The coming of age of artificial intelligence in medicine ☆

Vimla L. Patel [a,*], Edward H. Shortliffe [a,b], Mario Stefanelli [c], Peter Szolovits [d], Michael R. Berthold [e], Riccardo Bellazzi [c], Ameen Abu-Hanna [f]

[a] *Department of Biomedical Informatics, Arizona State University, ABC1, 425 North Fifth Street, Phoenix, AZ 85004, USA*
[b] *Department of Basic Medical Sciences, University of Arizona College of Medicine, ABC1, Room 227, 425 North Fifth Street, Phoenix, AZ 85004-2157, USA*
[c] *Dipartmento di Informatica e Sistemistica, University of Pavia, Via Ferrata 1, 27100 Pavia, Italy*
[d] *Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, The Stata Center, 32-254, 32 Vassar Street, Cambridge, MA 02139, USA*
[e] *Department of Computer and Information Science, University of Konstanz, BOX M712, 78457 Konstanz, Germany*
[f] *Department of Medical Informatics, Academic Medical Centre, University of Amsterdam, Meibergdreef 15, 1105 AZ Amsterdam, The Netherlands*

Summary   This paper is based on a panel discussion held at the Artificial Intelligence in Medicine Europe (AIME) conference in Amsterdam, The Netherlands, in July 2007. It had been more than 15 years since Edward Shortliffe gave a talk at AIME in which he characterized artificial intelligence (AI) in medicine as being in its "adolescence" (Shortliffe EH. The adolescence of AI in medicine: will the field come of age in the '90s? Artificial Intelligence in Medicine 1993;5:93—106). In this article, the discussants reflect on medical AI research during the subsequent years and characterize the maturity and influence that has been achieved to date. Participants focus on their personal areas of expertise, ranging from clinical decision-making, reasoning under uncertainty, and knowledge representation to systems integration, translational bioinformatics, and cognitive issues in both the modeling of expertise and the creation of acceptable systems.

1. Introduction

The earliest work in medical artificial intelligence (AI) dates to the early 1970s, when the field of AI was about 15 years old (the phrase "artificial intelligence" had been first coined at a famous Dartmouth College conference in 1956 [1]). Early AI in medicine (AIM) researchers had discovered the applicability of AI methods to life sciences, most visibly in the Dendral experiments [2] of the late 1960s and early 1970s, which brought together computer scientists (e.g., Edward Feigenbaum), chemists (e.g., Carl Djerassi), geneticists (e.g., Joshua Lederberg), and philosophers of science (e.g., Bruce Buchanan) in collaborative work that demonstrated the ability to represent and utilize expert knowledge in symbolic form.

There was an explosive interest in biomedical applications of AI during the 1970s, catalyzed in part by the creation of the SUMEX-AIM Computing Resource [3] at Stanford University, and a sister facility at Rutgers University, which took advantage of the nascent ARPANET to make computing cycles available to a national (and eventually international) community of researchers applying AI methods to problems in biology and medicine. Several early AIM systems including Internist-1 [4], CASNET [5], and MYCIN [6], were developed using these shared national resources, supported by the Division of Research Resources at the National Institutes of Health.

The general AI research community was fascinated by the applications being developed in the medical world, noting that significant new AI methods were emerging as AIM researchers struggled with challenging biomedical problems. In fact, by 1978, the leading journal in the field (*Artificial Intelligence*, Elsevier, Amsterdam) had devoted a special issue [7] solely to AIM research papers. Over the next decade, the community continued to grow, and with the formation of the American Association for Artificial Intelligence in 1980, a special subgroup on medical applications (AAAI-M) was created.

It was against this background that Ted Shortliffe was asked to address the June 1991 conference of the organization that had become known as Artificial Intelligence in Medicine Europe (AIME), held in Maastricht, The Netherlands. By that time the field was in the midst of "AI winter" [1], although the introduction of personal computers and high-performance workstations was enabling new types of AIM research and new models for technology dissemination. In that talk, he attempted to look back on the progress of AI in medicine to date, and to anticipate the major challenges for the decade ahead. A paper based on that talk was later published in *Artificial Intelligence in Medicine* [8]. Thus, when our panel of senior AIM researchers was constituted for the AIME conference in Amsterdam in July 2007, we chose to reflect on some of the assessments and predictions that had arisen from Shortliffe's presentation some 16 years earlier. This article summarizes those remarks from the AIME 2007 panel.

