Kemi-Tornio University of Applied Sciences

Applying Semantic Web Services Technologies in the eHealth Domain

Famakinwa Oluwasegun Michael

Bachelor's thesis of the Degree Programme in Business Information Technology

Bachelor of Business Administration

TORNIO 2013

ABSTRACT

KEMI-TORNIO UNIVERSITY OF APPLIED SCIENCES

Degree programme:	Business Information Technology
Writer:	Famakinwa, Oluwasegun Michael
Thesis title:	Applying Semantic Web Services Technologies in the eHealth Domain
Pages:	62
Date:	May 10, 2013
Thesis instructor:	Ryabov, Vladimir Ph.D.

This research focuses on exploring the application of Semantic Web Services technologies within the healthcare environment. More specifically, this research seeks to explore a possible integration of two separate disciplines, i.e. Semantic Web and eHealth, in order to create a functional system that will be beneficial to patients that seek health-related assistance on the Web.

The issue of integrating Semantic Web and eHealth is still quite novel, despite the fact that there have been significant research works done on the topics as separate disciplines. Terms such as Health 3.0, Medicine 3.0 and other similar terminologies have become popularised, as efforts to enhance healthcare services with Semantic Web Services technologies intensify. Therefore, the main objective of this research is to investigate and propose ways in which Semantic Web and eHealth can be integrated, particularly from the perspective of improving the quality, access and efficacy of healthcare worldwide.

Based on the research questions formulated to achieve the objectives of this research, the exploratory research methodology was employed. In order to ensure the academic validity of this thesis, this research derives from the theoretical materials of established research works in the disciplines of Semantic Web and eHealth. These theoretical materials, which are scientific sources in the form of printed publications and electronic books, were subjected to extensive critical analysis.

Therefore, on the basis of the findings obtained from the analysis of literature, this research proposes practical ideas to execute the integration of some Semantic Web Service technologies with aspects of eHealth systems. To ensure the continuity of this research work, future research directions are also suggested.

Keywords: Semantic Web, Semantic Web Services, eHealth, Ontology, Web 3.0, Health 3.0

CONTENTS

FIGURES

1 INTRODUCTION

The choice of this thesis topic is motivated here. The background of the thesis topic is also discussed. In addition, the objectives of the research work are explained, and the general structure of this thesis work is described in this chapter.

1.1 Motivation and background

The world today is increasingly dependent on technology for people's everyday activities. Since the advent of social media and social networking platforms such as Facebook and Twitter, the Internet has become a common means of communication for various activities, especially among the younger generation. This rise in Internet usage has also coincided with an era that has seen a steep rise in the number of mobile phone usage worldwide. Figure 1, adopted from Fox (2010), illustrates this change in Internet use by different age groups from 2000 to 2010.

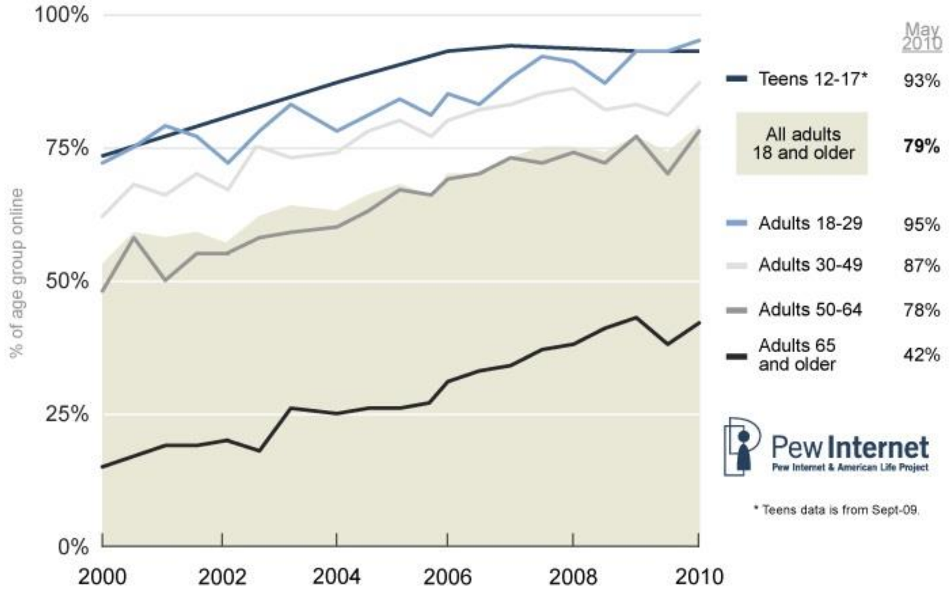

Figure 1. Change in Internet use by age (Fox 2010)

Information retrieval has now assumed an automated approach, with search engines such as Google and Bing, and online communities such as Wikipedia and YouTube being the dominant tools that are frequently consulted to obtain desired knowledge. A cursory glance at the Alexa top 500 global website rankings confirms this trend, as all of the aforementioned websites are currently listed in the top 20 based on the number of hits frequently received (Alexa Internet, Inc. 2013). Considering current Information and Communications Technology (hereinafter ICT) trends, the health sector stands to benefit substantially from integrating certain aspects of ICT in its operations, especially when one considers the immense opportunities that the Internet in particular offers in that regard. According to Jessen (2007), data and computational systems are being employed today to make ailments more predictable, to anticipate and prevent diseases, and to personalise treatments even further. Indeed, ICT has been deeply enmeshed in virtually all aspects of human activity in recent times.

The application of ICT in health care has grown rapidly in the last 15 years. Consequently, its potential to improve effectiveness and efficiency has been recognised by governments worldwide (Pagliari & Sloan & Gregor & Sullivan & Detmer & Kahan & Oortwijn & MacGillivray 2005.) The aim of ICT for Health, i.e. eHealth, is to improve the quality, access and efficacy of healthcare. eHealth describes the application of ICT across a whole range of functions that positively influence the health sector. To emphasise the importance and relevance of eHealth, it should be noted that the European Commission has been supporting research activities in the eHealth field for almost twenty years. These developments have since contributed to the emergence of an eHealth industry. (European Commission 2007, 8.)

It is paramount that healthcare is supported by systems that are consciously and accurately designed to produce care that is safe, effective, patient-centred, timely, efficient, and equitable. In view of this, it is believed that ICT possesses enormous potential to improve the quality of health care with regard to these six aims. There are various opportunities to apply ICT in the eHealth domain, such as facilitating access to clinical knowledge through understandable and reliable Web sites and online support groups, and the use of clinical decision support systems to tailor information according to the characteristics, genetic makeup, and specific conditions of an individual patient. Both patients and health personnel can also benefit through the use of Internet-based communication such as telemedicine and immediate access to automated clinical information, diagnostic tests, and treatment results to improve timeliness. It is equally of great importance to build ICT infrastructure to support evidence-based medical practice. This includes providing increasingly organised and reliable information sources on the Internet for both patients and health personnel, as this will bring about the much needed healthcare improvements. (National Research Council 2001, 164-165.)

According to Yu (2007, 3), the Internet has three main uses, i.e. search, integration and Web mining. Yu (2007, 7-8) however points out the limitations of the Internet in its current state by maintaining that the Internet is constructed in such a way that its documents only contain enough information for the computers to present them, not to understand them. In other words, the Internet only stores presentation of data elements but does not record their meanings in any form. As a result, the meaning is left to humans to decode and understand. Due to this limitation, the Semantic Web concept was conceived. The basic idea of the Semantic Web is to extend the current traditional Web by adding semantics into Web documents. The added semantics is expressed as structured information that can be read and understood by machines. Once this is accomplished, each Web page will contain not only information to instruct machines about how to display it, but also structured data to help machines to understand it. (Yu 2007, 8-9.)

This research topic was selected because there is a growing need for eHealth solutions around the world. The health sector has remained traditional for a long time, and it is high time the field embraced the technology that will spread medical assistance in a readily accessible manner. The call for integration of the Semantic Web with eHealth is justified when one considers that in health-related searches, context is highly important. If search results are poor, inaccurate, or incomplete, users can easily be misinformed and confused. The issue of information overload is also a problem. (Trzebucki 2008.) Therefore, the Semantic Web will move towards making content accessible by applications other than a web browser. It will become possible to build the next layer of intelligence into the Web, thereby allowing for both interaction and collaboration. One of the first to get involved in the development of the Semantic Web was the health industry. Recognising the potential benefit of integrating the Semantic Web with eHealth, the World Wide Web Consortium (hereinafter W3C) established the Health

Care and Life Sciences Interest Group (HCLSIG) to "develop, advocate for, and support the use of Semantic Web technologies across health care, life sciences, clinical research and translational medicine. These domains stand to benefit from intra- and inter-domain application of Semantic Web technologies as they depend on the interoperability of information from many disciplines." (Jessen 2007.)

According to a research conducted by Pew Research Centre's Internet & American Life Project, the Internet is trusted second only to a physician in health matters. More so, the Internet is frequently consulted as a first resource for health information. When one considers that the same Pew research established that 75% of individuals searching for health-related information do not verify their sources, this lack of verification measures becomes a problematic issue. (Fox 2011.) It is essential that information source and quality is verified for health information, since most users typically start their online health search through an Internet search engine before considering visiting a hospital or consulting a physician. Clearly there is a need for improved, accurate and in-depth health information retrieval online, which is what the Semantic Web assures. (Trzebucki 2008.) Considering this need, this thesis work emphasises the importance of integrating the Semantic Web with eHealth.

1.2 Objectives

The objective of this thesis is to explore the application of Semantic Web Services (hereinafter SWS) technologies within the healthcare environment. This objective is achieved by studying and understanding the underlying technologies behind the Semantic Web and eHealth concepts. Due to the exploratory nature of this research, it is important that core concepts are defined and understood adequately before attempting to assess their possible integration. The theoretical development and applications of the Semantic Web and eHealth are defined in detail in order to pave the way for the specific research objective, which is to integrate these concepts. This research achieves this objective by deriving from the theoretical materials of research works in the same disciplines.

More specifically, the core technologies that make up eHealth systems and the Semantic Web, as well as their functionalities are highlighted during the course of this research. Highlighting these technologies and their functionalities is necessary in order to investigate the possibilities of integration between eHealth and the Semantic Web. The question of what type of technologies can be employed in this integration, are addressed as a means to achieve this objective.

Based on the preceding steps, connections are drawn from the detailed study of the Semantic Web and eHealth which then allows for an investigation of how they can be integrated. As a practical outcome, this thesis identifies and analyses possible techniques to integrate SWS technologies with eHealth, such that the health sector is improved significantly, especially from the perspective of improving the quality, access and efficacy of healthcare, as previously discussed. This integration and subsequent improvement will ensure that credible health-related information and assistance are readily accessible to as many people that need them.

1.3 Structure of the thesis

This thesis is divided into seven chapters. The succeeding texts in this thesis are structured as follows. Chapter 2 defines the scope of this research, the research questions, and the research methodology in detail. Chapter 3 presents an introduction to the concept of eHealth, addresses the problem of eHealth literacy, and gives an overview of various eHealth tools. Chapter 4 introduces the Semantic Web and its underlying technologies, with a main focus on its application and impact on the healthcare industry. Chapter 5 discusses four scenarios demonstrating the application of SWS technologies to achieve a semantically-enabled eHealth domain. To demonstrate the practical relevance of this research, Chapter 6 describes existing projects that have successfully integrated the Semantic Web with eHealth. Finally, Chapter 7 discusses the conclusion of this research work.

2 RESEARCH SCOPE, QUESTIONS AND METHODOLOGY

The scope of this research work is discussed first in this chapter. The research questions formulated in order to achieve the objectives of this research are also defined. Lastly, the methodology employed for this research is explained.

2.1 Research scope

As was previously discussed, the general aim of this research is to focus on integrating SWS technologies with eHealth. The specific objectives of this research are to identify, analyse and propose ways to integrate the Semantic Web with eHealth by studying the technologies that form the basis of these disciplines.

For the purpose of this Bachelor's thesis, research was carried out with extensive consultations with relevant and most recent literature available. The scope of this research has been narrowed down to include basically the main concepts of eHealth and Semantic Web and the possible integration of the two disciplines. As this research can be categorised as a basic research, it excludes implementing any practical solutions from this integration, such as developing any software applications. The research work is entirely theoretical, with resources employed primarily from the analysis of literature and previous research works, and only goes as far as the extent of a Bachelor's thesis research work. Hence, this research does not assure that any proposed suggestions will work as predicted or assumed.

Additionally, the aim of this research is limited to exploring and providing suggestions for the integration of two separate disciplines, i.e. Semantic Web and eHealth. This research work was carried out with the assumption that the Semantic Web is already in full deployment. Therefore, any ideas, recommendations or suggestions provided therein worked consistently with that assumption.

2.2 Research questions

In order to achieve the objectives of this research, the following research questions are addressed.

1. What are the concepts of eHealth and the Semantic Web? What underlying technologies are these concepts comprised of?

   The eHealth and Semantic Web concepts need to be understood in order to explore and understand the potential ways these two disciplines can be integrated. Therefore, answers to this research question are discussed in Chapter 3 and Chapter 4 respectively. In these chapters, eHealth and the Semantic Web are both defined, followed by detailed descriptions and extended definitions of both concepts. The tools of eHealth and relevant underlying technologies governing the Semantic Web are also studied and analysed in details.

2. How can Semantic Web be integrated with eHealth? What technologies are needed to implement this integration?

   This research shows how Semantic Web tools can be used to support eHealth systems. At present, there is ongoing research into medical ontologies by established research centres, with a focus on how medical terminologies are represented efficiently. This kind of research reveals one way in which the Semantic Web can be integrated with eHealth, and consequently contribute towards improved healthcare. This particular research question is partially discussed in some subsections of Chapter 4 and fully answered in Chapter 5. With the assumption that the Semantic Web is in full deployment, this research investigates how such integration is possible and what technologies are needed to implement it.

3. What are the potential practical applications of the Semantic Web and eHealth integration?

This research illustrates how existing eHealth environments have incorporated SWS technologies by describing some health-related applications of the Semantic Web, and also identifying current projects that have successfully implemented this integration. One example is the creation of ontology-based applications for clinical usage. Another example is the deployment of Semantic Interoperability in eHealth systems. These examples are discussed in Chapter 6 of this thesis. Future research directions are also suggested for improving the integration of the Semantic Web and the eHealth domain.

