

18

MedCERTAIN/MedCIRCLE: Using Semantic Web Technologies for Quality Management of Health Information on the Web

GUNTHER EYSENBACH

The Semantic Web

What Is the Semantic Web?

The Semantic Web is a vision: the idea of having data on the web defined and linked in a way, that it can be used by machines—not just for display purposes, but for using it in various applications. (<http://www.semanticweb.org/introduction.html>)

This chapter examines how the World Wide Web might evolve in the near future, from an “information jungle” environment with largely narrative, human-understandable information, to a global knowledge repository, where much of the information is machine readable and directly processable by computers, enabling the use of advanced knowledge management technologies to steer consumers to trusted health information. This vision has been called the “Semantic Web” by its inventor, Tim Berners-Lee, who from the beginning had envisaged the Web to be a worldwide, distributed knowledge base rather than a medium with primarily narrative information targeted only for human consumption, as the web presents itself today [1–3]. The “Semantic Web” can be thought of an extension of the present Web, with Web authors publishing an additional layer of machine-processable data beneath the visible layer of human-readable information. It is essentially an attempt to create a global, decentralized knowledge base, represented as a “semantic net” that is woven by a large heterogeneous community of “authors.” A semantic net is a knowledge representation method with a long history in Artificial Intelligence ever since first introduced by Quillian back in the late 1960s [4]. Semantic nets use the idea that the semantics (meaning) of a concept comes from its relationship to other concepts. In other words, information—a collection of unrelated facts—becomes knowledge if it is contextualized by making links to related concepts explicit. In a semantic net, the concepts can be graphically depicted as nodes, and the relationships (links) between the nodes can be illustrated as labeled arcs. Turning the current World Wide Web into a global semantic net requires that (at least some) individuals and organizations who currently publish information on the Web will publish additional machine-processable documents (e.g., using XML) which describe the concepts and their relationships with other concepts (which may be defined on other Web sites) unambiguously, so that software can aggregate this knowledge and draw inferences. The idea of the MedCERTAIN project and its successor organization MedCIRCLE, an international nonprofit collaboration of consumer health information gateways and health Web sites, is that members of the collaboration make stan-

standardized machine-processable statements about health Web sites and health information providers using a common agreed-on vocabulary with agreed-on semantics, so that applications can aggregate, “understand,” and use this meta-information, for example, to match it against individual consumer needs, and to help consumers to make informed choices. This is the basis for decentralized quality management on the Web, without a single organization being in charge of “accrediting” health Web sites in a “top-down” fashion [5].

It could be argued that on today’s Web people are already expressing relationships with other chunks of information in a “machine-processable” way by using hyperlinks. However, hyperlinks are semantically ambiguous; that is, they can imply many different kinds of relationships, such as a reference (“see also”), an endorsement (“recommended reading”), or sometimes even something completely else, for example, pointing to contradictory information. In addition, a hyperlink usually links from a word or text phrase (which may be semantically ambiguous, in that others may use the same word but mean a different concept) to another Web page (which also is ambiguous, as it is not clear which concept is meant, that is, does the relationship refer to the individual behind that Web site, the topic discussed on the Web site, etc.). The sentence, “For further information see <http://www.healthfinder.org>” may be understandable to humans, but a piece of software cannot easily figure out which relationship between which entities this statement implies, unless some natural language processing software is employed. Indeed, Meric [6] has reported that there is no correlation between the number of links pointing to a site and “quality” as defined by health professionals (although it is debatable whether this gold standard of “quality” is the correct criterion). It is also obvious that an expression of trust or quality should not be simply binary (trust yes/no, quality yes/no), but needs to be more explicit in why and which aspects are trusted. The need for a more expressive “vocabulary” and language to express the meaning of relationships between sites (but also people, organizations, etc.) is obvious.

In contrast, on the Semantic Web, statements such as “see-also” or relationships between actors on the Web (such as “is-member-of” or “has-certified”) [7] would be published in an unambiguous, machine-processable way, and can be processed by software that can draw inferences from the knowledge chunks provided on different Web sites.

Building standards and tools for the Semantic Web is currently one focus of the activities of the World Wide Web Consortium (W3C), which describes the aim of its activities as follows: “The goal of the Semantic Web is to develop enabling standards and technologies designed to help machines understand more information on the Web so that they can support richer discovery, data integration, navigation, and automation of tasks. With the Semantic Web we not only receive more exact results when searching for information, but also know when we can integrate information from different sources, know what information to compare, and can provide all kinds of automated services in different domains from future home and digital libraries to electronic business and health services” [8].

