Preface

INTRODUCTION TO THE BOOK

Healthcare has always been both an opportunity and challenge for emerging technologies. For instance, the discovery of X-rays by Wilhelm C. R. Röntgen in 1895 created the opportunity for diagnostic imaging techniques to be developed in the 20th century. Specifically, it provided the basis for X-ray testing. Coupled with the mathematical foundation introduced by Johann Radon in 1917, the so-called Radontransformation, and digitization, it further enabled the invention of the computer tomography scanner by Godfrey N. Hounsfield and Allan M. Cormak in 1972 (Hounsfield and Cormak were awarded the Nobel award for medicine in 1979 for their invention). The implicit challenge lies in the long lag between discovery and invention.

Similarly, the incremental and still ongoing diffusion of information technology (IT) into healthcare has markedly changed the environment in which healthcare is studied, taught, and practiced. In addition, it has energized the establishment of the novel discipline of medical informatics. *Medical informatics* is evolving in two main directions. These are signal processing and information management. *Signal processing* in healthcare is applied, e.g., in equipment. *Equipment* is comprised of sensors and motors. For instance, *sensors* are used for monitoring the state of a patient (circulation, respiration, chemicals, etc.) during anesthesia or under emergency condition, while *motors* are applied in computer-supported (i.e., micro-) surgery.

Information management in healthcare is understood in a broad sense and includes a rich variety of activities. To mention a few: documentation of the patient's condition, diagnostic measures, treatment, etc.; monitoring of the patient's compliance during long-term treatment; diagnostic coding and other terminology-based application; clinical data warehousing; case-based reasoning; expert systems for drug prescription, second opinion, and education; information services for evidence-based medicine; and tools supporting the transfer and management of medical knowledge in practice.

With the advent of the Internet as a global communications infrastructure and of the World Wide Web as an interactive, ubiquitous, and multimedia information service the opportunities and challenges of IT in healthcare have gained a new dimension. It is within this context of an evolving "global knowledge medium" that the topics presented in this book must be considered. With respect to the rough classification, as introduced above, this book refers to the broad field of information management. The research question, which implicitly underlies this book's concept, can be sketched as follows: *How will the recent trend towards knowledge media change the way things are done in healthcare?* Because of the nature of the question a definite answer cannot be given at this time. However, the contributions presented here, coming as they do from the frontlines of research into IT and its application to healthcare, provide reasonable evidence on how healthcare *might* evolve in the future. Interestingly, the contributions to the book originate from quite different domains, i.e., computer science, medical science, and business sciences. This reflects the requirement of a multidisciplinary approach to the topic that is needed, in order to meet its challenges effectively.

CONTENTS OF THE BOOK

Foundations of Knowledge Media

Stanoevska-Slabeva provides a definition of knowledge media and an overview of their past and future. The term "knowledge media" was coined and their initial vision introduced by Mark J. Stefik. Stefik defines knowledge media "as an information network with semi-automated services for the generation, distribution, and consumption of knowledge." The first attempts to implement Stefik's concept of knowledge media were conducted at Stanford University by the Knowledge Sharing Effort of DARPA. The result of the joint effort is the concept for an Agent Communication Language (ACL), which enables interoperability and knowledge reuse among heterogeneous artificial agents. ACL consists of an inner language called Knowledge Interchange Format (KIF), an outer language called Knowledge Query and Manipulation Language (KQML), and common vocabularies, i.e., ontologies. Particularly, *ontologies* are considered the basic elements of knowledge media. *Stanoevska-Slabeva* then presents the concept of knowledge media as introduced by Beat F. Schmid from the University of St. Gallen. While adhering to Stefik's vision and the Stanford approach, it extends the original conception in several ways, particularly by introducing the concept of the *community*. The approaches of the Knowledge Media Design Institute at the University of Toronto and of the Knowledge Media Institute at the Open University in London are then presented. Even though these are not directly grounded in the original vision they reflect basic features of knowledge media. Finally, an overview of the Internet as the future global knowledge medium is provided. In this context, the recent initiative towards a *Semantic Web* is discussed.