2. Comments by Edward H. Shortliffe

There were three key points to my 1991 presentation, all of which I believe are equally pertinent today. First, I claimed that AI in medicine cannot be set off from the rest of biomedical informatics, nor from the world of health planning and policy. Realistic expectations of the field's influence in health care and biomedical sciences require that we draw upon AI as only one of the many methodological domains from which good and necessary ideas can be derived. This amounts to an argument that AIM researchers need to be willing to draw on other fields of computer science and informatics as necessary, ranging from principled approaches to human—computer interaction or database theory to numerical analysis and advanced statistics. It is the ultimate applications, and their value in biomedicine, that must drive our work, and this may mean being eclectic and as oriented to policy and sociocultural realities as we are to the technical underpinnings of a medical AI application.

Second, we need to realize that the practical influence of AIM in real-world settings will depend on the development of integrated environments that allow the merging of knowledge-based tools with other applications. The notion of stand-alone consultation systems had been well debunked by the late 1980s [9], and thus we must be looking for ways to combine "backend" AI notions with such ubiquitous systems as electronic medical records, provider order-entry systems, results reporting systems, e-prescribing systems, or (on the biological side) tools for genomic/proteomic data management and analysis. This reality creates challenges for researchers, because the implication is that we need breadth of knowledge and collaborations that go beyond our immediate AI roots.

Third, our ability to influence the delivery of health care, or the quality of biomedical research, will depend on vision and resources from leaders who understand that medical practice, and biomedical research, are inherently information-management tasks—and must accordingly be tackled and supported as such. To this day I find it remarkable how many leaders continue to view their IT investments as discretionary, and do not realize the key

strategic role that clinical and biological computing infrastructure has on quality, error reduction, efficiency, and even cost savings. Biomedical informatics researchers, including those who work in the AIM area, must learn to be effective missionaries, presenting their case effectively to key decision makers in ways that gradually effect the cultural change that will be necessary for the full impact of our technologies to be felt.

In the 1991 talk and subsequent article, I also laid out three key challenges for the field. First it seemed clear to me then, as it does now, that we need more professionals who are broadly educated regarding the interdisciplinary nature of biomedical informatics, including its AIM component. Having learned that there are too few individuals with focused training at the intersection of biomedicine and computer science (and the other informatics component sciences, such as decision science, cognitive science, and information science), we have tried to gear up with new formal and informal programs offering graduate degrees and certificate training, as well as continuing education courses for a variety of health professionals (physicians, nurses, dentists, pharmacists, etc.). But with growing demands for these interdisciplinary skills, there are still too few people capable of working effectively at the intersection, even in academic or industrial research roles, and we need more departments, more support for training positions, and more buy-in from institutions that instinctively eschew the formation of new academic units.

Second, in 1991 we identified the need to develop national and international biomedical networking infrastructures for communication, data exchange, and information retrieval. We were just beginning to embark on the "democratization" of the Internet in 1991, with the earliest forays into web concepts underway. Today, 16 years later, we see remarkable progress in this area, with growing dependence on electronic communication, e-publishing, and online collaborative activities based on Web 2.0 and related concepts. There is still much work to be done, but I believe that the community has met the challenge from the early 1990s and continues to expand its capabilities and activities in this important area.

Third, we identified the need for credible international standards for communications, data and knowledge exchange. Again there has been a great deal of work in this area in the intervening years, not the least of which has been a broadened acceptance of the importance of standards adoption to support system integration (including, of course, the integration of AIM decision support with biomedical and clinical data systems of various sorts). Certain standards have been widely adopted, such as HL7 for data exchange (http://www.hl7.org), but there continues to be much work to be done in this key area.

Against the backdrop of these issues from 1991, our panel at AIME-2007 encouraged me to consider issues such as (a) How has the field advanced?, (b) In what ways, and to what extent, has the field had a direct influence on clinical medicine or other biomedical fields?, and (c) How well is the field being supported (by funding agencies, by academic and research organizations, and by our biomedical or computer-science colleagues)? What follows is a summary of some of those observations.