2.3 Research methodology

The research methods will be discussed by finding support from McDaniel Jr. and Gates (2010). As the work is entirely theoretical, the research method employed here is the qualitative research method. While Semantic Web and eHealth by themselves are not entirely new concepts, an integration of both is a relatively novel idea. Therefore, an exploratory approach was used chiefly in this regard. McDaniel Jr. and Gates (2010, 43) state that the purpose of an exploratory research is "to obtain greater understanding of a concept or to help crystallise the definition of a problem. It is also used to identify important variables to be studied. Exploratory research is preliminary research, not the definitive research used to determine a course of action." On this basis therefore, exploratory research is relevant for this thesis work, as it aims to explore the disciplines of Semantic Web and eHealth, and consequently offer suggestions for their integration.

Furthermore, this research was conducted based on a collection of secondary data and in-depth analysis of literature from established scientific sources. These research activities were carried out extensively throughout the entire research process. Due to the nature of this research work, this technique is relevant. Tertiary sources were also employed to help shape critical analysis of the literature appropriately.

3 THE EHEALTH CONCEPT

To comprehend how eHealth can be integrated with the Semantic Web, it is imperative to understand the concept of eHealth in itself, as well as the tools it employs. This chapter presents an extended definition of the eHealth concept, explains the impact of the Internet in health-related searches, addresses the important issue of eHealth literacy, and describes the tools that make eHealth into a functioning system.

3.1 What is eHealth?

eHealth is a concept that only came to the fore in the year 2000, but has since become widespread (Pagliari et al. 2005). The concept is one that is widely used within various spheres, from individuals to academic institutions, professional bodies and funding organisations. Although there is yet to be a generally accepted clear definition, the concept has nonetheless become an acceptable coinage. (Oh & Rizo & Enkin & Jadad 2005.)

Oh et al. (2005) acknowledge that while it is impossible to find a "universally acceptable" and "applicable formal definition", the eHealth concept could still be better understood by reviewing a range of proposed meanings. This impossibility was also noted by Pagliari et al. (2005) who maintain that since there is a lack of consensus on the meaning of the concept especially among academics, policymakers, providers and consumers, definitions of eHealth vary with respect to the functions, stakeholders, contexts and theoretical issues targeted. Therefore, Oh et al. (2005) seek to address this lingering issue of uncertainty and determine the contextual usages of the concept by exploring various pieces of scientific literature for definitions. The result of this exploration is a systematic review of published definitions, three of which are presented below.

The most concise, if not simplistic, definition of eHealth is "the integration of the Internet into health care" (Watson 2004). However, McLendon (2000 as cited by Oh et al. 2005) offers a more expansive definition by stating that "eHealth refers to all forms of electronic health care delivered over the Internet, ranging from informational, educational and commercial 'products' to direct services offered by professionals, non-professionals, businesses or consumers themselves. eHealth includes a wide variety of the clinical activities that have traditionally characterised telehealth, but delivered through the Internet. Simply stated, eHealth is making healthcare more efficient, while allowing patients and professionals to do the previously impossible."

Perhaps, an even more encompassing definition than the one above is the one offered by Eysenbach (2001) when he defines eHealth as "an emerging field in the intersection of medical informatics, public health and business, referring to health services and information delivered or enhanced through the Internet and related technologies. In a broader sense, the concept characterises not only a technical development, but also a state-of-mind, a way of thinking, an attitude, and a commitment for networked, global thinking, to improve health care locally, regionally, and worldwide by using information and communication technology." This definition by Eysenbach (2001) is regarded as a "global definition" that is "well represented" (Pagliari et al. 2005). This is further affirmed by the fact that it is the most commonly cited definition on the Internet, as it has been adopted or referred to by at least 87 websites on the Internet (Oh et al. 2005).

It should be noted that a set of divergent concepts such as health, technology, and commerce are encompassed in the eHealth definitions above. It should also be noted that Health, as used in these definitions, refers expressly to healthcare as a process, rather than as an end result. Furthermore, in the definitions of eHealth presented, technology is viewed both as a tool to enable a process and as the paradigm of eHealth itself. Particularly noteworthy is that technology is represented as a means to expand, to assist, or to enhance human activities, rather than as an alternative to them. Overall, an attitude of optimism is reflected in the collective understanding of eHealth, as none of the published definitions suggests that eHealth may have any adverse effects. (Oh et al. 2005.)

Although there remains a latent understanding of the meaning of eHealth, its widespread usage suggests that it is an important concept indeed.

3.2 Internet use for health-related information

There has been a surge in the use of the Internet as a source of health information. Sorensen and Andreassen (2010) argue that the potential of using the Internet and other electronic media in promoting health and health care seems propitious, considering the large group of people that can be reached fast and cheaply. The perceived importance of the Internet is allegedly rising. According to Fox (2011), "eight in ten Internet users look online for health information, making it the third most popular online activity among all those included in the Pew Internet Project's surveys." Figure 2, adopted from Fox (2011), depicts varying Internet activities among the different age groups surveyed.

Activity	Millennials Ages 18-34	Gen X Ages 35-46	Younger Boomers Ages 47-56	Older Boomers Ages 57-65	Silent Gen. Ages 66-74	G.I. Gen. Age 75+	All online adults Age 18+
Go online	95%	86%	81%	76%	58%	30%	79%
For the following activities, the youngest and oldest cohorts may differ, but there is less variation between generations overall:							
Email	96	94	91	93	90	88	94
Use search engine	92	87	86	87	82	72	87
Look for health info	78	84	80	83	73	69	80
Get news	76	79	76	76	67	54	75
Buy a product	68	66	64	69	59	57	66

Figure 2. Pew Research Centre's Internet and American Life Project surveys (Fox 2011)

Although this particular study was conducted with a survey of a number of adults living in a specific geographical region i.e. the United States, it could also be indicative of the general Internet habits of Internet users in other parts of the world. The search for health-related information on the Internet often takes various forms. Users typically search for information on a specific disease or medical problem, certain medical treatment or procedure, hospitals or other medical facilities, health insurance, food safety, drug safety, environmental health hazards, pregnancy and childbirth, and many more (Fox 2011).

Considering this increasingly popular usage of the Internet for health-related purposes, Sorensen and Andreassen (2010) warn that there is still fragmented valid knowledge on how eHealth is influencing health care services and health users. Therefore, due to the growing trend of this particular Internet use, and the criticality of readily accessible health-related information on an unguarded territory such as the Internet, the issue of eHealth literacy needs to be addressed and elaborated upon.

3.3 eHealth literacy

Norman and Skinner (2006) define eHealth literacy as “the ability to seek, find, understand, and appraise health information from electronic sources and apply the knowledge gained to addressing or solving a health problem.” While the extent of information on the Web has skyrocketed, the ability to access the right information pales in comparison. Retrieving data in proper context is particularly crucial with health information, and on a general basis, the Internet is chaotic in this regard. (Trzebucki 2008.) As was discussed in the Introduction chapter, the Pew Research Centre’s Internet and American Life Project reported that three-quarters of individuals searching for health-related information on the Internet do not verify their sources (Fox 2011). This information is especially disconcerting when one considers the types of users that mostly employ the Internet as a source of health information. Bundorf, Wagner, Singer and Baker (2006) report that, “individuals with reported chronic conditions were more likely than those without to search for health information on the Internet. The uninsured, particularly those with a reported chronic condition, were more likely than the privately insured to search. Individuals with longer travel times for their usual source of care were more likely to use the Internet for health-related communication than those with shorter travel times.” Clearly, there is a need to make Internet users more aware of the implication of the tools they use to solve their various health problems.

To address these associated problems of eHealth literacy, Norman and Skinner (2006) assert that unlike other forms of literacy, eHealth literacy comprises six types of literacy skills which all combine to promote eHealth. These literacy skills are traditional literacy, information literacy, media literacy, health literacy, scientific literacy and computer literacy. Among these, media and computer literacies are peculiar to the Internet context. When synthesised, these six literacy types form the required skills to fully optimise the experiences of eHealth users. Figure 3, adopted from Norman and Skinner (2006), illustrates the relationship between these literacy skills in a lily model.

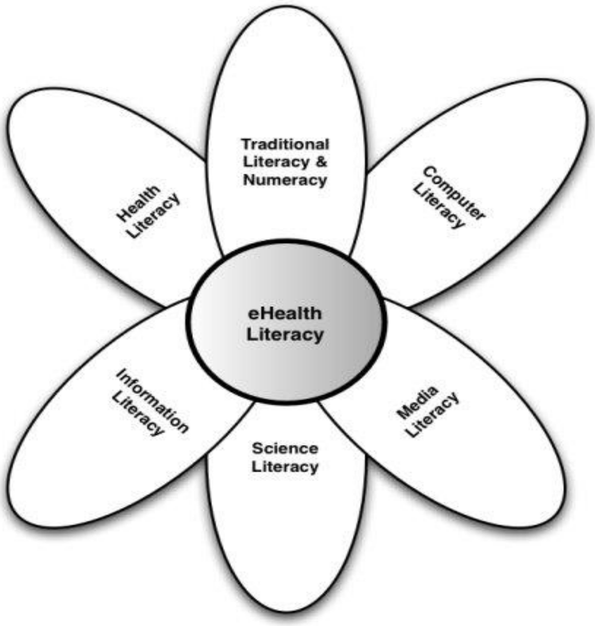

Figure 3. A lily model of eHealth literacy skills (Norman & Skinner 2006)

Norman and Skinner (2006) further assert that eHealth literacy of an individual is influenced by a number of factors. These factors include their educational background, health status, and motivation for seeking the information. Furthermore, Norman and Skinner (2006) argue that eHealth literacy is a "process-oriented skill that evolves over time as new technologies are introduced and the personal, social, and environmental contexts change." The aim of eHealth literacy is to empower individuals by allowing them to partake in health decisions that are contingent on eHealth resources. Norman and Skinner (2006) further advocate matching eHealth technologies to the skills of their intended users. By augmenting the users' working knowledge of computers to a level that tends to achieve health-related goals, and by designing systems with the users in mind, matching eHealth technologies to users' skills can be accomplished.

As was previously discussed, users with "serious health needs" and those facing "significant barriers" in accessing health care the traditional way often turn to the Internet for health information (Bundorf et al. 2006). Therefore, there is indeed a need to educate susceptible and chronically ill people, and to design technology in a manner conforming to more consumers (Neter & Brainin 2012). With this need recognised, the demand for the integration of eHealth with the Semantic Web becomes more valid.

3.4 eHealth tools

A discourse on the concept of eHealth will not be complete without an overview of the tools that support the primary elements of eHealth systems. In a report on Global Observatory for eHealth (2006), the World Health Organization (hereinafter WHO) segmented the tools of eHealth into different categories based on their respective functions. The categories are as follows: tools for professionals, tools to support health care provision, tools for health care and financial administration, tools for policy and population health care, tools for technical requirements and tools for citizens (World Health Organization 2006). Based on this categorisation, some specific eHealth tools are therefore described.

Electronic Health Records (eHR), which are also referred to as Electronic Medical Records (eMR), are used to support clinical actions by health professionals. They include information such as test results, medication and general clinical history of a patient. Through ICT, these records can be made promptly available to the authorised personnel providing care for the patient. (World Health Organization 2006.) In the eHR, clinical and administrative health care information about an individual's lifetime of health experiences are digitally stored. The purpose this digital storage serves is to support continuity of care, education and research, while also ensuring confidentiality. The eHR is a tool that is used to support the delivery of health care throughout all stages of care and linked through health telematic networks. (Silber 2003.)

Patient Information Systems (PIS) contain information about a hospitalised patient and are used to support both the administrative and clinical activities in a hospital. They are usually hospital-wide, but may be restricted to single or multiple departments. They do not usually contain multimedia data, thereby distinguishing them from an electronic health record system. They contain numeric and textual data about the patient in addition to the basic administrative data, which distinguishes them from hospital information systems. Hospital Information Systems (HIS) are computer-based information systems that support information processing within a hospital in areas such as administration, appointments, invoicing, planning, budgeting and personnel (World Health Organization 2006.)

General Practitioner Information Systems (GPIS) are ICT-based systems that support the work of a general practitioner (GP) or a primary health care practitioner. The variation in health care models makes functions required by countries quite different. Where the GP is part of a primary health care team, the system may also be known as a Primary Care Information System. The prime functions of this system are to manage and share data about patients. They often link to other health care systems such as invoicing, GP reimbursement or laboratory results reporting systems. (World Health Organization 2006.)

Some other tools include National electronic registries which are electronic databases of related records on specific medical subjects. They contain data on births, mortality, cancer, diabetes or other subjects of medical or epidemiological interest. Registries can be accessed by authorised users through the use of ICT. Similarly, National drug registries are electronic databases which contain national pharmaceutical information. The content varies depending on the purpose of the registry. Examples include databases of risks of exposure to drugs during pregnancy and potential drug interactions. In addition, Directories of healthcare professionals and institutions is another eHealth tool. They are electronic databases of individuals and institutions providing health care. These are usually searchable by location, specialisation, professional association or credentials. They are often associated with registration and accreditation status. (World Health Organization 2006.)

Other tools such as Decision Support Systems are employed in eHealth as well. They are automated or semi-automated systems that support decision-making in a clinical environment. Telehealth refers to the use of ICT to either support the provision of health care or as an alternative to direct professional care. It encompasses telemedicine and the use of remote medical expertise. Lastly, Geographical Information Systems are computer-based applications for capturing, integrating, analysing and displaying data related to geographic coordinates. (World Health Organization 2006.)

4 THE SEMANTIC WEB

In order to explore the ways in which eHealth and Semantic Web can be integrated, it is necessary to provide foundations for this integration by extensively defining the Semantic Web concept, as has been done with eHealth in the previous chapter. On this basis, this chapter discusses the Semantic Web technologies in detail. In keeping up with the objectives of this thesis, this chapter further explains Semantic Web Services (SWS) as used in the eHealth domain. Lastly, the impact of the Semantic Web on healthcare is discussed.