Metadata

One prerequisite for the Semantic Web is that authors of Web sites and Web documents provide richer machine-processable information, essentially metadata. Metadata are “data about data.” The vision of using standardized metadata on health Web sites can be compared with food labels: Similar to producers of food, who have

to display ingredients on standardized labels, telling consumers, for example, the amount of fat and sodium contained in their products, health information providers on the Web should use standardized labels to disclose certain facts about their information, so that consumers can make informed decisions [5,9,10]. Until 1999, there had been many different ways to link metadata to Web documents, for example, using META tags in HTML or using PICS (Platform for Internet Content Selection) for self- and third-party descriptions of information. Although PICS was developed primarily with description and rating adult Web sites in mind, a vocabulary to describe and “label” health Web sites was developed in 1997 [11,12]. The W3C subsequently unified different approaches and the result of these efforts is the Resource Description Framework (RDF)—the current standard based on XML (extensible Markup Language) to transport metadata and a major pillar of the Semantic Web [13–15]. One feature of RDF is that (other than, for example, the HTML META-tag) it allows people to describe concepts and resources other than a Web document. In contrast, by using the HTML META-tag and a set of keywords the developer implicitly makes a statement about the document or Web site (but often it is not even clear whether the keywords refer to the document or the entire Web site), but cannot make more broad statements, for example, about other resources or concepts, as with RDF. Further, RDF provides a mechanism for giving unambiguous meanings to metadata keywords. In contrast, keywords used in META tags are essentially just ambiguous “words” that have no meaning (semantics) for software as they are not linked to other concepts. Words can be ambiguous in that they may have different meanings. For example, the word “virus” can refer to a computer virus or a biological virus. RDF provides a mechanism to define what kind of “virus” is meant by referring to the RDF statement or site where this concept is defined [again, through its relationship to other concepts, for example “virus (as defined in this statement) is-a software,” linking the word software again to another RDF document on the Web that defines “software,” etc.], thereby creating “meaning.”

As noted earlier, RDF can be expressed in XML syntax [15]. Although RDF is basically an XML file, the difference between an RDF document and a “plain” XML document is significant: Whereas XML-Schemas only tell computers (and us) how, for example, an application form for a driver’s license *looks like*, RDF is able to explain to a machine what a driver’s license *is*, by providing the meaning of the concepts used in a driver’s license. This is done by providing the relationships of the concepts to other concepts. As the RDF developers point out, RDF is a simple frame system, that is, a format for knowledge representation, where objects (concepts) and their relationships to each other are specified. The RDF specification does not contain a reasoning system; this needs to be built on top of it.

Unfortunately, the uptake of providing metadata on the Web—even in its simplest, nonsemantic form, the META tag—has been slow so far: Web content is still largely devoid of metadata labels [13] and a critical mass of metadata has to be generated before applications can be developed making use of it. The MedCERTAIN/MedCIRCLE projects (explained in detail later) developed some open source tools for health information providers to enter disclosure information deemed ethical (see Chapter 4, this volume) as machine-processable metadata. The health information provider does not need to understand RDF—all he or she needs to do is to fill in a questionnaire for self-disclosure and description, and his or her answers will be translated into metadata [16]. Finally, existing tools for creating knowledge bases, such as Protégé-2000, can be used to create RDF statements [17], and future Web editors may provide additional functionalities to model knowledge and build knowledge bases.

Application Scenarios

If the vision of the Semantic Web becomes reality, this will have a profound impact on how people will interact with the Web and obtain information. The first and most obvious change will include the markedly improved abilities of search engines to conduct accurate and relevant searches on the Web, and to guide users to trusted and relevant health information. Search engines will not only better “understand” what a user is looking for, but also what the Web pages they are indexing are about. They can, for example—if a user looks for “SARS in Canada”—recognize that the user is likely looking for information on severe acute respiratory syndrome rather than the South African Revenue Service, and then list only those Web sites that contain information about the disease and not the Revenue Service. The results will even include links to relevant Web pages that do not use any of the search terms—for example, if a Web page contains the word SRAS (for syndrome respiratoire aigu sévère), it will be found as well, because the search engine looks for semantic rather than syntactic matches, and the Web crawler has previously “understood” the context in which the word has been used and what the Web page on which it appeared is all about.

