013_proceedings.pdf. Accessed February 7, 2014.

171. Youtube. About YouTube. 2012. Available at: http://www.youtube.com/t/about_youtube.

172. Zeng QT, Tse T, Divita G, et al. Term identification methods for consumer health vocabulary development. *J Med Internet Res*. 2007;9(1):e4. Available at: 10.2196/jmir.9.1.e4.

PAPERS

RQ1.Paper1

RQ1.Paper2

RQ1.Paper3

RQ1.Paper 4

RQ1.Paper 5

RQ2.Paper 1

RQ3.Paper 1

RQ3.Paper 2

RQ4.Paper 1

RQ4.Paper 2