4.1 Semantic Web technologies

The Semantic Web is an extension of the current Web in which information is given well-defined meaning, facilitating computers and people to work in cooperation (Berners-Lee et al. 2001 as cited by Yu 2007, 8). Its exact aim is to harmonise semantic discrepancies in software systems by providing machine-interpretable semantics, and to “understand” ambiguous descriptions – thus achieving a new quality of intelligent and automated information processing in the web (Davies et al. 2006 as cited by Studer & Grimm & Abecker 2007, 3). In the words of Berners-Lee et al. (2001 as cited by Matthews 2005, 2), “the Semantic Web will bring structure to the meaningful content of Web pages, creating an environment where software agents roaming from page to page can readily carry out sophisticated tasks for users.” Fundamentally, the idea of the Semantic Web is to include a mechanism that will define semantics about resources and links on the Web, i.e. to imbue the Web with meaning. This will therefore make possible the automatic processing of the Web by the aforementioned software agents, rather than intervention by users. (Matthews 2005, 2.)

To implement the Semantic Web, a set of technologies, tools and standards which will form the basic building blocks of the system are needed. A variation of the Semantic Web layered architecture adopted from Matthews (2005, 4) showing these technologies, tools and standards is depicted in Figure 4.

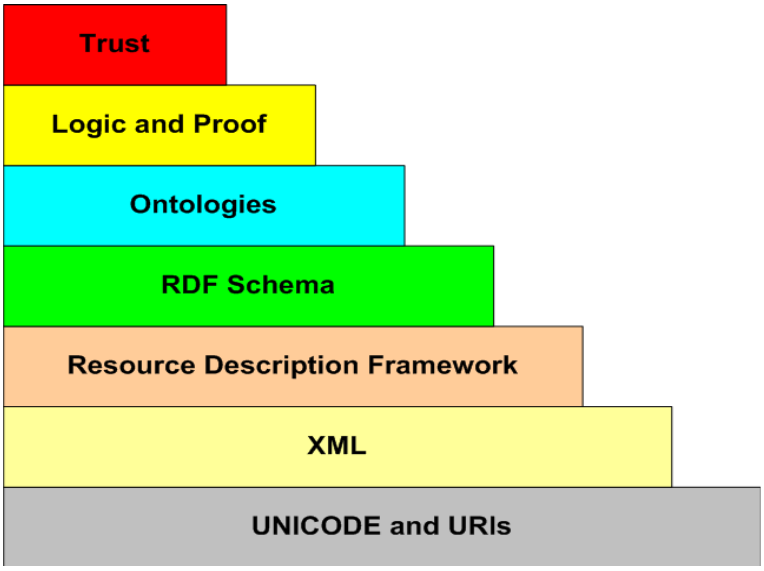

Figure 4. Semantic Web Layered Architecture (Matthews 2005, 4)

Of the various Semantic Web components contained in these layers, three are key technologies that form the basis of the Semantic Web. They are as follows: Resource Description Framework (RDF), which is a metadata representation framework that encodes meanings defined through ontologies; Ontology, which defines terms and their relationships; and the Extensible Markup Language (XML) which provides syntax and data structure on the Web (Bose & Sugumaran 2007, 223). These three technologies, along with the other components of the Semantic Web are discussed below.

The Resource Description Framework (RDF) is a language that is used to represent information about resources in the World Wide Web. It is used especially to represent metadata about Web resources, such as the title, author, and modification date of a Web page, copyright and licensing information about a Web document, or the availability schedule for some shared resource. (Manola & Miller 2004.) RDF, being the first layer of the Semantic Web "proper", is a simple metadata representation framework that makes use of Universal Resource Identifiers (URIs) to identify Web-based resources. RDF also uses a graph model to describe relationships between resources. (Matthews 2005, 4-5.) Practically, the intended use of RDF is for circumstances in which information needs to be processed by the application, rather than just being displayed.

Basically, RDF provides a common framework which allows applications to exchange information on the Web without losing any meaning. (Manola & Miller 2004.)

Additionally, the Resource Description Framework Schema (RDFS), which is a Vocabulary Description language, is an extension to RDF. RDFS enables vocabularies to be formed for RDF metadata. (Grimm & Hitzler & Abecker 2010, 83.) According to Matthews (2005, 5), RDFS provides a simple reasoning framework for analysing types of resources, while also describing classes of resources and properties between them in the basic RDF model. When RDFS is combined with RDF, the resultant term becomes RDF(S). This combination provides a simple ontology language for conceptual modelling with basic capabilities for inference. (Grimm et al. 2010, 83.)

However, despite its usefulness as an ontology language, RDFS also has its limitations. For this reason, it is sometimes classified as a representation language for supposed "lightweight" ontologies. As a result, more expressive representation languages such as the Web Ontology Language (OWL) are used for more sophisticated applications. (Hitzler & Krötzsch & Rudolph 2010, 47.)

Ontology is defined as "a formal explicit specification of a shared conceptualisation of a domain of interest" (Gruber 1993 as cited by Grimm et al. 2010, 69). This is the "dominating definition" of ontology within the Semantic Web community, as it captures several features such as formality, explicitness, being shared, conceptuality and domain specificity (Grimm et al. 2010, 69). Within the context of the Semantic Web, Ontology is said to be domain-specific, as it only defines a group of terms in a given domain and the relationship among them (Yu 2007, 90). Ontology is an important component of the Semantic Web architecture because it provides a way to reuse domain knowledge. It also makes domain assumptions explicit, and together with ontology description languages, it provides a way to encode knowledge and semantics such that machines are able to understand. More importantly, it makes automatic large-scale machine processing to be possible. (Yu 2007, 92.)

In 2004, the Web Ontology Language (OWL) was adopted as a World Wide Web Consortium (W3C) recommended standard for the modelling of ontologies. Since then, it has become an increasingly widespread language for creating ontologies in many application domains. One reason for its popularity is that OWL is an expressive representation language based on formal logic. Thus, it is able to model complex knowledge. In addition, OWL allows Semantic Web developers to perform logical reasoning on the knowledge, hence enabling access to implicitly modelled knowledge. (Hitzler et al. 2010, 111.) It should be noted that OWL and RDFS have the same purpose, i.e. to define classes, properties and their relationships. However, unlike the RDFS, OWL allows for the capability to express much more complex and richer relationships. In other words, tools and software agents with greatly enhanced reasoning ability can be constructed with OWL. (Yu 2007, 95.)

In 2009, a newer version of the Web Ontology Language, OWL 2, was introduced as an improvement to OWL. OWL 2 adds several new features to the first version of OWL, such as increased expressive power for properties, extended support for data types, simple meta-modelling capabilities, extended annotation capabilities, and keys. OWL 2 also defines several profiles, i.e. sub-languages that may better meet certain performance requirements or may be easier to implement. (Golbreich & Wallace 2012.) The structure of OWL 2, adopted from the W3C OWL Working Group (2012), is seen in Figure 5.

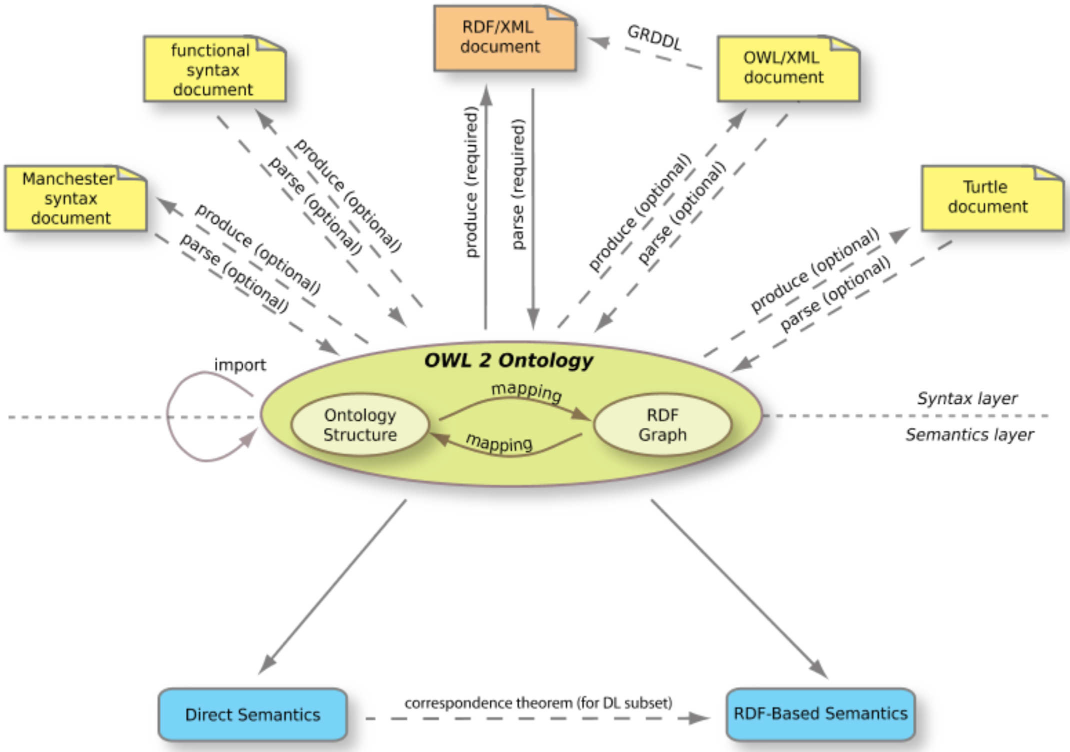

Figure 5. The Structure of OWL 2 (W3C OWL Working Group, 2012)

Extensible Markup Language (XML) “describes a class of data objects, called XML documents, and partially describes the behaviour of computer programs which process them.” These XML documents are composed of storage units called entities, which contain either parsed or unparsed data. Parsed data is composed of characters, some of which form either character data or markup. The purpose of the markup is to encode a description of the document’s storage layout and logical structure. (Bray & Paoli & Sperberg-McQueen & Maler & Yergeau 2008.) While Hyper Text Markup Language (HTML) was designed to display data, XML was designed to transport and store data. Along with its related standards such as Namespaces and Schemas, XML forms a means to structure data on the Web. However, they do not convey the meaning of the data. (Matthews 2005, 4.)

The Unicode is the standard for representing computer characters, while URI is the standard for identifying and locating resources. URI also provides a baseline for representing characters used in most of the languages in the world. (Matthews 2005, 4.)

Resources identified and denoted by URIs can be any object within the context of a specified application which maintains a clear identity such as books, cities and humans, as well as their relationships to one another (Hitzler et al. 2009, 9).

Logic and Proof is a reasoning system for software agents that is automatic and is added to the ontology structure to make new inferences. Its purpose is such that a software agent using such system will be able to infer the satisfaction level of the requirements of a resource. Lastly, Trust, which is the final layer of the Semantic Web architecture, addresses the issues of trust that the Semantic Web is able to support. However, this layer has not been standardised yet. (Matthews 2005, 5.)

4.2 Semantic Web Services and Semantic Interoperability in the eHealth Domain

As was previously discussed, the technologies of the Semantic Web seek to implement a Web that is machine-interpretable, i.e. a Web where computer algorithms are able to process and reason with information that only humans are currently able to interpret. Meanwhile, Web Service technologies are tending towards a system in which organisations are able to make some of their resourcefulness accessible from the Internet. To achieve this, some computational capability is “wrapped” with a Web Service interface, therefore enabling other organisations to locate and interact with it. The vision of SWS is to combine these two technologies i.e. Web Services and Semantic Web. Through this combination, automatic and dynamic interaction between software systems will become possible. Since Web Service technology already allows an interface to be described in a standard way, but does not specify what the software system does in machine-interpretable form, this lack of specification can be fixed by using Semantic Web technology, thus giving rise to Semantic Web Services (Preist 2010, 159.)

Within the context of healthcare, adopting and applying SWS technologies will greatly benefit eHealth systems. The peculiar nature of healthcare, being its life-and-death connotation, suggests that adoption of new processes involving any kind of technology must meet the highest standards of efficacy and precision (Valle & Cerizza & Celino & Dogac & Laleci & Kabak & Okcan & Gulderen & Namli & Bicer 2010, 381). Thus, the issue of semantic interoperability in healthcare takes on great importance in applying SWS technologies in the eHealth domain.

The CEN/ISSS eHealth Standardization Focus Group (as cited by Valle et al. 2010, 384) defines interoperability as “a state which exists between two application entities when, with regard to a specific task, one application entity can accept data from the other and perform that task in an appropriate and satisfactory manner without the need for extra operator intervention.” Within the context of the Semantic Web, interoperability refers to the transmission of data between machines such that there is a shared and distinct meaning between them. This process is termed Semantic Interoperability. In order to actualise semantic interoperability, the application entities that are communicating with each other need to have a common ontology. This will ensure that the information being transmitted is correctly interpreted and understood by all the parties involved. From a theoretical perspective, applying semantic interoperability to the healthcare field seems ideal given that it is such a distributed field. However, Valle et al. (2010, 384) point out that defining an application protocol for interoperability in healthcare is currently a major challenge for eHealth. Nevertheless, there have been some standardisation efforts to aid in curtailing this challenge. Some of the standards include HL7 (Health Level Seven), GEHR (Good European Health Record)/openEHR, CEN/TC 251 (CEN Technical Committee 251), and IHE (Integrating the Healthcare Enterprise). (Valle et al. 2010, 384.)

Although there are similarities in concept and functionalities of these standardisation processes, such as the uniform use of ontologies, some challenges still persist. One challenge is that it is difficult to deal with systems that commit to different ontologies. Another challenge is that there are no extensive models that can automate the usage of services such mediation at data and process levels. (Valle et al. 2010, 385-386.) Despite these challenges, some projects have been able to make considerable progress. Chapter 6 details some specific instances where semantic interoperability has been attempted and successfully implemented in eHealth.

4.3 Impact of the Semantic Web in healthcare

Eysenbach (2003, 17) points out that health information is one of the most sought after on the web, as it constitutes about 4.5% of all queries in search engines. Therefore, the Semantic Web provides an enhanced platform to search for health products and services, as well as the attributes and reputation of these health products and services.

As was previously discussed, since health matters are critical, the issues of accessibility and quality of health information on the Web are paramount. Eysenbach (2003, 4) argues that the Semantic Web has the potential to have a profound influence on how people will interact with the web and obtain information. The infusion of Semantic Web technologies into search engines will enable users to conduct accurate and relevant searches on the Web for health-related information. Furthermore, search results will be better ranked not only by relevance, but also by quality. Ranking by quality reveals the extent to which a health resource is trusted within the healthcare community. Additionally, search engines will become more intelligent and will be able to provide accurate answers to direct questions that a user queries. Apart from such enhanced search engines, new types of software agents will be able to conduct searches independently. These software agents will also have the ability to analyse and combine fragments of knowledge published by different sources, while also performing some autonomous reasoning on these fragmented knowledge. (Eysenbach 2003, 4–5.)