Lanzola and Boley report on their experience with applying a functional-logic language to distributed, multi-agent medical problem solving. RELFUN, used both as the agent-implementation and the communication-content language, cross-extends Horn relations and call-by-value functions just enough to yield a unified operator concept. Particularly, relations acquire application nesting and higher-order notation; functions acquire non-groundness and non-determinism. Relations are defined by Horn-like clauses; functions are defined by rules with an additional returned-value premise. The chosen multi-agent framework relies on the communication protocol of KQML. The selected medical problem concerns task planning for patients affected by AML (acute myeloid leukemia). The medical professionals involved in AML treatment are modeled as agents communicating in a well-defined fashion. On the basis of a RELFUN representation of KQML performatives, both cooperation and domain knowledge are formalized in a functional-logic style. Its declarative power is shown to be useful for this application. However, for communicating clausal knowledge, update operations turn out to be unavoidable, and RELFUN is extended towards assert/retract. The efficiency of the interpreter has been sufficient in this communication-intensive prototype.

Boley examines mutual relationships between logic programming and XML. XML documents are introduced as linearized derivation trees and mapped to Prolog structures. Conversely, a representation of Herbrand terms and Horn clauses in XML leads to a pure XML-based Prolog (XmlLog). The XML elements employed for this are complemented by uses of XML attributes like id/idref for extended logics. XML's document type definitions are introduced on the basis of ELEMENT declarations for XmlLog and ATTLIST declarations for id/idref. Finally, queries in languages like XQL are treated functional-logically and inferences on the basis of XmlLog are presented. All concepts are explained via knowledge-representation examples from the discourse domain *E-health*.

Knowledge Representation and Human-Computer Interaction

Straub analyzes four different types of classification models. Classification models are intensely used in healthcare, e.g., for the coding of diagnoses or indexing of medical subjects. Starting from the problem of mapping clinical facts onto an artificial representation he introduces step by step more complex models, thereby evaluating each of them against the real world. With the fourth generation of classification models, i.e., the *multifocal, multi-point model*, which approximates real world best, he surpasses existing research and makes an original contribution to the topic. The presentation is supported by a rich number of examples from the practice of diagnostic coding.

Wagner argues that pharmaceutical expert information is characterized by steady increase in extent and complexity as well as by strong interweavement and penetration with interdependent references. These properties suggest the utilization of modern hypermedia technology, organizing knowledge as a network of linked pieces and enabling more intuitive, faster and easier navigation within the growing information space. He concludes that there is a need for a suitable, role-oriented *external* representation of pharmaceutical expert information. *Wagner* further argues that the application of pharmaceutical expert information when analyzing medications for contraindications and interactions is mainly mechanical in nature. It requires the ability to follow references and compare codes. The electronic availability of medication information, which is often forced by drug dispensing systems, is not exhausted when these procedures are performed manually at the screen. He concludes that there is a need for automation of these procedures. This automation requires a suitable *internal* representation of pharmaceutical expert information. *Wagner* further argues that the universality of pharmaceutical information systems leaves open their conceptual environment. The different kinds of application scenarios pose strong requirements on the formalization of expert information. He concludes that a *universal representation* has to be found which captures the nature of pharmaceuticals including common properties of drug classification systems. The diversity of the latter within the different healthcare organizations is to be regarded as the most serious barrier aggravating global communication within the healthcare system.

Kapetanios presents a semantically advanced query language for medical knowledge and decision support which enables a system-guided construction (formulation) of meaningful queries based on *meaning* of terms as they appear in medical application discourses. Representation of meaning and/or semantics is achieved by the structure of a knowledge space, i.e., a constraints-based graph formalism. This knowledge space is consulted by an inference engine which drives the query construction process. Both the knowledge space and the inference engine underlie the specification of MDDQL as the *Meaning Driven Data Query Language* at the core of the system. MDDQL requires neither learning of a particular language-based syntax for querying a data repository nor understanding of the underlying database schema in terms of adequate interpretations of data constructs. It embeds natural language based interpretation of the meaning of acronyms used as attributes and/or values and considers the context which relates to the data as expressed in terms of measurement units, definitions, constraints, etc. The system has been and is currently in use in various cardiological and gynecological hospitals and clinics in Switzerland.