The idea of the MedCERTAIN/MedCIRCLE approach (described later) is that results in search engines will be better “ranked” not only by relevance but also by “quality,” for example, the degree of how trusted a health resource is in a community.

Accessibility and quality issues of health information on the Web are especially hot topics in the medical literature and subject of hundreds of empirical “descriptive infodemiology” [18] studies. These studies mostly suggest that it is hard for consumers to find high-quality health information among a flood of dubious or commercially driven information [19]. Surveys such as the Pew Internet Survey also show that 86% of consumers are concerned about getting low-quality health information on the Web [20]. While empirical studies now provide more than sufficient evidence on the inadequacies of the current Web, there is a surprising lack of debate in the medical world discussing the possibilities of technology to address these problems—presumably as many of the current developments in the field are unknown or remain not understood. The current MedCIRCLE Collaboration for Internet Rating, Certification and Labelling of Health Information, a global collaborative network of health information gateways described in detail later, is working toward this aim by enriching the current Web with machine-processable evaluation and trust data.

Knowledge Translation for Consumers on the Semantic Web

From Information to Knowledge

The Web as it exists today has played a significant role in fostering consumerism in health care [21,22]. The current Web provides an abundance of information, but giving “information” to a patient is certainly not enough. The ultimate goal is to enhance “knowledge”: the 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 (e.g., in the health record) have to be made explicit. This is the difference between “information” and “knowledge.” The Semantic Web enhances the possibility of supporting “knowledge translation” for consumers, the translation of information into knowledge. Doctors who are confronted with “Web-informed” patients often

complain that patients often find irrelevant information on the Web—information the patient (and the clinician) have to sift through and evaluate, and that is often not applicable to the individual situation [23]. Many patients do not even know the correct names of their diagnoses and are therefore unable to enter the correct terms into search engines. The vision for the future is that people will use their Web-based personal health record as a starting point that may be enriched by all kinds of information gathered by intelligent agents from trusted sources on the Web that are specifically relevant to the patient [24]. For example, if the Web-based health record contains a certain diagnosis, and on the same day the *British Medical Journal* publishes new research results published about this disease, the agent (which would be a part of the electronic health record software) could automatically generate a link to that article. It doesn't matter if the *British Medical Journal* article uses a different terminology than the doctor in the health record, as the agent will be able to link the terminologies. The Web-based electronic health record would be a dynamic entry point and knowledge management platform for patient and health professionals alike. Challenges, described in detail elsewhere, include privacy and disintermediation [25].

Using the Semantic Web for Steering Patients to Best Quality Health Care

Perhaps most challenging for healthcare providers is the prospect that people will use the Web not only to locate the least expensive used car in their neighbourhood, but also to search for the best quality healthcare providers, taking into account their own preferences and decentralized data from different sources such as hospital report cards, specialized providers of healthcare performance data such as *healthgrades.com*, and—perhaps most significantly—also based on ratings given by fellow patients with the same conditions and similar demographic background [26]. The Semantic Web makes relationships between things explicit and computable, and therefore further increases the transparency for consumers, much as the current Web has already made it easier to compare prices and offers, revolutionizing other areas such as the travel industry. The Semantic Web will make it even easier to compare things, as software can, for example, map different terminologies and aggregate decentralized knowledge dispersed all over the Web. For example, software agents would roam the Web and return information on who has the best offer of a certain car model in a given community. Similarly, software could be used to aggregate experiences of people with all kinds of health services and products, including, for example, their experience with over-the-counter or prescription drugs, hospitals, or individual physicians. While today patients use primarily mailing lists, newsgroups, and chat rooms to exchange anecdotal and narrative information and experiences about health products, services, and providers, patients could publish their experiences about virtually anything and everything in RDF on homepages—from experiences with a new dishwasher to experiences with healthcare professionals, hospitals, or drugs. Patients could rate their treatments and services directly in the Web-based electronic health record and feed them (in anonymized form) into the Semantic Web (e.g., hospitals and doctors provide RDF dumps of their patients on their sites), so that agents can aggregate this information. Such “knowledge” evolving on the Web could also be used systematically for post-marketing surveillance efforts to monitor the ongoing safety of marketed drugs on a global scale.