The Semantic Web offers some opportunities within the healthcare community. One opportunity is the translation of knowledge, in which patients have increased possibilities to access information that is actually relevant to their health. In addition, the Semantic Web enables health consumers and health personnel to better identify quality and trusted health information on the Web. More so, the integration of metadata in the RDF format enables information to be accessible to different target audiences in unique ways. However, the Semantic Web poses some challenges as well to healthcare. Firstly, integrating the Semantic Web into healthcare leads to an increased danger of disconnect between patients and health personnel, since accurate information is readily accessible on the Web. Secondly, there are privacy concerns to be addressed, as there could be increased possibilities to collect information about individuals on the Web. (Eysenbach 2003, 7.)

Nonetheless, if the prospects of the Semantic Web are considered for the field of knowledge management and knowledge translation in consumer health informatics, its impact on healthcare could prove beneficial eventually. Perhaps, the most significant application of the Semantic Web in healthcare is trust management. With the Semantic Web, consumers are better able to identify high quality trustworthy health resources on the web. (Eysenbach 2003, 1.)

The ultimate aim of the Semantic Web is to enhance knowledge. Eysenbach (2003, 8) states that "information has to be put into context, the concepts have to be explained and defined, and their relationships to other concepts and to personal information have to be made explicit." The possibility to guide consumers to trustworthy health information using Semantic Web technologies is perhaps the most significant impact of the Semantic Web in healthcare (Eysenbach 2003, 8). Indeed, the Semantic Web has the potential to impact positively on consumer health informatics.

5 A SEMANTICALLY-ENABLED EHEALTH DOMAIN

Having elaborated on the concepts of eHealth and the Semantic Web in Chapters 3 and 4 respectively, this chapter discusses an eHealth environment that is semantically-enabled. The various attempts and some suggestions to integrate eHealth with SWS technologies are discussed. The technological advancements in the areas of medical ontologies, health 3.0, medical search engines, and clinical decision support systems are also analysed to illustrate how SWS technologies fit into these systems.

5.1 Medical ontologies

Ontology is one technique that is used to enable semantic interoperability of health-related information across various domains. Stenzhorn, Schulz, Boekern and Smith (2008, 3769) find that logically defined and precise formalisms are normally relied upon in order to build ontologies. This makes it possible to describe concepts without the intervention of human interpretations. According to Pisanelli (2007), ontology refers to a declarative model of a domain that defines and represents the concepts existing in that domain, as well as their attributes and the relationships between them. Ontologies are usually expressed as a knowledge base that becomes accessible to applications so that knowledge of a particular domain is shared and applied. Within the eHealth community, ontology refers to a formal description of a health-related domain.

The use of ontologies in the medical field is mainly focused on the representation of medical terminologies. Specialised languages and lexicons are developed by health personnel, which enable them to store and communicate general health-related knowledge and patient-related information efficiently. One major benefit of ontologies, as was discussed previously, is their applicability in constructing highly efficient semantic interoperability within healthcare systems. Ontologies are also able to support the need of the healthcare process to transmit, re-use and share patient data. Furthermore, one very significant benefit of ontologies in healthcare systems is their ability to support the vital integration of knowledge and data. However, there remains some scepticism about the impact that ontologies may have on the design and maintenance of healthcare information systems in practice. (Pisanelli 2007.)

From a practical perspective, ontologies have become one of the most prominent and important resources in ongoing biomedical informatics research. Due to their major objective of advancing semantic interoperability, ontologies are applied in large clinical research projects that are interconnected. It is believed that ontologies are able to offer stable and language independent vocabulary that can aid in standardising and explaining the actual meaning of domain terms. (Stenzhorn et al. 2008, 3778.) Current research works on medical ontologies include projects such as ACGT (Advancing Clinico-Genomic Trials on Cancer) and @neurIST (Integrated Biomedical Informatics for the Management of Cerebral Aneurysms). The ACGT project focuses on nephroblastoma and breast cancer, while the @neurIST project focuses on estimating the risk levels of intracranial aneurysms and subarachnoid haemorrhage. As these projects present different environments, designing ontologies that integrate all associated elements and data poses a challenge. Therefore, these projects entail the development of customised ontologies to create a common basis for applications to semantically mediate between the disparate software components within the projects. The overall aim of ACGT and @neurIST is to create integrated ICT infrastructures by implementing common software platforms. The specific objective is to improve the management of diseases by creating an efficient environment that enables existing knowledge to be combined with newly generated data. The ultimate goal of the developed platforms is to integrate the highly fragmented and heterogeneous data from different sources and disciplines within the projects. (Stenzhorn et al. 2008, 3772.)

Some other works on medical ontologies include The Foundational Model of Anatomy (FMA), which is a domain ontology that represents a coherent body of explicit declarative knowledge about human anatomy. The Gene Ontology Consortium is another research work which aims to produce a controlled vocabulary that can be applied to all organisms. In addition, the Medical Ontology Research program aims to develop an efficient medical ontology to enable various knowledge processing applications to communicate with one another. Lastly, the Language and Information Engineering Lab in Germany focuses on automatic text analysis in order to service various applications such as information extraction, text mining, cross-language document retrieval, and text summarisation. Most of these applications are embedded in the biomedical domain. (Pisanelli 2007.)

Furthermore, in addition to the existing medical ontologies currently in use, some other ontologies exist which can be potentially applied in healthcare systems. Knowledge Interchange Format (KIF) is one of those ontologies. It is a computer-oriented language for the interchange of knowledge among various programs. It has declarative semantics and is logically comprehensive. KIF also helps define objects, functions, and relations. Ontolingua is another ontology which might prove useful in the eHealth domain. It is a reusable ontology which makes use of the Web to enable wide access and provide users with the ability to publish, browse, create, and edit ontologies stored on an ontology server. Lastly, there is the General Ontology Language (GOL) which is a conceptual modelling language. (Pisanelli 2007.)

5.2 Health 3.0

Since the surge in the use of the Internet in recent times, there have been efforts to use the semantic web to enhance healthcare services. This has led to the emergence of the Health 3.0 concept. Health 3.0 is essentially a concept that is used to imply the integration of Web 3.0 and eHealth. Web 3.0 is defined as “a supposed third generation of Internet-based services – such as those using Semantic Web, micro-formats, natural language search, data-mining, machine learning, recommendation agents, and artificial intelligence technologies – that emphasise machine-facilitated understanding of information in order to provide a more productive and intuitive user experience” (Spivacks 2006 as cited by Cheung & Yip & Townsend & Scotch 2008, 8). Health 3.0, more explicitly, therefore refers to a health-related component of the Web 3.0 concept in which the user experience is optimised by the personalisation of the user's interface with Web data. As Nash (2008) remarks, Health 3.0 will present an opportunity for individuals to be able to better retrieve and possibly contribute to personalised health-related resources, particularly within a social networking context. Some potential benefits of Health 3.0 include the establishment of supportive virtual centres where individuals can help one another with various health-related issues. More so, therapeutic healing could be improved as nurses and doctors are better able to reach out to their patients through personalised social networking platforms. From a Semantic Web perspective, enriched access to health-related information on the Internet could also be facilitated by Health 3.0. This will not only aid a greater understanding of health issues, it will equally revolutionise disease management due to proper awareness of disease prevention. (Nash 2008.)

However, there remains the challenge of centralising information and data from multiple sources in order to actualise the objectives of the Health 3.0 concept. Realising this challenge, Cheung et al. (2008, 3) point out that the idea of mashups emerged as a means to address the problem of integrating data access from diverse sources. Cheung et al. (2008, 3) define a mashup as “a Web application that combines multiple third-party services over the Web.” In other words, a mashup uses information from one or more sources and presents it in a unique way to create a new service. According to Cheung et al. (2008, 3), one of the most recent use of data mashup is seen in the Health Care and Life Sciences (HCLS) domain, where the Google Earth application was integrated geographically and visually with different kinds of data, such as public health data to help track the spread of avian influenza around the world. To understand how such integration works, the following scenario is presented by Cheung et al. (2008, 4) in which some data on cancer is geographically incorporated with environmental data using Yahoo! Pipes, Google Maps, and GeoCommons.

The study of the correlation between human diseases such as cancer, and environmental factors often requires that different data sources such as population census, quality of air, and environmental pollution are integrated together. However, an automated integration of these data poses a challenge, as the data are usually produced by different agencies. Thus, mashups offer the possibility to automate the integration of diverse health care data in order to facilitate the environmental health research. To perform this integration, a cancer profile dataset is identified at the agency website. The needed data is then extracted using Yahoo! Pipes. Subsequently, the output is fed to a widget that displays this data on Google Maps. The map is further exported and uploaded to GeoCommons. At the GeoCommons website, users are allowed to annotate their uploaded maps, including those uploaded by other users. (Cheung et al. 2008, 6.) In Figure 6, a GeoCommons interface that shows the cancer profile map superimposed with water pollution map in the United States is shown. The bright colours indicate the polluted areas, while the white dots indicate the cancer profiles. The image is adopted from Cheung et al. (2008, 23).

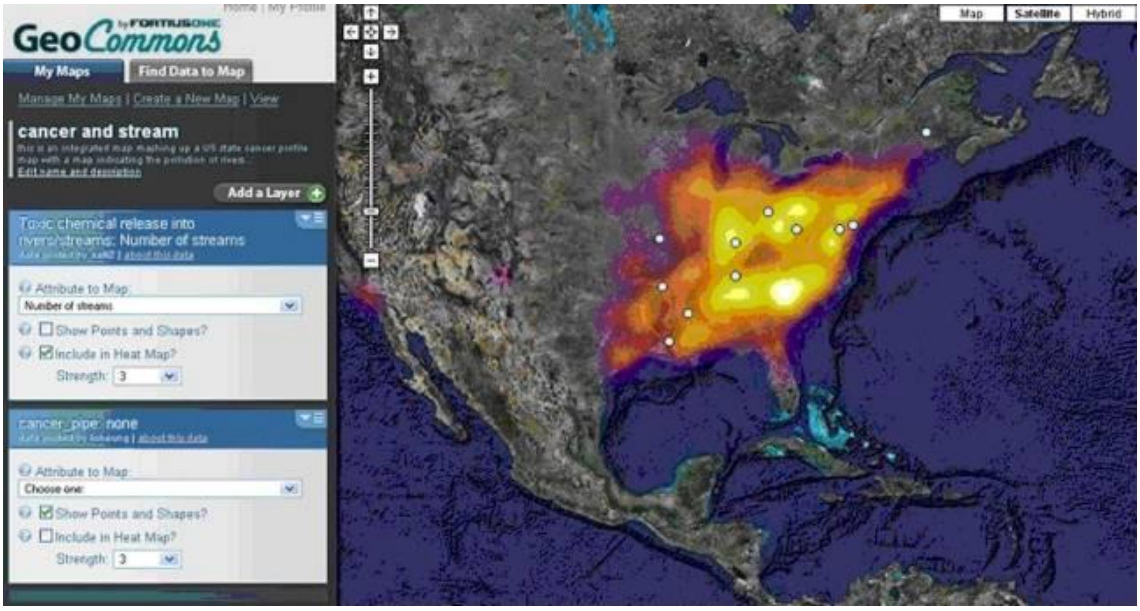

Figure 6. A mashup of the state cancer profile map and water pollution map (Cheung et al. 2008, 23)

Mashup technologies have proven to be beneficial within the HCLS domain to a large extent. However, despite their advantages, they have some limitations as well. As a result, the W3C launched the Health Care and Life Sciences Interest Group (HCLS IG) to develop and support the use of Semantic Web technologies to improve the HCLS domain. Cheung et al. (2008, 8-10) point out that there have been attempts to convert different HCLS data sources into the standard Semantic Web data formats endorsed by W3C, such as RDF and OWL. As in the case of the cancer mashup, a semantic mashup based on locations could be built. For instance, ontology may be defined in which a city, e.g. Tornio, is located in a province, e.g. Lapland, which is located in a region, e.g. Northern Finland. With such ontology, deductions based on location may be carried out when mashing up data. To create improved human-computer interaction capabilities, Semantic Web applications could be built such that semantic mashup of HCLS data will be supported in enhanced user-friendly ways.

5.3 Medical search engines

Medical search engines are used to find answers to medical questions, to search for solutions to health problems, and to obtain information about various health topics. Today, there are many medical search engines that cater to the various needs of users. Some of the most commonly used medical search engines include Healthline, PubMed and OmniMedicalSearch. Although these medical search engines all have the same basic purpose of providing users with health-related search results on the Web, they do so in unique forms. For example, OmniMedicalSearch not only provides users with authoritative health-related search results, it also provides users with health news and images. PubMed is a service of the United States National Library of Medicine that has an extensive medical database with millions of health articles and peer-reviewed journals. Meanwhile, Healthline offers users medically filtered search results that are developed by experienced health professionals.

Search engines today are an integral part of people's online activities and experience, as millions of Internet users frequently search the Web for health information. Consequently, decisions about health and healthcare are affected by these pieces of information obtained. (Leroy 2009.) Leroy (2009) observes that the input method for searching information on search engines has remained the same for several years. This input method, i.e. a single textbox, has subsequently impacted on the search behaviour of users, as they are often compelled to use few and inaccurate keywords. In turn, obtaining the wrong information from improper use of keywords could prove detrimental to the health of the user. As a result, Leroy (2009) seeks to persuade users to make precise search engine queries by proposing and evaluating a change of search engine user interface. This user interface change will act as a buttress to the more efficient algorithms which have always been the focus of most search engine advancements.

In order to implement the kind of new search engine system described above, the concept of affordance should be understood. Leroy (2009) defines affordance as "a property of an entity or object that allows interaction with that object in a specific way."

Daily human interactions with the physical environment have always been dependent on the manipulation of affordances. For example, the type or design of a chair will influence whether one should recline or sit upright. A chair without a back and arm rest will certainly force an individual to sit upright. In the same manner, such affordances that compel users to behave in fixed ways abound on the Internet. Examples include "pushing" buttons and "checking out" when purchasing items. Likewise, search engines that provide only a single search box, in which users input strings of text, ensure that no other type of search input is possible. As a possible improvement to this traditional search engine system, Leroy (2009) proposes the use of diagram queries. Rather than a single search box, a two-dimensional interface can be used. This interface will consist of multiple search boxes, links between those searches boxes, and the capability to type in the search box. There will also be the possibility to add additional boxes. In Figure 7 which is adopted from Leroy (2009), the question of "What medication treats depression in teenagers?" is asked as a sample query. As can be seen in the diagram, each box represents a search term, while the label and direction of the arrows depict how the search terms are related to one another.