At first blush, AI in medicine is alive and well, with AIM researchers using a wide array of AI-inspired methods to tackle a broad range of important clinical and biological problems (see Table 1). However, although AI issues are ubiquitous in biomedicine, many people who are doing AIM research do not label it as AI. What was once a catchy, respected label has lost much of its luster—a casualty of AI winter and the general societal sense that AI had somehow overpromised and failed to deliver. Yet I see AI broadly represented in the biomedical informatics field, in areas such as knowledge representation and ontology development, terminology and semantic modeling of domains, decision support and reasoning under uncertainty, model-based image processing, and many others.

Table 1 Topics and themes at AIME 2007.

Computer-based knowledge generation	Data and knowledge representation
Clinical data mining	Knowledge-based health care
Probabilistic and Bayesian analysis	Feature selection/reduction
Visualization	Classification and filtering
Information retrieval	Agent-based systems
Temporal data mining	Machine learning
Knowledge discovery in databases	Text processing
Natural language processing	Ontologies
Decision support systems	Image processing
Pattern recognition	Clinical guidelines
Workflow	

Ironically, whereas many researchers in these areas do not call their work AI, even though the historical and methodological roots are clearly in the AI area, those commercial systems that claim they offer "artificial intelligence" almost never do—at least by the technical standards that we would tend to use in determining whether a piece of work draws on AI methods. With the diffusion of AI research throughout biomedical informatics, the biennial AIME conference, and the international journal *Artificial Intelligence in Medicine* stand out as the two remaining forces for defining and recognizing AI in medicine as a subfield of biomedical informatics and computer science.

Another observation is the fascinating transition to an emphasis on guideline-based decision support. This parallels what is happening in clinical medicine, where clinical guidelines have been introduced as a proposed way to reduce unjustified clinical variability among providers and to enhance error reduction efforts. Clinical guidelines are sometimes viewed simply as a resurgence of interest in the "clinical algorithm" notions that were popular in the late 1960s and early 1970s. Guidelines are often accompanied by algorithms or flow charts that provide declarative information about how to diagnose, work up, or treat patients with certain conditions or complaints. Implementing guidelines is accordingly quite different from the classical patient-specific decision-support efforts that had emerged for diagnosis and therapy planning from researchers in the AIM community. Thus the shift to guideline issues has in part been at the expense of ongoing work on statistical aspects of medical diagnosis, Bayesian belief networks, ontology development to support reasoning under uncertainty, or complex planning approaches applied in clinical domains. This is not to say that guideline work has been simple. As always, the devil is in the details, and researchers on clinical guidelines have uncovered important challenges in knowledge representation, standardization, integration, and presentation of advice.

Meanwhile there has been impressive progress in several AIM research areas: knowledge representation (and the associated tools, including the remarkable worldwide impact of Protégé, itself a product of AIM research at Stanford [10]), machine learning and data mining for knowledge discovery (including in text databases), and temporal representation and reasoning (to mention only a few). Yet progress has been slow, albeit real, in the adoption of key standards needed for integration and knowledge sharing (e.g., controlled terminologies and their semantic structuring, standards for representing clinical decision logic to enhance its sharability, and incorporation of AI concepts into robust, well-accepted clinical products). Many of the barriers to progress in these latter areas have been political, fiscal, or cultural rather than purely technical.

A particularly welcome transition has been the gradual tendency of traditional computer science departments to embrace biomedical applications work. Two decades ago, it was a significant barrier to computer scientists' careers if they were viewed as being "too applied" in any single domain. Today, recognizing the stimulation of cutting edge computer science that can come from work on biomedical applications (and the new sources of grant funding that accompany such work), academic computer science has begun to embrace biomedical applications as valid areas of emphasis for computer science faculty members. This has been especially true for faculty who work in the bioinformatics domain, many of whom draw on artificial intelligence methods in their work.

My summary assessment, then, is that the AI in medicine field is robust, albeit less visible than it was in AI's heyday. There is clear evidence of progress, and a community of talented researchers that would benefit from more growth in numbers and in research grant funding. What began largely in the United States in the late 1960s and early 1970s is now a worldwide field, with important contributions from around the globe, but with special acknowledgement to our European colleagues who continue to lead us with their biennial AIME conferences and the highly regarded international journal *Artificial Intelligence in Medicine*.