Overcoming Quality Issues as Opportunity

When people write and talk about the Semantic Web today, they mainly stress the advantages for information retrieval. However, the Web is an information space that reflects not just human knowledge but also human relationships; thus the Semantic Web can also represent trust relationships among people and organizations.

“Trust management” is a prerequisite for successful knowledge management on the web. Without the possibility for people to filter information or for agents to make semiautomated decisions on which knowledge chunks, ontologies, or sources to trust, the jewels on the Web will be lost in a “noise” of imperfect, cheaply produced or commercially motivated, biased information.

Although central authorities to regulate, control, censor, or centrally approve information, information providers, or Web sites are neither realistic nor desirable [5], health professionals are still interested in making systems available that direct patient streams to the best available information sources.

MedCERTAIN and MedCIRCLE

The author of this chapter has argued for many years that on a decentralized, electronic medium such as the Web, a global metadata infrastructure is the most appropriate answer to the current debate on the “quality of health information on the Web.” One has to think along the lines of a collaborative “Semantic Web of trust” when it comes to the question on how consumers can be steered (or can steer themselves!) to the best available health information on the Web [5,7,12,27,28]. A “Collaboration for Critical Appraisal of Health Information on the Web”—a loose community of health information providers and health gateways using metadata to describe and annotate health Web sites—had been proposed as early as in 1997, and mentioned in two seminal articles in 1998 and 1999 [5,12]. Today, such a collaboration is known as the MedCIRCLE Collaboration, a loose nonprofit umbrella organization for health information gateways and health Web sites, inspired by the model of the Cochrane Collaboration. Membership is open to any organization using a standardized metadata vocabulary to express evaluative and descriptive statements about health information resources. The basic idea is that quality management on the Web should be based on a collaborative model with many actors (including health professionals and consumers) being able to say different things about anything in a machine-processable way (i.e., using metadata). This would enable software to analyze the trust relationships and would enable “downstream filtering” at the client computer or positive selection of trusted content using agents, instead of relying on upstream filtering approaches such as kitemarks [5] or even such well-intended but misguided proposals for (ab-)using top-level domains to centrally approve health information providers [29]. It would also allow search engines to rank their results according to quality and trust criteria of the individual user.

A metadata vocabulary for this purpose, MedPICS [based on the W3C PICS (Platform for Internet Content Selection Standard)] was first proposed in 1997, and also contained metadata elements that could be used by third parties to express evaluative statements about other sites [12]. The MedPICS proposal later led to the MedCERTAIN (2000–2001) and MedCIRCLE (2002–2003) projects, both of which aimed to implement such metadata on health Web sites and third-party organizations. With the PICS standard being superseded by XML/RDF [13], the projects became early “Semantic Web” projects, using RDF to transport and exchange metadata. As the

PICS standard became obsolete, MedPICS was renamed into HIDEEL (Health Information Disclosure, Description and Evaluation Language) [16]. Unlike other initiatives in this field, such as Health on the Net Foundation (HON), Centre for Health Information Quality kitemark (CHiQ), URAC Health Web Site Accreditation program, MedCERTAIN is *not* a traditional “kitemark” (i.e., seal of approval) project, but instead tried to develop an infrastructure and common ontology to link existing approaches, to make them interoperable, and to generate a critical mass of health-related descriptive and evaluative metadata on the Web. Unfortunately, the ideas behind MedCERTAIN/MedCIRCLE are not easy to communicate and the projects were consistently and repeatedly misunderstood and misrepresented as a “kitemarking” or third-party certification program [30], while the main goal—to develop and demonstrate a decentralized Web-of-trust infrastructure using of metadata—were not widely understood.

The constant misunderstandings concerning MedCERTAIN were one reason to change the project name to MedCIRCLE (Collaboration for Internet Rating, Certification, Labeling and Evaluation of Health Information), stressing the collaborative idea. The Collaboration involves a wider medical community to assess health information, demonstrating the power of collaborative and interoperable evaluations in a Semantic Web environment.