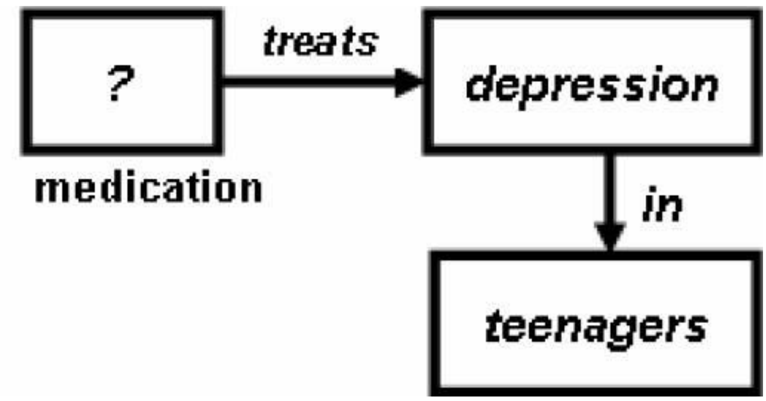

Figure 7. A sample query diagram (Leroy 2009)

Apart from the more structure queries, the use of two affordances is also seen in the interface illustrated in Figure 7. One is the question mark which tends to compel a user to be specific rather than descriptive, and the other is the option to include metadata, which in this case is "medication", for more accurate results. The diagram queries will not only be intuitive, but will act as a more powerful search interface than what is currently obtainable. However, to implement this sort of search system, different processing techniques and data structures will be required. (Leroy 2009.) In this sense, the Semantic Web technologies could prove useful.

Within the context of the Semantic Web, a medical search engine is primarily concerned with adding intelligence to the existing system. By applying Semantic Web technologies, an ideal medical search engine which has added effectiveness of catering to the specific health needs of users could be developed. As current search engines are only able to analyse keywords and retrieve documents based on those keywords, applying the technologies of the Semantic Web will enable search engines to further understand and respond to search queries in more specific ways (Dietze & Schroeder 2009). For example, a gynaecologist queries a search engine with the following questions: (a) which disease can be linked to placenta infection of an expectant mother? (b) What are the negative implications of In vitro fertilisation (IVF) treatments? The answers to these questions are readily and widely available on the Web. However, the traditional keyword-based search is unable to provide straightforward answers because the Web does not understand the questions. It only analyses the keywords based on the individual strings of text contained in the search and presents documents that match those keywords. Therefore, by applying the standards of the Semantic Web such as OWL, XML, RDF and RDF(S) to enhance machine-readability and knowledge processing, search queries could be improved significantly (Dietze & Schroeder 2009). Recognising this possibility, Dietze and Schroeder (2009) introduce an approach which integrates the traditional keyword-based search with text-mining and ontologies. The purpose of this integration is to manage large sets of results and enable the answering of search queries. The result of this approach has led to the development of GoWeb.

GoWeb is an internet search engine that is based on the use of ontologies. GoWeb filters long lists of search results based on the categories provided by the GeneOntology (GO) and the Medical Subject Headings (MeSH). GO and MeSH are semantic health standards. GoWeb offers efficient search and result set filtering mechanism, and semi-automatic question answering with the ontological background knowledge. When a user submits a query on the GoWeb website, the server pre-processes the query and the search request is sent to the search service. Subsequently, initial results are returned in the search service. These results are then annotated, while the concepts and keywords are highlighted, rendered and sent to the user. Once the initial results are processed, the server begins to fetch the remaining results. The server does this for up to 1000 results, which are all eventually annotated as well. Figure 8, adopted from Dietze and Schroeder (2009), illustrates the GoWeb workflow. It shows the main components of the system and the interactions between the external services.

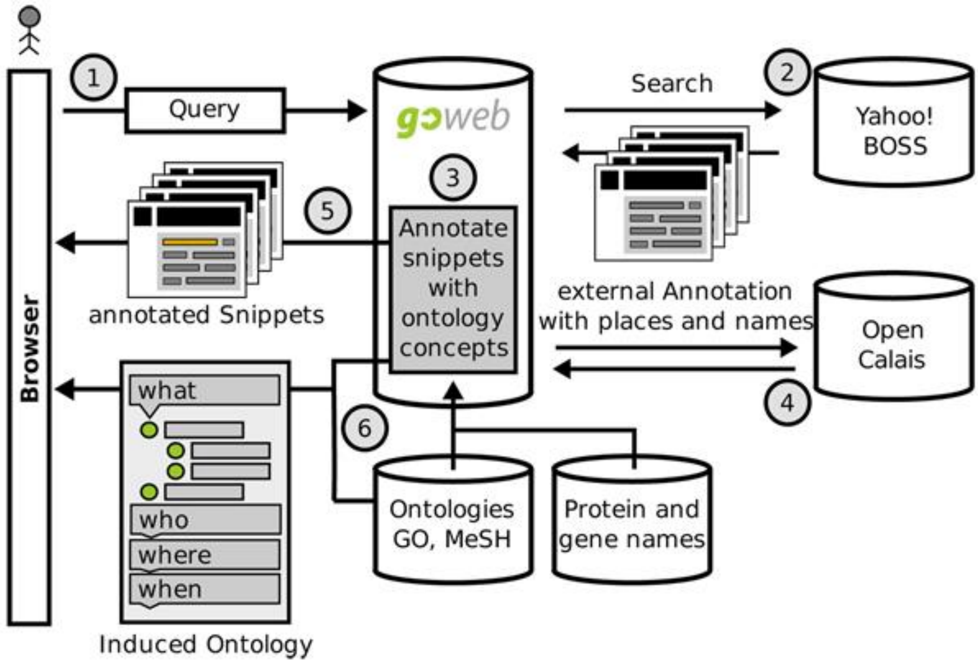

Figure 8. A general GoWeb workflow (Dietze and Schroeder 2009)

Luo (2008, 1201) goes a step further than Leroy (2009) and Dietze and Schroeder (2009) in these integration attempts by introducing a unique concept called iMed, an intelligent medical search engine. iMed was specifically built to address the issue of uncertainty that frequently plagues users in their search for health-related information. iMed uses medical knowledge and an interactive questionnaire to help users form their queries. The search results of these queries are combined and returned to the user in a traditional sequential order. The uniqueness of this search engine is its ability to automatically offer users what they want, rather than wait until they ask explicitly. The search results are structured into a multi-layered hierarchy with clearly marked medical meanings. This structure is to ensure that users are able to efficiently navigate among all the search results and obtain desired information promptly. (Luo 2008, 1201.)

Figure 9 demonstrates the diagnostic decision tree of the iMed system for a chest pain symptom. The diagram is adopted from Luo (2008, 1201).

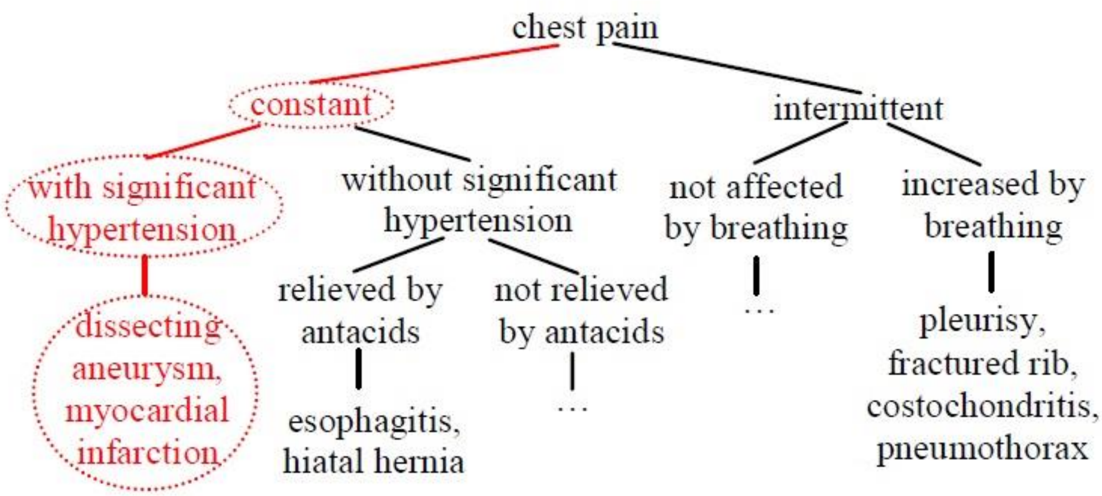

Figure 9. iMed diagnostic decision tree (Luo 2008, 1201)

When the user inputs a query, e.g. chest pain, iMed asks: “Is the pain constant or intermittent?” When the user selects an option, e.g. constant, iMed asks the following question: “With significant hypertension or without?” By default, depending on the user selection, the appropriate queries will be formed by iMed until a solution is found. (Luo 2008, 1201.) This kind of system ensures that the user inputs more specific queries in order to be able to receive accurate answers.

In addition to all the aforementioned integration examples, there are several other medical search engine innovations that cannot be covered in this research due to time and space constraints. Among them are MedicoPort and MedSearch. MedicoPort is a next generation domain search engine designed for users with no medical expertise. It is enhanced with the domain knowledge obtained from Unified Medical Language System (UMLS) to increase the effectiveness of searches. The strength of the system is based on its ability to understand the semantics of web pages and the user queries. MedicoPort aims to generate maximum output with semantic value using minimum input from the user, such that the retrieved answers from Web are relevant to the user request. MedSearch is very similar to MediPort except that it does not specify the intended users of the system. There have also been attempts to build search engines to meet the needs of Public Health Information so that documents are indexed formally to make them easier to locate on the Web. Furthermore, Zheng, Mei and Hanauer (2011) propose a full-text search engine for Electronic Health Records (eHR) in order to aid collaborative search. Such collaboration will enable more efficient and quality information retrieval in healthcare.

5.4 Clinical Decision Support Systems (CDSSs)

Due to the error-prone traditional methods currently used by physicians to diagnose patients, the need for a system that produces optimal diagnostic results arose. Hence, the idea of a Clinical Decision Support System (hereinafter CDSS) was conceived. CDSS is defined as "software that is designed to be a direct aid to clinical decision-making, in which the characteristics of an individual patient are matched to a computerised clinical knowledge base and patient-specific assessments or recommendations are then presented to the clinician or the patient for a decision" (Sim & Gorman & Greenes & Haynes & Kaplan & Lehmann & Tang 2001, 528).

CDSS is built to enable the integration of a medical knowledge base, patient data and an inference engine in order to aid better decision making on health matters. This system assists health personnel in diagnosing patients, determining ideal treatment methods, and offering suggestions on substitute procedures. Basically, CDSS solves the problem of how to manoeuvre the network of clinical actions and decisions in an optimal way. A CDSS works by searching for similarities between clinical methods applied on a current patient with the methods applied on previous patients with similar ailments. The system does this with the help of an inbuilt software application that analyses the stored electronic health records of similar patients in order to possibly administer the same diagnosis or therapy for the current patient. The CDSS receives the health records of the patient, and compares with the treatments of similar former patients. Based on these records, the system then generates classifiers to be paired with potential clinical measures in the future. A quality value indicating the potential success of a particular treatment on the current patient, based on its implementation on a similar patient, is calculated by the CDSS as well. Subsequently, the system indicates the clinical measures that are associated with the highest quality value, which are displayed on the

CDSS graphical user interface. In essence, the basic idea of the similarity search of the CDSS is to determine if the same results can be achieved when the methods are replicated on another patient. (Schmidt & Schaepe & Heydler & Rinecker & Binnig 2012, 1.) In Figure 10 below adopted from Schmidt et al. (2012), the structure of the Clinical Decision Support System is illustrated.

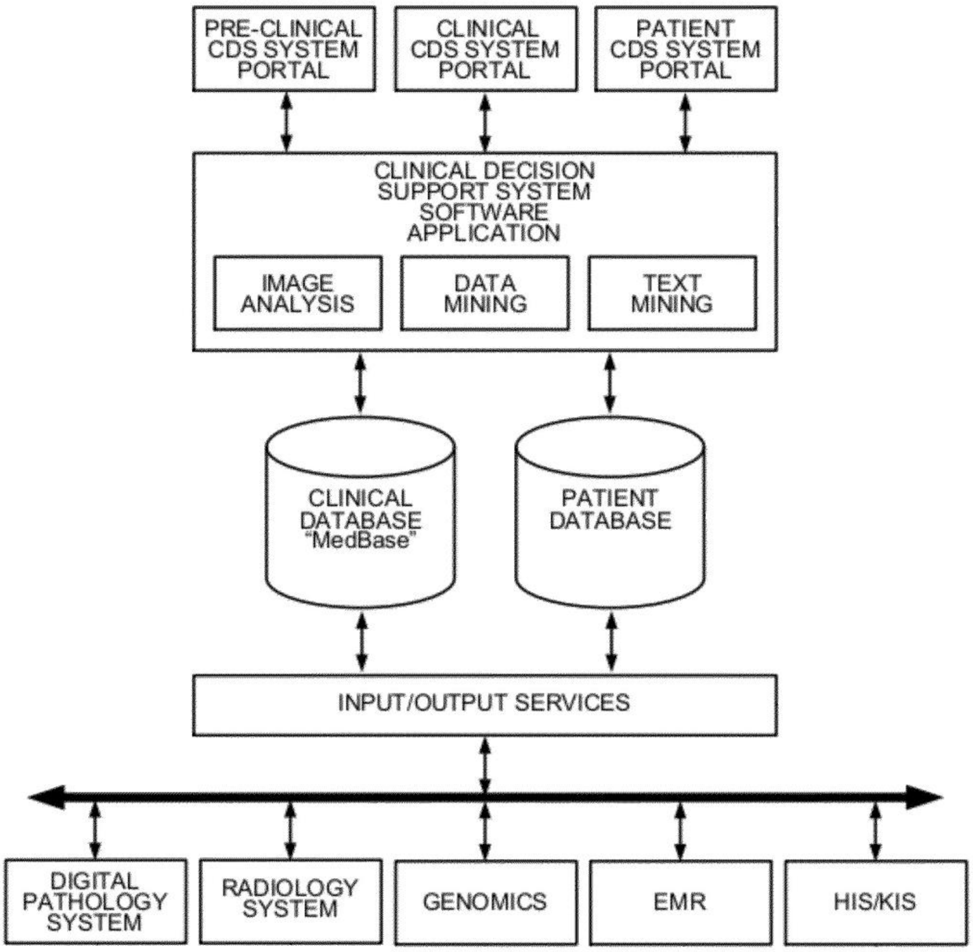

Figure 10. A Clinical Decision Support System (Schmidt et al. 2012)

As can be seen in Figure 10 above, the CDSS has a layered architecture of data and software. There is a CDSS software application that runs on a server processor. The large numbers of patient records are stored in a clinical database, represented as MedBase in the diagram. For each patient in this MedBase, the system stores the sequence of clinical actions steps in the patient database. These steps include measurements, assessments and therapies. The clinical action steps also include the previous decisions and corresponding costs. The input/output services enable the final layer of the data network to access data sources. These data sources are the digital pathology system, radiology system, genomics, electronic medical records, and Hospital Information Services (HIS). (Schmidt et al. 2012, 3-5.)