3. Comments by Vimla L. Patel

It was Mario Stefanelli, and the AIME program committee, who asked me to present an address at the 1991 conference at which Ted Shortliffe gave his "Adolescence of AI in medicine" speech. I was asked to discuss studies in human intelligence (thinking and reasoning) and their relationship to medical artificial intelligence [11]. Today I would like to ask whether, in the evolution of AIM research, we have forgotten about the human mind as we perform our work. Since the early days of AI, there has been a debate about the extent to which people who build AI systems should be modeling how human beings think and solve problems. The debate is exemplified by two nicknames for AI researchers, those who are the "scruffies" (pragmatists in the sense that a system's performance on tasks is more important to them than whether the system solves problems as human beings would) and the "neats" (formalists, theoreticians, or psychologists who argue that true AI requires modeling and insights into human intelligence). In today's

world, we need both types of people, or people who effectively move between the extremes, since the two approaches serve different purposes in the AI in medicine community.

Issues that concerned the AIM community in the 1980s were different from those in the current decade. In the past, there was an emphasis on the development of stand-alone AI systems, using computer science/engineering approaches, aiming for accurate and reliable decision-making performance, regardless of whether the system solved problems in the same way that human experts do. Thus our AIM traditions have tended to be derived from the "scruffy" branch of AI. Today we have moved away from these stand-alone systems [9] to the development of integrated systems in clinical environments, interfacing with medical record and order-entry systems, thereby using a wide variety of computational methods. Given that there is a difference in the way knowledge is organized in performance-oriented systems from the way in which that same knowledge is organized in the minds of human beings [12], there is also generally no attempt to model human reasoning processes. There is also a greater emphasis now on clinical workflow and socio-technical considerations among the design issues for the AIM community.

Yet one of the lessons of informatics work in recent decades has been that even the performance-oriented "scruffies" need to build systems with insights into the human mind if they are going to achieve the outcomes desired. System users are, after all, human beings, and their modes of reasoning and mental models of domains will determine how they utilize and respond to advice or guidance provided through AIM systems. As in most domains, there has always been a gulf between technologic artifacts and end users. Since medical practice is a human endeavor, there is a need for bridging disciplines to enable clinicians to benefit from rapid technologic advances. This in turn necessitates a broadening of disciplinary boundaries to consider cognitive and social factors related to the design and use of technology. A large number of health information technologies fail. Our evaluations today tell us that most of these failures are due not to flawed technology, but rather to the lack of systematic considerations of human issues in the design and implementation processes. In other words, designing and implementing these systems is not as much an IT project as a human-centered computing effort, dependent on topics such as usability, workflow, organizational change, and process reengineering.

All technologies mediate human performance. Technologies, whether they be computer-based or in some other form, transform the ways individuals and groups behave. They do not merely augment, enhance or expedite performance, although a given technology may do all of these things. The influence of technology is not best measured quantitatively since it is often qualitative in nature. Technology, tools, and artifacts not only enhance people's ability to perform tasks but also change the way in which they do so. In cognitive science, this ubiquitous phenomenon is called the representational effect, which refers to the phenomenon that different representations of a common abstract structure can generate dramatically different representational efficiencies, task complexities, and behavioural outcomes. These are the current challenges that we in the AIM community face and will require some understanding of the cognitive factors that influence design [13].

The importance of cognitive factors that determine how human beings comprehend information, solve problems, and make decisions cannot be overstated. Investigations into the process of medical reasoning have been one area where advances in cognitive science have made significant contributions to AI. In particular, reasoning in a medical context involving high throughput and high degree of uncertainty (such as critical care environments), compounded with constraints imposed by resource availability, leads to increased use of heuristic strategies. The utility of heuristics lies in limiting the extent of purposeful search through data sets. By reducing redundancy, such strategies have substantial practical value. A significant part of a physician's cognitive effort is properly selecting and utilizing pertinent heuristic approaches. However, the use of heuristics introduces considerable bias in medical reasoning, often resulting in a number of conceptual and procedural errors. These include misconceptions about laws governing probability, flawed instantiation of general rules to a specific patient at the point of care, misunderstanding prior probabilities, as well as falsely validating a hypothesis. Much of physicians' reasoning is inductive, with attached probability. Human thought is fallible and we cannot appreciate the fallibility of our thinking unless we draw on an understanding of how physicians' thinking processes operate in the real working environment. Such level of understanding will be necessary as AIM research further evolves [14].