Secure Healthcare Application and Data

Knorr and Röhrig argue that even though security requirements in healthcare are traditionally high, most computerized healthcare applications lack sophisticated security measures or focus only on single security objectives. Their chapter describes special security problems that arise when processing healthcare data using public networks such as the Internet. *Knorr and Röhrig* propose a structured approach using a context-dependent access control mechanism over the Internet as well as other security mechanisms to counter the threats against the major security objectives: *confidentiality, integrity, availability, and accountability*. The feasibility of the proposed security measures is shown through a prototype, which has been developed in a research project focussing on security in healthcare.

The concept of *context-dependent access control* has emerged during the last years: Information about the state of a process model of a working environment is combined with general knowledge about a person to grant or revoke access to protected data. Being understood very well in principle, different problems arise when implementing context-dependent access control, in particular on an open network. *Ultes-Nitsche and Teufel* report on an ongoing project on context-dependent access control to support distributed clinical trials. Centrally stored data will be accessed from contributors to the clinical trial over the Internet. They present in their chapter how context-dependent access control can be implemented on the Internet in a secure way. Technically they use Java servlets to implement the access control and SSL to secure communication. The whole framework is built around the Java Web server. *Ultes-Nitsche and Teufel* emphasize the technical aspects of this scenario in their chapter.

Knowledge Transfer and Management

Shaughnessy, Slawson and Fischer argue that traditional university-based teaching methods equip students with vast amounts of information but they fail to teach the skills for continuous learning. While computers and the Internet have made information readily available at everyone's fingertips, little consideration has been given to how information is delivered. The primary focus of evidence-based medicine is to identify and validate written information; for most doctors this is a too time-consuming process. They conclude that if best available evidence is to be used at the point of care, sophisticated *filters* are needed that increase the yield on relevance of the information. Doctors need an *alert method* for becoming aware of relevant new information that implies the need to update their knowledge; these systems should be tailored to the doctor's individual needs.

Similarly, *Khan, Bachmann, and Steurer* argue that with the enormous, rapid and exponential expansion in the medical literature, there is a need for effective and efficient strategies to keep abreast of relevant new knowledge. The *medical journal club*, traditionally used as an educational tool in postgraduate and continuing medical education, can be designed for acquisition and appraisal of relevant, current best clinical evidence by using a systematic approach to both acquisition and appraisal of evidence in a context directly related to patient care. According to *Khan, Bachmann, and Steurer*, incorporation of computer technology in the journal club helps with acquisition, appraisal and refinement of knowledge, but more importantly it allows for knowledge transfer by making possible the storage and instant retrieval of appraised topics in the future.

As mentioned, there is an enormous amount of new medical knowledge generated every year as a result of research. However to make a practical difference, this knowledge needs to be disseminated and used in everyday medical practice. *Moody and Shanks* describe a knowledge management project which provides medical staff with on-line access to the latest medical research at the point of care, in order to improve the quality of clinical decision making and to support evidence-based practice. The project has been highly successful, and a survey of medical staff using the system found that over 90% felt that it had improved the quality of patient care. They describe how the system was developed and implemented, its functional characteristics and organizational impact. The theoretical significance of their work is that it is one of the first empirical studies of a knowledge management project in the public sector. The practical significance of their research is that it provides a model for other similar organizations to follow in implementing such a project. Finally, *Moody and Shanks* draw some wider lessons from this case study for the practice of healthcare.