Figure 18.1 illustrates the operational model of health information providers collaborating in the MedCIRCLE. MedCIRCLE members are primarily trusted health information gateways, government portals, medical societies, accrediting organizations,

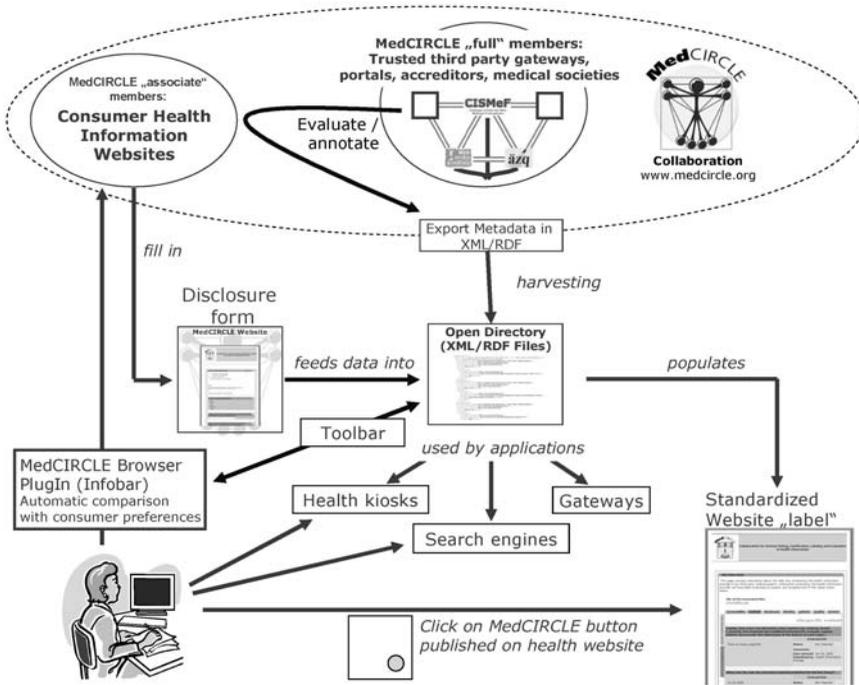


FIGURE 18.1. Operational model of the MedCIRCLE collaboration.

and libraries. What they have in common is that all are “third parties” that are in the business of describing, annotating, or making statements about other organizations, health information providers, or consumer health Web sites. For example, a medical society offering “recommended links for consumers” is a “gateway.” Rather than offering unspecific hyperlinks to “recommended sites,” the gateway can semantically enrich the endorsements by using a standardized vocabulary HIDDEN (Health Information Disclosure, Description and Evaluation Language) [16], expressed in XML/RDF, to report evaluation results in detail. Similarly, an organization in the business of “accrediting” health Web sites would use the vocabulary to express accreditation results. Among the current MedCIRCLE members are, for example, three major European gateway sites for consumer health information, two of which are backed by official professional physician associations. Other health subject gateways, accreditation, or rating services are encouraged to join the Collaboration simply by implementing HIDDEN on their gateways. The hope is to eventually establish a global Web of trust for networked health information.

As illustrated in Fig. 18.1, MedCIRCLE members export HIDDEN/XML/RDF data into an Open Directory. In addition, participating consumer health information Web sites can export disclosure and self-descriptive data into the Open Directory. Data in the Open Directory can be used by various applications and other Web sites under an Open Directory license, that is, free of charge, as long as the originator of the data and MedCIRCLE are acknowledged, and the integrity of the data is left intact. For example, MedCIRCLE gateways can display the data of other MedCIRCLE members, search engines can use data to rank their results, health kiosks can use the data to facilitate access to trusted Web sites, and client-side software, for example, browser plug-ins or “toolbars,” such as the MedCIRCLE infobar (Fig. 18.2), can make use of the data.

Conclusion

“Consumer health informatics” is the emerging science at the crossroads of health informatics and public health that deals with investigating determinants, conditions, elements, models, and processes to design, implement, and maximise the effectiveness of computerised information and telecommunication and network systems for consumers [31]. Nobel laureate economist Herbert A. Simon (quoted in Coiera’s paper on “information economics” [32]) once stated that “Information consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention, and a need to allocate that attention efficiently among the overabundance of information sources that might consume it.” One of the central topics of consumer health informatics is how to guide consumers to quality health information. Technology for producing and distributing information is useless without some way to locate, filter, organize, and summarize it. In that sense the Semantic Web remains a double-edged sword. The main opportunities lie in the fact that consumers will have even better possibilities to find, aggregate, and appraise health information than today. On the other hand, one might fear that this may lead to a further overreliance on external information, a process of disintermediation between patients and healthcare professionals, and erosion of the patient–physician relationship. Such concerns may not, however, stop the development of the Semantic Web, as the possibilities for e-commerce can be mind-boggling, in that search engines such as Google may evolve into marketplace managers and personal assistants to find, buy, and sell articles on the Web [33]. As health infor-