Finally, given the current trend in managing medical errors, the future work in AI that relates to human beings working within a socio-cognitive context becomes even more salient. Early research on clinical errors included studies of human reliability in the process, with the human component being considered as just one more element in the system,

viewed as more or less equivalent to the technical components. Just as technical safety is improved through the reduction of technical breakdowns, it seemed intuitive that one could improve safety through the elimination of human errors. However, we now know that mistakes are inevitable and cognitively useful phenomena that cannot be totally eliminated. This raises an issue of having suitable goals for management and recognition of these errors plus proper responses of the systems (and individual) when they occur. These issues require research so that we can better understand the boundaries of human errors and risk taking and apply these lessons in the design of safe systems that are resilient [15]. Such resiliency should become a key element in the design and implementation of future AIM systems.

4. Comments by Mario Stefanelli

I would like to direct my remarks to the socio-organizational approach in the development of health care systems. Although machines are not yet showing general intelligent behaviours, AI is nowadays much more than a promise. AI has profoundly and paradigmatically changed computer science by introducing the separation between knowledge representation and inference. Rather interestingly, albeit without spotlights, the major achievements of AI are going to be reached in the current days. AI is now part of current software technology solutions in the areas of logistics, data mining and image processing. Moreover, AI is boosting discovery in genetics and molecular medicine, by providing machine learning algorithms, knowledge representation formalisms, biomedical ontologies, and natural language processing tools [16,17].

As far as medicine is concerned, knowledge management (KM) is one of the most interesting AI fields [18]. The goal of KM is to improve organizational performance by enabling individuals to capture, share and apply their collective knowledge to make optimal "decisions in real time". Such approach is completely coherent with the current vision of the role of health care organizations (HCOs) in the 21st century [19]. The new main goals of HCO are safety, efficiency and effectiveness, centrality of the patient, continuity of care, care quality and access equity. As a consequence, medical KM and health care process management are crucial to achieve the desired quality. The first goal of KM in medicine is therefore the definition of effective tools for supporting communication between all the actors involved in patients' care. Such communication aims at developing shared meanings of what is happening outside and inside the HCO in order to plan and make decisions. Shared interpretations are needed to define the organization intent or vision about what new knowledge and capabilities the organization needs to develop.

Managing knowledge in HCOs, however, does not merely focus on improving the availability of instruments for improving communication. On the contrary, KM aims at transforming information into actions; this transformation is the basic premise to knowledge creation, which amplifies the knowledge acquired or discovered by individuals and makes it available through the organization [20]. From an organizational viewpoint [21,22], knowledge creation is the result of a *social interaction* between two fundamental types of knowledge, tacit knowledge and explicit knowledge [23]. Tacit knowledge is characterized by the fact that it is personal, context specific and therefore hard to formalize and communicate. Explicit knowledge is transmittable through any formal or systematic representation language, from a text written in natural language to a (more or less) complex computer-based formalism. The transformation between explicit and tacit knowledge process has been called *knowledge conversion*. Four different modes of knowledge conversion have been postulated: *socialization, externalization, combination, and internalization. Socialization* is the process of sharing experiences that creates tacit knowledge as shared mental models and technical skills. Newly trained physicians and nurses successfully learn by imitating the behaviours of experienced practitioners.

Externalization is the process of conversion of tacit into explicit knowledge through the development of models, protocols or guidelines. *Combination* is the process of recombining or reconfiguring bodies of existing explicit knowledge that leads to the creation of new explicit knowledge. *Internalization* is the process of learning by repetitively doing a task applying the explicit knowledge so that the achieved outcomes become absorbed as new tacit knowledge of the individual. All four phases may effectively be supported relying on AI methods and tools. Intelligent data analysis and data mining support the extraction of patterns and regularities from the process data collected during HCO activities [24]. The transformation of such patterns into explicit knowledge requires knowledge representation formalisms and tools. Guidelines, protocols and decision models are derived as the final part of the *externalization* activity. Once knowledge is acquired and formalized, it is effectively exploited thanks to knowledge management methods and tools [25,26]. The high level *combination* of information and processes may lead
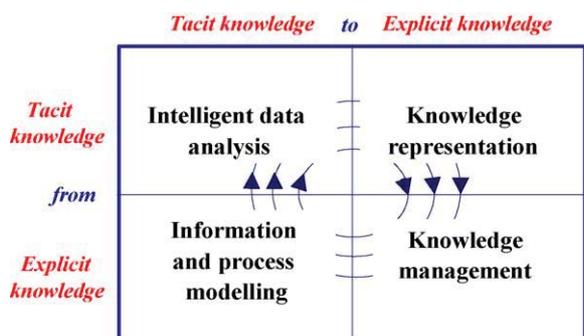
Figure 1 The knowledge cycle implemented with AI methods and tools.

to the definition of new knowledge that, once *internalized* and diffused with *socialization*, is mirrored by the actions of HC providers, collected by process data. The entire knowledge cycle is thus implemented with AI technologies (see Fig. 1).