According to Sim et al. (2001, 529), one crucial measure to develop a more efficient CDSS is to create improved and beneficial data that is recent, unrestricted, and machine-interpretable. Such measure suggests integrating SWS technologies with CDSSs. One way to apply SWS technologies in a CDSS is to use ontologies. Bodenreider (2008, 73) points out that ontologies enhance CDSSs in two fundamental ways. Firstly, ontologies enable biomedical entities to possess a standardised vocabulary, as they support the integration of knowledge and data. For example, a system for pain medications ought to be capable of parsing different types of pain into standard codes, and integrating the pain coding systems with the knowledge base of suitable medications. Secondly, ontologies provide a "computable domain knowledge" that can be used to aid decision support. For example, in a system for pain medications, an efficient representation can be provided for chest pain medications if the system can access a classification of pain, rather than a direct access to specific pains. Bodenreider (2008, 73) further cites an example of the application of the FMA ontology in anatomical practice to illustrate the relevance of ontologies, not only in clinical decision support, but in application reasoning as well. The cited example highlights the use of the FMA, through semantic reasoners, to predict the repercussion of penetrating injuries. The FMA identifies the proximity of the injury path to the vital organs of the body. As a result, the decision on specific clinical actions to perform is greatly improved.

Furthermore, there have been substantial attempts to execute Clinical Practice Guidelines (hereinafter CPG) through CDSS. A CPG aims to guide the decision-making of health personnel regarding healthcare practices. According to Hussain, Abidi S. and Abidi S.S. (2007, 451), "CDSS can offer the functionality to (a) execute the CPG at the point of care; (b) guide healthcare practitioners to make evidence based decisions, actions and recommendations; (c) standardise the delivery of care at a particular healthcare setting; and (d) collect all necessary and relevant patient data." In addition, a

Semantic Web framework can be applied to develop a CPG-enabled CDSS. The application of semantic formalisms, such as CPG ontology, and the implementation of semantic interoperability between various knowledge resources can equally enhance CDSS as a whole. (Hussain et al. 2007, 451-452.)

The model represented in Figure 11 below further demonstrates how CDSS can benefit from Semantic Web integration. Khan and Hederman (2012, 1) propose the idea of a universal CDSS that is powered entirely by Semantic Web Services. This vision is based on a Web which will offer various clinical decision support services to users over the Internet. Such services will include diagnosis, drug prescriptions and an avenue to converse with health personnel.

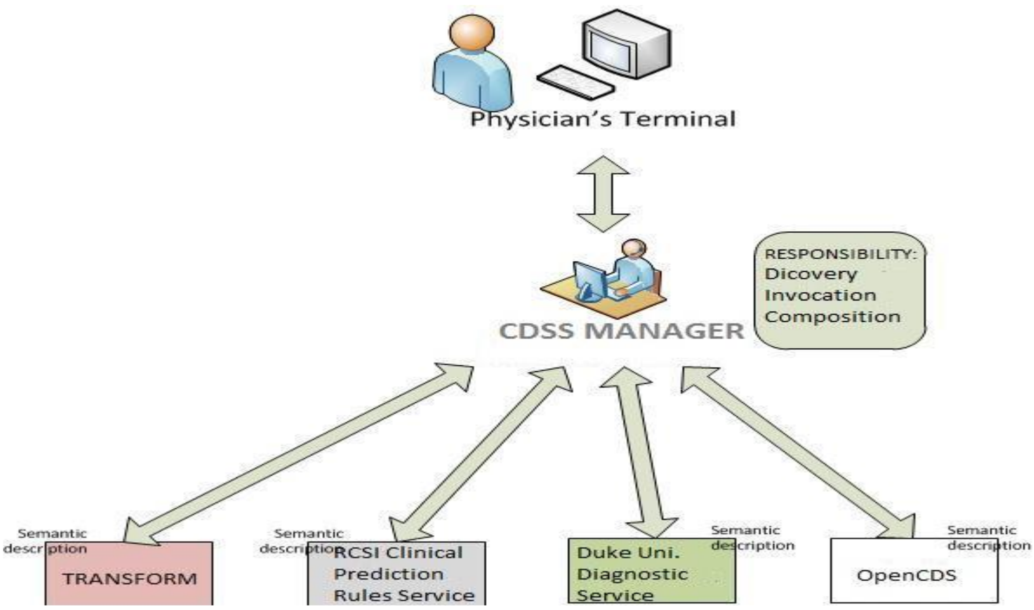

Figure 11. A semantically-enabled CDSS (Khan & Hederman 2012, 2)

As can be seen in Figure 11, a CDSS manager will be present at each healthcare institute to regularly identify, examine and revise clinical decision support components. The system will also integrate these components from various providers such as OpenCDS and TRANSFORM. OpenCDS, for instance, facilitates the appendance of keywords to clinical knowledge. Therefore, in order to implement such integration efficiently, rich semantic descriptions will be applied. Applicable ontologies will be embedded to construct a semantically-powered CDSS. The role of the CDSS manager in this system is essentially to automate the processes between the physician and the services to be provided. This automation will ensure minimal human intervention, thus enabling higher healthcare efficiency. A universal CDSS will ensure more effective retrieval of clinical decision support services through metadata, as opposed to keywords. A universal CDSS will also enable relevant services to be analysed by the CDSS manager. These services can then be applied based on the healthcare needs. (Khan & Hederman 2012, 2-6.)

The Internet in particular has offered vast possibilities for patients to obtain health information. However, it has also exposed patients to the risks of misinformation and misinterpretation of search results. Consequently, patients are now less dependent on health personnel for information, but still entrust them with the assessment and approval of health decisions. Interactive tools that enable patients to examine relevant information can be implemented through CDSS. Considering the increased involvement of patients in the decision-making process, a CDSS can ensure that decision-making becomes more collaborative. By providing both patients and health personnel with appropriate information, a CDSS may result in healthcare decisions that are in consonance with approved recommendations. More so, healthcare decisions that are better suited to individual patients can be achieved through CDSS, thus resulting in enhanced healthcare. (Sim et al. 2001, 529-530.)

6 HEALTH-RELATED APPLICATIONS OF THE SEMANTIC WEB

Rovan, Jagušt and Baranović (2011, 245) define Semantic Web application as “a Web application that depends on the Semantic Web standards for its successful execution.” In essence, a Web application can only be classified as semantic when Semantic Web technologies are applied in at least one of its functional components. As was previously discussed in the Introduction chapter, context is very important when one conducts health-related searches on the Internet (Trzebucki 2008). More so, applications that apply the related technologies of the Semantic Web, i.e. metadata, ontologies and knowledge representations are needed to help improve the quality and efficacy of healthcare access on the Web. As a result, there have been ongoing research works in an attempt to integrate the Semantic Web with eHealth, particularly with a focus on semantic interoperability. Clearly, the need for improved, accurate and in-depth healthcare and health information is being recognised. Some successfully implemented health-related applications of the Semantic Web are discussed in this chapter.

6.1 ARTEMIS

ARTEMIS is a Semantic Web Service-based peer-to-peer (hereinafter P2P) infrastructure for the interoperability of Medical Information Systems. In other words, it develops a SWS-based interoperability framework for the healthcare domain. The functionality of ARTEMIS is made possible by the extensions it provides to P2P architectures in order to enable discovery of Web services based on their semantic descriptions. One problem that has afflicted healthcare informatics over the years is the inability to share patient records across enterprises. Although there have been several standardisation efforts to digitally represent clinical data which aim to structure and markup clinical content for the purpose of exchange, the presence of more than one healthcare standard has made it difficult to achieve interoperability. As a result, ARTEMIS was developed to address this difficulty. The purpose of ARTEMIS is to provide the healthcare industry with an ideal platform to exchange meaningful clinical information among healthcare institutes through semantic intervention. ARTEMIS provides an interoperability platform where organisations keep their proprietary systems, but expose the functionality through Web services. To achieve interoperability, an ontology-based description of these data exchange is proposed within the scope of the ARTEMIS infrastructure. The result of this proposition is that ARTEMIS enables medical practitioners to access patient records securely and seamlessly through a low-cost P2P infrastructure, regardless of where the patients or their records are situated. (European Commission 2007, 28-29.) Figure 12, adopted from Valle et al. (2010, 396), depicts the P2P architecture of ARTEMIS. In the diagram, the healthcare institutes are represented as peers. Each peer is able to communicate with the rest of the network through the super peers, also known as mediators. (Valle et al. 2010, 395.)

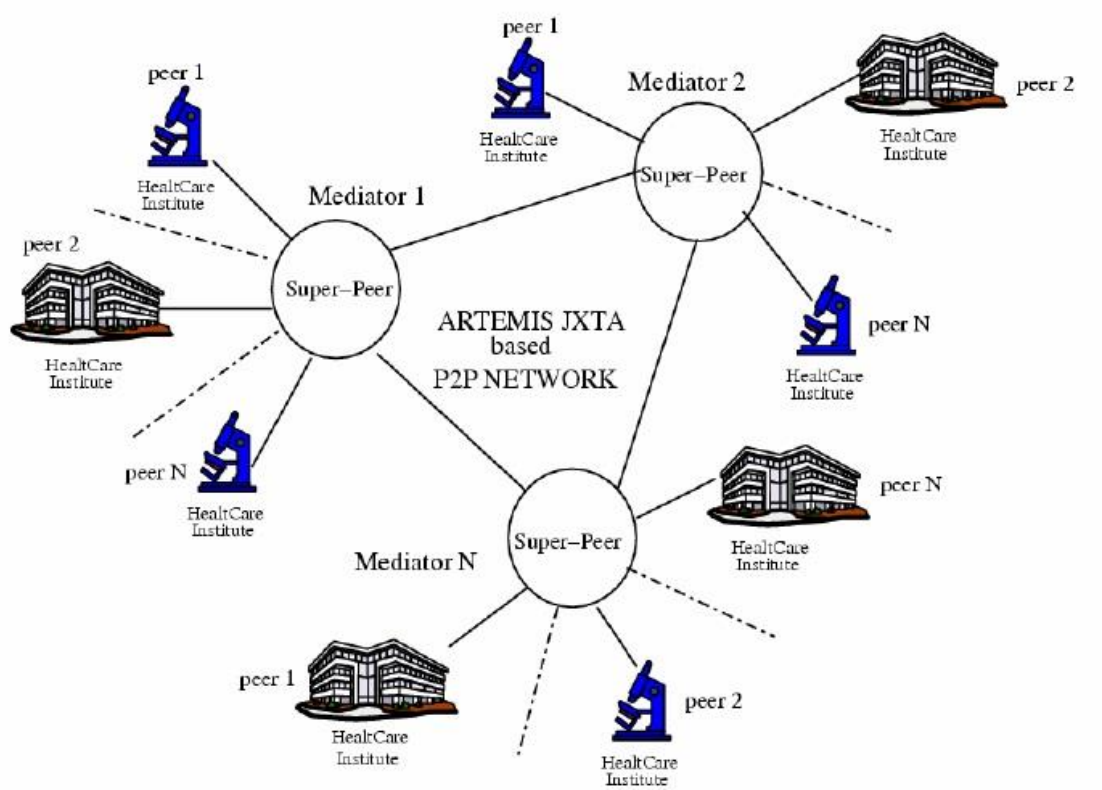

Figure 12. ARTEMIS P2P structure (Valle et al. 2010, 396)

To further clarify how ARTEMIS works, a practical scenario is presented briefly. A patient is admitted to a nearby hospital from an ambulance after an accident. On arrival at the hospital, the hospital admission service automatically searches for relevant healthcare records of the patient from the ARTEMIS P2P network. The patient information is subsequently presented to the doctor who then proceeds to work on the patient. (European Commission 2007, 28.) A sample graphical user interface of an ARTEMIS peer is shown in Figure 13 below. The medical services provided by the hospital, based on their functionalities, are seen in the P2P network. The image is adopted from Valle et al. (2010, 397).