Towards a Global Knowledge Medium

According to *Fierz*, medical information processing by machines depends on the accessibility of medical information in electronic form. The manner in which such information is expressed and perceived, both by computers and by human beings, is fundamental to the success of any attempt to profit from modern IT in healthcare. He argues that medical information resides not, to a large extent, in the data content of information itself but is rather distributed within the connections (links), as well as within the structured context of the data elements. To enable a computerized processing of such information, some basic requirements for a structured and linked data model have to be fulfilled. These include the *granularity* of data elements; the way to attach *semantic information* to the data elements, links and structures; the *storage* of the data together with their structure and the connections between the data elements; a *query system* for the extraction of the information contained within the structure and connectivity as well as from the data proper; and the *display* of the query result in a way that structure and connectivity are intuitively and usefully expressed and can be stored again in a structured, machine-accessible way. *Fierz* argues that the developing Internet technology provides a suitable model for how these requirements can be fulfilled. The *hypertext* paradigm together with the *markup* technology for structuring information provides all the necessary ingredients for developing information networks that might be called *Medical Data Webs*. On top of that, Topic Maps and the Resource Description Framework (RDF) can be employed to semantically navigate through these Medical Data Webs.

RDF(S) constitutes a newly emerging standard for metadata that is about to turn the World Wide Web into a machine-understandable knowledge base. It is an XML application that allows for the denotation of facts and schemata in a Web-compatible format, building on an elaborate object-model for describing concepts and relations. Thus, it might turn up as a natural choice for a widely useable ontology description language. However, its lack of capabilities for describing the semantics of concepts and relations beyond those provided by inheritance mechanisms makes it a rather weak language for even the most austere knowledge-based system. *Staab, Erdmann, Maedche, and Decker* present an approach for modeling ontologies in RDF(S) that also considers axioms as objects that are describable in RDF(S). Thus, they provide flexible, extensible, and adequate means for accessing and exchanging axioms in RDF(S). Their approach follows the spirit of the World Wide Web, as they do not assume a global axiom specification language that is too intractable for one purpose and too weak for the next, but rather a methodology that allows (communities of) users to specify what axioms are interesting in their domain.

With the objective to facilitate integration of information from distributed and heterogeneous sources *Grütter, Eikemeier, and Steurer* describe the integration of an existing Web-based ontology on evidence-based medicine into the RDF/RDFS framework. Their contribution rests upon the application of methods which are to some extent already in place to a real-world scenario. Based on this application, the scope of the term "ontology" within the RDF/RDFS framework is redefined, particularly by introducing a *Simple Ontology Definition Language* (SOntoDL). This redefinition contributes to the implementation of the Semantic Web and to ontology modeling in general.

CONTRIBUTION OF THE BOOK

Knowledge Media in Healthcare: Opportunities and Challenges identifies areas of current research into knowledge media in healthcare on the basis of relevant examples from the fields of knowledge engineering, knowledge representation, human-computer interaction, application and data security, knowledge transfer, and knowledge management. Results from this research show *opportunities* for improved management of medical data, information, and knowledge, thereby supporting the provision of high quality and cost-efficient patient care.

The *challenges* rest upon the degree to which these opportunities are exploited, the ease with which technological innovations are made available, transferred into, and adopted from practice. Possible obstacles thereby include a mismatch between applications and users'–i.e., healthcare professionals'–needs, technology-adverse behavior, and technological illiteracy. In order to anticipate the first two, a lot of investigation into human needs is still required. Such investigations must not only focus on information needs but also on basic needs and how they can be met in an increasingly changing environment. Technology must never hinder but must further the development of human personality and social competency. In order to overcome technological illiteracy, investment into education, at all levels, is necessary. Hopefully, this book can contribute to this endeavor.

ACKNOWLEDGMENTS

It is particularly due to the significant research into knowledge media by Professor Dr. Beat F. Schmid, head of the Institute for Media and Communications Management, University of St. Gallen, that the editor was inspired to adopt the generic concept of knowledge media to the herein presented application domain. The editor wishes to thank the authors for their valuable insights and excellent contributions to this book. Likewise, the help of those involved in the review process, without whose support the project could not have been satisfactorily completed, is gratefully acknowledged. A further special note of thanks goes to the staff at Idea Group Publishing, who provided professional support throughout the whole process from inception of the initial idea to final publication.

Rolf Grütter, DVM, MSc
St. Gallen, Switzerland
March 2001