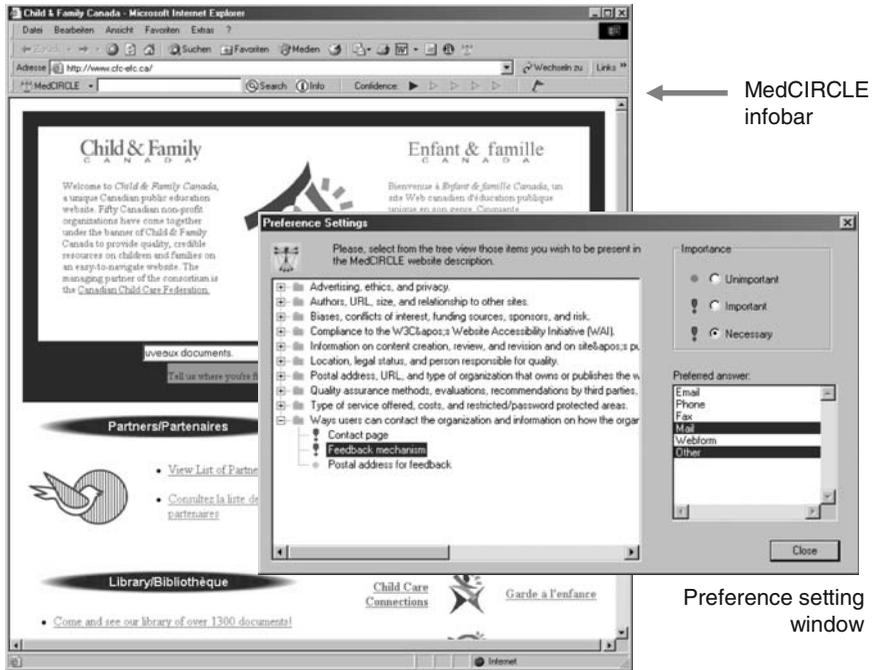


FIGURE 18.2. The MedCIRCLE infobar, a browser plug-in displaying a confidence score based on the user's preference setting concerning the presence of certain quality criteria, and what MedCIRCLE members say about the accessed health Web site.

mation is still some of the most sought after content on the Web, constituting of about 4.5% of all queries in search engines [34], people will not stop short of using these technologies for health products and services, researching the attributes and reputation of health products and services with a far greater sophistication than on today's Web. The World Wide Web as it exists today might be just the beginning of yet another consumer health informatics revolution.

References

1. Berners-Lee T. Weaving the Web. San Francisco: Narper, 1999.
2. Berners-Lee T, Hendler J, Lassila O. The Semantic Web. *Sci Am* 2001;284:34-43.
3. Berners-Lee T. Information Management: A Proposal. CERN, 1989.
4. Quillian MR. Semantic memories. In: Minsky MM, ed. *Semantic Information Processing*. Cambridge: MIT Press, 1968, pp. 216-70.
5. Eysenbach G, Dieppen TL. Towards quality management of medical information on the internet: evaluation, labelling, and filtering of information. *Br Med J* 1998;317:1496-500.
6. Meric F, Bernstam EV, Mirza NQ, et al. Breast cancer on the World Wide Web: cross sectional survey of quality of information and popularity of websites. *Br Med J* 2002;324:577-81.
7. Eysenbach G. An ontology of quality initiatives and a model for decentralized, collaborative quality management on the (semantic) world-wide-web. *J Med Internet Res* 2001;3:E34.
8. Koivunen MR, Miller E. W3C Semantic Web activity. In: *Proceedings of the Semantic Web Kick-off Seminar in Helsinki, Finland, November 2, 2001*.
9. Eysenbach G, Köhler C, Yihune G, Lampe K, Cross P, Brickley D. A framework for improving the quality of health information on the World-Wide-Web and bettering public (e-)health:

- the MedCERTAIN approach. In: Medinfo 2001, Proceedings of the Tenth World Congress on Medical Informatics Amsterdam: IOS Press, 2001, pp. 1450–4.
10. Eysenbach G, Yihune G, Lampe K, Cross P, Brickley D. Website labels are analogous to food labels [Letter]. *Br Med J* 2001;322:794.
 11. Eysenbach G, Yihune G, Lampe K, Cross P, Brickley D. Quality management, certification and rating of health information on the Net with MedCERTAIN: using a medPICS/RDF/XML metadata structure for implementing eHealth ethics and creating trust globally. *J Med Internet Res* 2000;2:2E1.
 12. Eysenbach G, Diepgen TL. Labeling and filtering of medical information on the Internet. *Methods Inform Med* 1999;38:80–8.
 13. Swick RR, Brickley D. PICS Rating vocabularies in XML/RDF. <http://www.w3.org/TR/rdf-pics> (2000).
 14. Klyne G, Carroll JJ. Resource description framework (RDF): concepts and abstract syntax. <http://www.w3.org/TR/2003/WD-rdf-concepts-20030123/> (accessed February 12, 2003).
 15. Lassila O, Swick RR. Resource description framework (RDF) model and syntax specification. <http://www.w3.org/TR/REC-rdf-syntax/> (1999).
 16. Eysenbach G, Köhler C, Yihune G, Lampe K, Cross P, Brickley D. A metadata vocabulary for self- and third-party labeling of health web-sites: Health Information Disclosure, Description and Evaluation Language (HIDDEL). *Proc AMIA Symp* 2001;169–73.
 17. Noy NF, Sintek M, Decker S, Crubezy M, Fergeson RW, Musen MA. Creating Semantic Web contents with Protege-2000. *IEEE Intell Syst* 2000;16:60–71.
 18. Eysenbach G. Infodemiology: the epidemiology of (mis)information. *Am J Med* 2002;113:763–5.
 19. Eysenbach G, Powell J, Kuss O, Sa ER. Empirical studies assessing the quality of health information for consumers on the world wide web: a systematic review. *JAMA* 2002;287:2691–700.
 20. Fox S, Rainie L. Pew Internet and American Life Project. The Online Health Care Revolution: How the Web helps Americans take better care of themselves. 2000. http://www.pewinternet.org/pdfs/PIP_Health_Report.pdf.
 21. Eysenbach G. Consumer health informatics. *Br Med J* 2000;320:1713–16.
 22. Eysenbach G, Jadad AR. Evidence-based patient choice and consumer health informatics in the Internet age. *J Med Internet Res* 2001;3:e19.
 23. Potts HW, Wyatt JC. Survey of doctors' experience of patients using the Internet. *J Med Internet Res* 2002;4:e5.
 24. Eysenbach G. Theme issue for medics and health informed public. What the future might hold for the BMJ in 2013. *Br Med J* 2003;326:166.
 25. Eysenbach G. The Semantic Web and healthcare consumers: a new challenge and opportunity at the horizon? *Int J Healthcare Technology and Management* 2003;5:194–212.
 26. Bates DW, Gawande AA. The impact of the Internet on quality measurement. *Health Aff (Millwood)* 2000;19:104–14.
 27. Eysenbach G, Yihune G, Lampe K, Cross P, Brickley D. MedCERTAIN: quality management, certification and rating of health information on the Net. *Proc AMIA Symp* 2000;230–4.
 28. Eysenbach G, Diepgen T, Lampe K, Brickley D. EU-project medCERTAIN: certification and rating of trustworthy and assessed health information on the Net. *Stud Health Technol Inform* 2000;77:279–83.
 29. Frenk J. Proposal to ICANN for health Internet Top Level Domain. ICANN. 2000. <http://www.icann.org/tlds/health1/WHO-C-SO-Proposal.htm>.
 30. Purcell GP, Wilson P, Delamothe T. The quality of health information on the Internet. *Br Med J* 2002;324:557–8.
 31. Eysenbach G, Jadad AR. Consumer health informatics: health information for consumers in the Internet age. In: Edwards A, Elwyn G, eds. *Evidence-Based Patient Choice*. Oxford: Oxford University Press, 2001.
 32. Coiera E. Information economics and the Internet. *J Am Med Inform Assoc* 2000;7:215–21.
 33. Ford P. This is August 2009: How Google beat Amazon and Ebay to the Semantic Web. *FTRAIN.com*. (2002).
 34. Eysenbach G, Kohler C. Health-Related Searches on the Internet. *JAMA* 2004;291:2946.