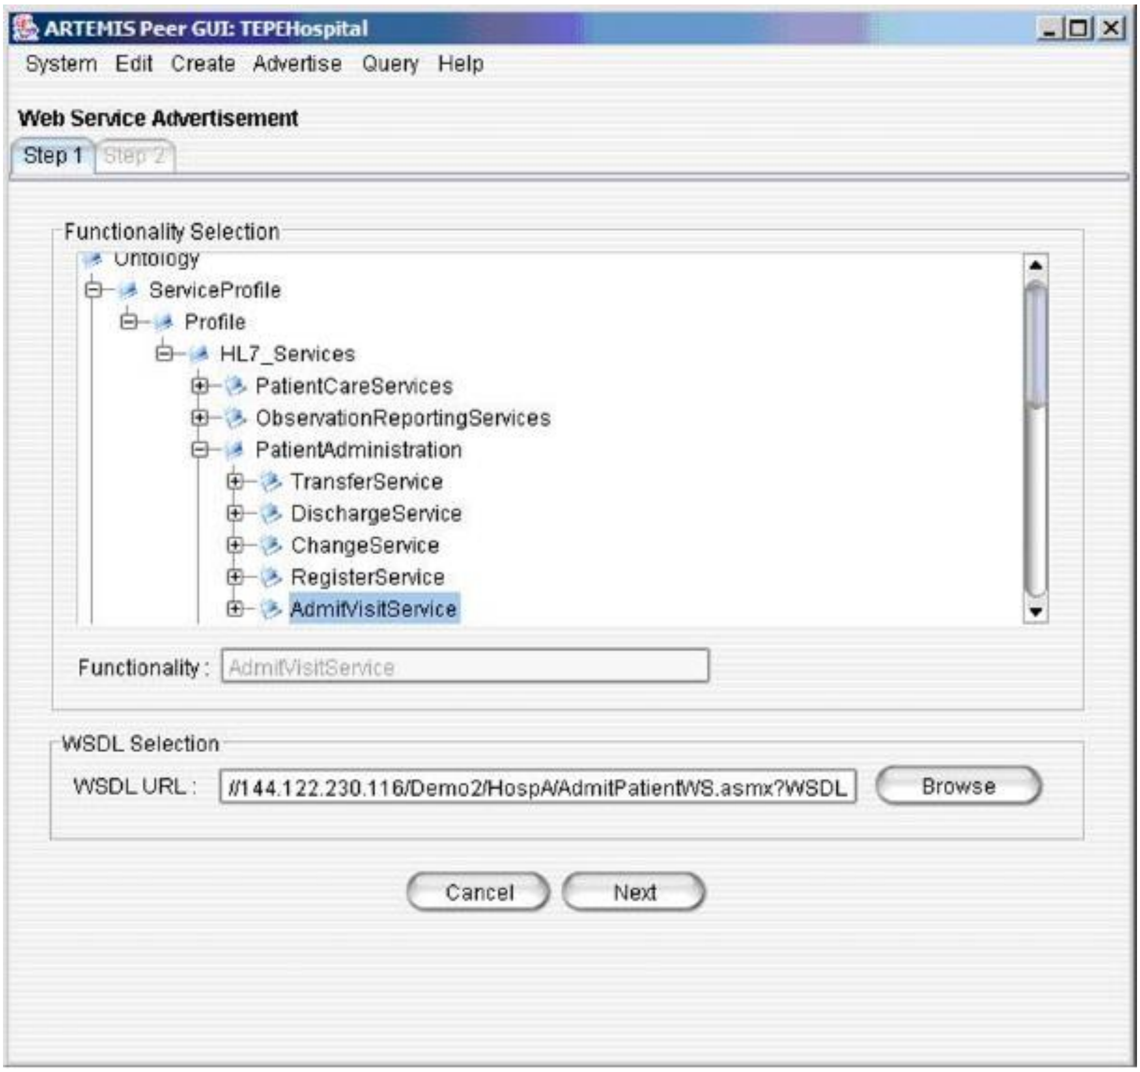

Figure 13. ARTEMIS peer interface (Valle et al. 2010, 397)

In this scenario described above, different hospital information systems with different messaging and coding standards are used to demonstrate the semantic-based interoperability platform. (European Commission 2007, 28.)

6.2 RIDE

RIDE is another health-related application of the Semantic Web. It is a roadmap for interoperability of eHealth systems with a special emphasis on semantic interoperability. RIDE was developed to address the perceived unrealistic expectations of having a single universally accepted clinical data model that will be adhered to all over Europe. The aim of RIDE is to lay a roadmap by coordinating various clinical efforts across the European continent. (European Commission 2007, 100.)

One scenario that best describes the applicability of RIDE is as follows. A family doctor in Finland, Oluwatosin Daniel wishes to refer a patient named Michael to a dental specialist named Elizabeth Priscilla in Eko Hospital, Nigeria. The referral note of Oluwatosin Daniel should be available to Elizabeth Priscilla to continue the care process. For this process to go smoothly within the RIDE system, the Patient Identifiers used by the document source and document consumer should be matched; the communication protocol used by these parties should be fixed; and the interoperability of the messaging and eHR standards used by the parties should be facilitated. (European Commission 2007, 100-101.) Based on these processes, RIDE enables a seamless continuity of healthcare treatment of the patient in a separate geographical region.

6.3 SemanticHEALTH

SemanticHEALTH is a semantic interoperability deployment and research roadmap which aims to develop a European and global roadmap in eHealth. SemanticHEALTH also focuses on the semantic interoperability issues of eHealth systems and infrastructures. Its aim is to deliver safe and effective healthcare. SemanticHEALTH particularly aims to avoid medical errors by providing adequate clinical documentation on patients. (European Commission 2007, 108.)

The vision of SemanticHEALTH is to transform current paper-based medical records into electronic medical records that are accessible to all necessary providers, and possibly to the patient as well. To achieve this vision, interoperability at the technical, syntactic and semantic level is mandatory. The semantic interoperability especially is vital to the seamless flow of data and consistency in meaning on the medical conditions of patients globally. In turn, this will form the basis on which future global health research, patient care and evaluation of public health management can be effectively implemented. (European Commission 2007, 108-109.)

6.4 SemanticMining

SemanticMining, similar to the other applications described above, also concerns itself with semantic interoperability, as well as data mining in biomedicine. SemanticMining is still in the developmental phase and has not been fully implemented yet. SemanticMining aims to preserve meanings in communication between information systems. Although this preservation of meanings should be characteristic of information systems, it has proven difficult to achieve particularly in the complex healthcare domain. The long term goal of SemanticMining is to develop generic methods and tools to support the critical tasks of the field such as abstraction and indexing of information, data mining, knowledge discovery, knowledge representation, and semantic-based information retrieval in a complex and high-dimensional information space. The challenges faced by healthcare systems concerning quality and cost-effectiveness will be equally addressed by SemanticMining. (European Commission 2007, 110.)

Furthermore, SemanticMining will be used to distribute healthcare services in ways that allow the patient to take an active part in relevant decisions. SemanticMining will also provide evidence-based medicine at all levels within the healthcare system, while equally using information effectively in the delivery of healthcare. One scenario that demonstrates the functionality of this system is the possibility of patients to have access to their own health records over the Internet. However, to make this process effective, the online facilities that will help patients without medical knowledge to access relevant information in the health records is necessary. With semantically well-defined eHR and language technology within SemanticMining, patients will be able to receive their records in a generally understandable form. (European Commission 2007, 110-111.)

6.5 HealthFinland

HealthFinland is a semantic health information publishing system. Due to the inadequacy of reliable, up-to-date and individually relevant health information on the Web, HealthFinland was developed. The main reasons for developing HealthFinland using Semantic Web technologies are to facilitate cost-effective distributed content creation in an interoperable way, to aggregate contents automatically based on semantics, and to provide the end-users with intelligent services. The aim of HealthFinland is to provide Finnish citizens with a comprehensive single access to reliable and up-to-date health information and eHealth services. HealthFinland particularly focuses on improving collaboration between all actors in the field of health promotion and health services. (Suominen & Hyvönen & Viljanen & Hukka 2009, 287-288.)

One problem that has plagued the health industry is the minimal coordination and cooperation between publishers of health information. As a result, published works become duplicated, while further expenses are incurred from publishing the same information more than once on different websites. HealthFinland addresses this problem by means of collaboration. The collaborative production network of HealthFinland ensures that health-related information is published only once by an authoritative body. With the application of Semantic Web technologies, content is annotated locally with semantic metadata based on shared ontologies. This annotation enables the published content to be accessible from different Web portals by other organisations. Consequently, the content maintenance costs of the portals are minimised, as the semantic link maintenance and aggregation of various content are automated. Lastly, the end-user is provided with intelligent services to locate the right information based on semantic relations. (Suominen et al. 2009, 288.)

There are three main components of HealthFinland that work in conjunction with one another in order to actualise the processes described above. A diagrammatic representation of these components, adopted from Suominen et al. (2009, 288), is shown in Figure 14.

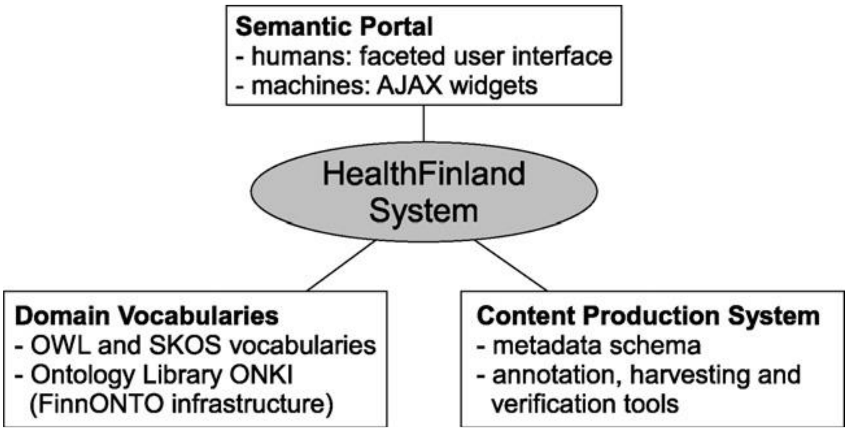

Figure 14. Three main components of HealthFinland (Suominen et al. 2009, 288)

As seen in Figure 14 above, HealthFinland comprises a centralised content infrastructure of health ontologies and services with tools. There is also a distributed semantic content production system that contains specifications and tools for annotating, harvesting and verifying content. Finally, there is an intelligent semantic portal that aggregates and presents the contents from end-user perspectives for human users and other websites and portals. (Suominen et al. 2009, 288.)

While Figure 14 depicts a basic overview of the components that constitute HealthFinland, a more detailed overview of the HealthFinland system architecture is illustrated in Figure 15.

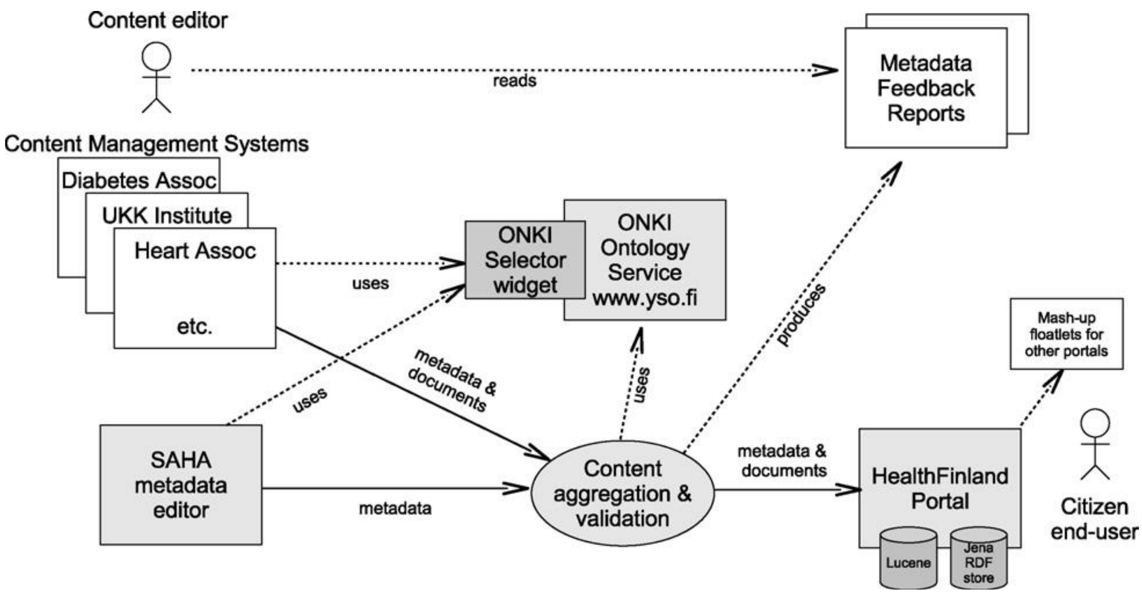

Figure 15. System architecture of HealthFinland (Suominen et al. 2009, 289)

In Figure 15, it can be seen that the content publishers retrieve metadata and documents by harvesting content from their content management systems. Content is also annotated manually with the SAHA metadata editor that is linked with the ONKI ontology services. Additionally, the content is validated, while the content providers receive reports of potential problems encountered. As a practical outcome, the validated metadata is then published through the portal for the use of humans and machines. (Suominen et al. 2009, 288.)

Essentially, the HealthFinland system demonstrates how heterogeneous content sources from different publishers can be aggregated through shared meanings with the application of ontologies. The collaborative publication of health content and reduction of duplicate works are enabled by the process of content creation, validation and aggregation infrastructure. For the end-user, the underlying Semantic Web technologies enable a search user interface that is citizen-centred, thereby providing health information that is actually relevant to the needs of the citizens. (Suominen et al. 2009, 296.)

7 CONCLUSIONS

In today's information-driven society, the Web presents immense opportunities for information retrieval. However, considering the plethora of resources on the Web, retrieving the right kind of information poses a problem. Despite the best efforts of Web developers, the Web, in its current state, is not intelligent enough to process information. When one considers that health information is of utmost importance especially since health issues are usually critical in nature, this perspective becomes further appreciated. The Semantic Web undertakes to enhance the current Web by ensuring that information becomes machine-interpretable. This enhancement will be achieved by the injection of semantics to the Web, thereby effectively creating an intelligent Web. This incorporation of intelligence to the Web offers immense potential, particularly to the healthcare industry. The eHealth domain in particular, has begun to adopt some Semantic Web technologies in recent years. The distributed nature of the healthcare industry and recent computerisation of healthcare practices suggest that this adoption has become inevitable. This is further validated by the increased use of ontologies in the medical field.

The focus of this research is to understand how Semantic Web Services technologies can be integrated into the eHealth domain. To demonstrate how such integration can be implemented, the Semantic Web and eHealth were studied as separate concepts. Furthermore, the different technologies of the Semantic Web and the tools of eHealth were described. These tools and technologies form the basis of the integration of the Semantic Web with eHealth. The implementation of medical search engines and CDSSs substantiate the applicability of these tools and technologies. To further illustrate the practical relevance of this research, some specific Web applications being used in the healthcare industry were presented. Therefore, while this research is entirely theoretical in its approach, the existence of eHealth systems integrated with Semantic Web Services technologies indicates that the content of this research has a practical value.

Applying SWS technologies in the eHealth domain is an important step in improving healthcare services. Such integration rids healthcare practices of errors to a minimum, while ensuring accurate services. The use of ontologies ensures that medical terminologies are efficiently represented. Furthermore, when these ontologies are applied to medical search engines, relevant health-related information becomes more readily accessible to users. A semantically-enabled CDSS also ensures that more informed and appropriate healthcare decisions are made by clinicians. These are vital exploits that have the potential to revolutionise the healthcare industry. Once improved healthcare services become accessible, these exploits have a tendency to improve the overall well-being of humanity.

However, the application of SWS technologies in the eHealth domain poses some challenges. Due to the relative novelty of the technologies, these challenges are expected. One of such previously identified challenges is that implementing semantic interoperability in healthcare is still problematic (Valle et al. 2010, 384). The transitioning of current eHealth systems into semantically-enabled infrastructure is a laborious exercise that requires conscientious planning and execution. Furthermore, as health matters are delicate, there are the issues of privacy, security and trust regarding patient data. eHealth systems that employ Semantic Web infrastructure, aiming for the interoperability of medical data, exemplify this particular challenge. The privacy and security of patient data shared across healthcare enterprises is one that must be addressed. More so, as pointed out in Chapter 4, the trust layer of the Semantic Web architecture is yet to be standardised. Consequently, there is no assurance that all health information offered by the Semantic Web platform will be infallible. As it is with most technological advancements, the efforts and financial resources invested in developing new systems often come at a cost. Therefore, increased healthcare expenses are one probable consequence of integrating the Semantic Web with eHealth.

Despite these challenges identified above, there are substantial benefits to be reaped from the integration of the Semantic Web with eHealth. The seamless flow of data between trusted healthcare institutes is one that can considerably facilitate healthcare services. Such rapid and easy data access is particularly important in cases of emergency. Additionally, a semantically-enabled healthcare system, whether in the form of medical search engines, or in the form of CDSS, results in a high level of flexibility in delivering healthcare services. Such flexibility entrusts patients with more responsibility concerning their own health. Furthermore, integrating SWS technologies into eHealth induces minimal human intervention. Consequently, the automated processes will provide a more efficient eHealth system.

The integration of SWS technologies into eHealth systems offers a wide range of possibilities. While these integration efforts have a major focus within Europe, they are equally relevant on a global perspective. Therefore, non-European countries are encouraged to intensify their involvement in such integration projects, especially regions with substandard healthcare systems. Further research may be conducted into the use of Semantic Web technologies in Emergency Management Systems, which requires a high level of collaboration and interaction at different levels of healthcare. Another research direction that could be explored is the development of a Semantic Web-based medical social networking platform, in which authorised health personnel are able to deliver sound healthcare services over the Web through interactive methods. Within the framework of this research, further research into Semantic Web Services can be conducted to overcome the already identified challenges. Furthermore, the current and potential benefits of Semantic Web Services can be exploited to further enhance the eHealth domain.

REFERENCES

Alexa Internet, Inc. 2013. Alexa Top 500 Global Sites. Retrieved April 17, 2013. <http://www.alexa.com/topsites>

Bodenreider, O. 2008. Biomedical Ontologies in Action: Role in Knowledge Management, Data Integration and Decision Support. IMIA Yearbook of Medical Informatics. Retrieved April 20, 2013. 67–79. <http://www.schattauer.de/de/magazine/uebersicht/zeitschriften-a-z/imia-yearbook/imia-yearbook-2008/issue/special/manuscript/9821/download.html>

Bose, Ranjit & Sugumaran, Vijayan 2007. Semantic Web Technologies for Enhancing Intelligent DSS Environments. In Kulkarni, Uday & Power, Daniel J. & Sharda Ramesh (eds.) 2007. Decision Support for Global Enterprises. Annals of Information Systems. Vol. 2. New York City, US: Springer, 221–238.

Bray, Tim & Paoli, Jean & Sperberg-McQueen, C. M. & Maler, Eve & Yergeau, François 2008. Extensible Markup Language (XML) 1.0 (Fifth Edition). W3C Recommendation. Retrieved April 16, 2013. <http://www.w3.org/TR/xml/>

Bundorf, M. Kate & Wagner, Todd H. & Singer, Sara J. & Baker, Laurence C 2006. Who Searches the Internet for Health Information? Health Services Research. Retrieved February 2, 2013. <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1713205/>

Cheung, Kei-Hoi & Yip, Kevin Y. & Townsend, Jeffrey P. & Scotch, Matthew 2008. HCLS 2.0/3.0: Health Care and Life Sciences Data Mashup Using Web 2.0/3.0. Retrieved December 29, 2012. <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2933816/>

Dietze, Heiko & Schroeder, Michael 2009. GoWeb: a semantic search engine for the life science web. BMC BioInformatics. Retrieved April 9, 2012. <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2755828/>

European Commission 2007. ICT for Health: Portfolio of the eHealth Projects in the FP6. Luxembourg: Office for Official Publications of the European Communities.

Eysenbach, G. 2001. What is e-health? Journal of Medical Internet Research. Vol. 3, No. 2. Retrieved March 21, 2013. <http://www.jmir.org/2001/2/e20>

Eysenbach, Gunther 2003. The Semantic Web and healthcare consumers: a new challenge and opportunity on the horizon? Int. J. Healthcare Technology and Management. Toronto: Inderscience Enterprises Ltd. Retrieved April 24, 2013. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.69.1618&rep=rep1&type=pdf>

Fox, Susannah 2010. The Power of Mobile. Pew Internet & American Life Project. Retrieved March 14, 2013. <http://pewinternet.org/Commentary/2010/September/The-Power-of-Mobile.aspx>

Fox, Susannah 2011. Health Topics. Pew Internet & American Life Project. Retrieved March 14, 2013. <http://pewinternet.org/Reports/2011/HealthTopics/Part-1/59-of-adults.aspx>

Golbreich, Christine & Wallace, Evan K. (eds.) 2012. OWL 2 Web Ontology Language. New Features and Rationale (Second Edition). W3C Recommendation. Retrieved March 16, 2012. <http://www.w3.org/TR/owl2-new-features/>

Grimm, Stephan & Hitzler, Pascal & Abecker, Andreas 2010. Knowledge Representation and Ontologies. In Abecker, Andreas & Grimm, Stephan & Studer, Rudi (eds.) 2010. Semantic Web Services: Concepts, Technologies and Applications. Berlin Heidelberg: Springer, 51–106.

Hitzler, Pascal & Krötzsch, Markus & Rudolph, Sebastian 2010. Foundations of Semantic Web Technologies. Boca Raton, Florida: Chapman & Hall/CRC, Taylor & Francis Group.

Hussain, Sajjad & Abidi, Samina Raza & Abidi, Syed Sibte Raza 2007. Semantic Web Framework for Knowledge-Centric Clinical Decision Support Systems. In Abu-Hanna, A. & Bellazzi, R. & Hunter J. (eds.) 2007. Berlin Heidelberg: Springer, 451–455. Retrieved April 20, 2013. <http://www.researchgate.net/publication/221450273_Semantic_Web_Framework_for_Knowledge-Centric_Clinical_Decision_Support_Systems/file/d912f5118b8a570f55.pdf>

Jessen, Walter 2007. Web 3.0 and Predictive, Preventive and Personalized Medicine. Retrieved September 12, 2012. <http://www.highlighthealth.com/healthcare/web-30-and-predictive-preventive-and-personalized-medicine/>

Khan, Haseeb & Hederman, Lucy 2012. A Universal Clinical Decision Support System using semantic web services. School of Computer Science and Statistics, Trinity College Dublin. Retrieved April 5, 2013. <http://hdl.handle.net/2262/64246>

Leroy, Gondy 2009. Persuading Consumers to Form Precise Search Engine Queries. AMIA Annual Symposium Proceedings Archive. Retrieved March 28, 2013. <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2815418/>

Luo, Gang 2008. Intelligent Output Interface for Intelligent Medical Search Engine. Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence. Retrieved March 19, 2013. <http://www.aaai.org/Papers/AAAI/2008/AAAI08-190.pdf>

Manola, Frank & Miller, Eric 2004. RDF Primer. W3C Recommendation Retrieved March 16, 2013. <http://www.w3.org/TR/rdf-primer/>

Matthews, Brian 2005. Semantic Web Technologies. JISC Technology and Standards Watch. Retrieved March 15, 2012. <http://www.researchgate.net/publication/30408878_Semantic_Web_Technologies/file/9fcfd50cb208e0360f.pdf>

McDaniel Jr., Carl & Gates, Roger 2010. Marketing Research Essentials with SPSS. 7th Edition. Hoboken, NJ: John Wiley & Sons Inc.

Nash, David B. 2008. Health 3.0. A Peer-Reviewed Journal for Managed Care and Hospital Formulary Management. Vol. 33, No. 2. Retrieved April 1, 2013. <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2730068/>

National Research Council 2001. Crossing the Quality Chasm: A New Health System for the 21st Century. Washington, D.C.: Institute of Medicine. The National Academies Press. Retrieved March 8, 2013. <http://www.nap.edu/openbook.php?record_id=10027>

Neter, Efrat & Brainin, Esther 2012. eHealth Literacy: Extending the Digital Divide to the Realm of Health Information. Journal of Medical Internet Research. Vol. 14, No. 1. Retrieved February 2, 2013. <http://www.jmir.org/2012/1/e19/>

Norman, Cameron D. & Skinner, Harvey A. 2006. eHealth Literacy: Essential Skills for Consumer Health in a Networked World. Journal of Medical Internet Research. Vol. 8, No. 2. Retrieved February 2, 2013. <http://www.jmir.org/2006/2/e9/>

Oh, Hans & Rizo, Carlos & Enkin, Murray & Jadad, Alejandro 2005. What is eHealth (3): A Systematic Review of Published Definitions. Journal of Medical Internet Research. Retrieved February 2, 2013. <http://www.jmir.org/2005/1/e1/>

Pagliari, Claudia & Sloan, David & Gregor, Peter & Sullivan, Frank & Detmer, Don & Kahan, James P. & Oortwijn, Wija & MacGillivray, Steve 2005. What Is eHealth (4): A Scoping Exercise to Map the Field. Journal of Medical Internet Research. Retrieved February 5, 2013. <http://www.jmir.org/2005/1/e9/>

Pisanelli, Domenico M. 2007. Ontologies. CNR-ISTC Laboratory of Applied Ontology, Rome. Retrieved February 24, 2013. <http://www.openclinical.org/ontologies.html>

Preist, Chris 2010. Goals and Vision: Combining Web Services with Semantic Web Technology. In Abecker, Andreas & Grimm, Stephan & Studer, Rudi (eds.) 2010. Semantic Web Services: Concepts, Technology and Applications. Berlin Heidelberg: Springer. 159–178.

Rovan, Lidia & Jagušt, Tomislav & Baranović, Mirta 2011. Defining Categories and Functionalities of Semantic Web Applications. International Journal of Systems Applications, Engineering & Development, 245 – 253. Issue 2, Vol. 5. Retrieved February 15, 2013. <http://www.universitypress.org.uk/journals/saed/19-826.pdf>

Schmidt, Guenter & Schaepe, Arno & Heydler, Thomas & Rinecker, Markus & Binnig, Gerd 2012. Clinical Decision Support System. United States Patent Application Publication. Retrieved April 9, 2012. <http://www.google.com/patents/US20120232930>

Sim, Ida & Gorman, Paul & Greenes, Robert A. & Haynes, R. Brian & Kaplan, Bonnie & Lehmann, Harold & Tang, Paul C. 2001. Clinical Decision Support Systems for the Practice of Evidence-based Medicine. Journal of the American Medical Informatics Association, 527 – 534. Vol. 8, No. 6. Retrieved April 9, 2013. <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC130063/pdf/0080527.pdf>

Sorensen, Tove & Andreassen, Hege 2010. E-health Trends in Europe: New Tools – New Practices? Norwegian Centre for Integrated Care and Telemedicine, Tromso, Norway. Oral Presentation. Retrieved February 20, 2013. <http://www.medicine20congress.com/ocs/index.php/med/med2010/paper/view/510>

Silber, Denise 2003. The Case for eHealth. Presented at the European Commission's first high-level conference on eHealth. Retrieved December 29, 2012. <https://www.eipa.eu/UserFiles/File/Free%20Publications/2003_E_01.pdf>

Stenzhorn, Holger & Schulz, Stefan & Boeker, Martin & Smith, Barry 2008. Adapting Clinical Ontologies in Real-World Environments. Journal of Universal Computer Science. Vol. 14, No. 22. Retrieved March 30, 2013. <http://jucs.org/jucs_14_22/adapting_clinical_ontologies_in/jucs_14_22_3767_3780_stenzhorn.pdf>

Studer, Rudi & Grimm, Stephan & Abecker, Andreas 2010. Introduction. In Abecker, Andreas & Grimm, Stephan & Studer, Rudi (eds.) 2010. Semantic Web Services: Concepts, Technologies and Applications. Berlin Heidelberg: Springer, 1–11.

Suominen, Osma & Hyvönen, Eero & Viljanen, Kim & Hukka, Eija 2009. HealthFinland – A national semantic publishing network and portal for health information. Web Semantics: Science, Services and Agents on the World Wide Web. Vol. 7, No 4. North America: Elsevier.

Trzebucki, Alex 2008. E-Health 3.0: How Can The Semantic Web Change The World Of Internet Health Information? Healthcare Innovation and Technology Lab. Retrieved September 12, 2012. <http://www.medicalnewstoday.com/articles/111201.php>

Valle, Emanuele Della & Cerizza, Dario & Celino, Irene & Dogac, Asuman & Laleci, Gokce B. & Kabak, Yildiray & Okcan, Alper & Gulderen, Ozgur & Namli, Tuncay & Bicer, Veli 2010. An eHealth Case Study: Applying Semantic Web Service Technology in the Healthcare Environment. In Abecker, Andreas & Grimm, Stephan & Studer, Rudi (eds.) 2010. Semantic Web Services: Concepts, Technologies and Applications. Berlin Heidelberg: Springer, 381–402.

W3C OWL Working Group 2012. OWL 2 Web Ontology Language Document Overview (Second Edition). W3C Recommendation. Retrieved April 16, 2012. <http://www.w3.org/TR/owl2-overview/>

Watson, Rory 2004. EU wants every member to develop a "roadmap" for ehealth. Electronic Communication and Health Care. Retrieved March 21, 2013. <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC411133>

World Health Organization 2006. eHealth Tools and Services. Needs of the Member States. Report of the WHO Global Observatory for eHealth. Retrieved March 14, 2013. <http://www.who.int/kms/initiatives/tools_and_services_final.pdf>

Yu, Liyang 2007. Introduction to the Semantic Web and Semantic Web Services. Florida, London, New York: Chapman & Hall/CRC, Taylor & Francis Group.

Zheng, Kai & Mei, Qiaozhu & Hanauer, David A. 2011. Collaborative search in electronic health records. Journal of the American Medical Informatics Association. Retrieved March 19, 2013. <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3078661/>