Knowledge creation is one of the basic components of *organizational learning*, which refers to the skills and processes of creating new knowledge by doing within a working organization [27]. To reach this goal, medical knowledge, organizational knowledge and clinical information must be effectively represented and integrated to assist patient and citizen care. From a technological viewpoint, KM can be implemented within a careflow management system (CfMS) [28,29] or a service-flow management system [30]. A CfMS acts as a component of the health information system (HIS) to completely define, create and manage the execution of careflows. A CfMS involves dedicated procedures through which administrative and supervisory tasks, such as sharing documents and information or assigning commitment for task execution, are passed from a care giver to another one according to a process definition. This consists of a network of activities and their relationships, criteria to indicate the start and termination of the process, and information about the individual activities.

CfMS are now implemented in running HIS. For example, within the stroke active guideline evaluation (Stage) project, which involves 27 neurological units in Italy, a CfMS was implemented at the stroke unit ''IRCCS Mondino'' in Pavia. Currently about 250 patients have been treated with the CfMS and its effectiveness has been shown [31].

A service-flow management system applies organizational learning concepts to chronic and sub-acute patient care. Several models of distributed care services have been recently defined. They range from case management, intensive case management, assertive community treatment and community-based practices. The latter model seems

particularly suited for implementing socio-technical learning strategies [32]. Community-based research attempts to improve academic research by valuing the contribution that community groups make in the development of knowledge. To this end, researchers and practitioners share goals, problems and interests on specific issues, solve new problems using their knowledge and find innovative solutions for new problems. This requires the development of a ''distributed'' team identity by facilitating the conversion of implicit into explicit knowledge and vice versa. As an example, the Italian Amyloidosis Network is implementing community-based research strategies to deal with amyloidosis, a rare severe disease which refers to a variety of conditions in which amyloid proteins are abnormally deposited in organs and/or tissues. The Italian network for amyloidosis involves 62 biomedical centers and the diagnostic and therapeutic guidelines are approved each year during the annual society meeting [33]. A national portal with all information and contacts related to amyloidosis has been implemented. The goal of the portal is to provide all participating communities to share the latest development of research, the latest treatment protocols and a shared health care record management system, based on standard terminologies and domain specific ontologies.

The number of successes of AI in medicine is likely to grow in the near future. On the opposite side of the general perception that AI is in its winter time, we fully agree with Rodney Brooks [1,34]:

''there's this stupid myth out there that AI has failed, but AI is around you every second of the day.''

The new generation of health care information systems and the current biomedical informatics research are constantly proving the truth of this sentence.

5. Comments by Peter Szolovits

This panel has presented a great opportunity to review the past 15 years of progress and changes since Shortliffe's influential talk and publication regarding ''AIM's adolescence''. My own take on the major changes that have happened over that period is that AI in medicine is viewed today much less as a separate field and more as an essential component of biomedical informatics and one of the methodologies that can help to solve problems in health care. Although this change was already occurring in the early 1990s and is foreshadowed by Shortliffe's article, I think the field has continued to generalize and to merge with larger concerns.

Today's "systems" thinking about health care focuses not only on the classical interactions between patients and providers but takes into account larger-scale organizations and cycles. We can, of course, still focus on short-term interactions such as those that occur during an office visit or hospitalization, or even during shorter-term interactions such as those arising during a surgical procedure or intensive care. In addition, however, we now also pay attention to the continuous and repetitive nature of clinical care, much of which occurs in the community rather than in a hospital, and involves sources of knowledge coming from family members, groups of patients suffering from similar conditions, various home-care programs, and especially web-enabled searches and remote communications. In addition, we are coming to recognize that the health care system is not a static background for our efforts but must learn from its own experiences and strive to implement continuous process improvements that can significantly improve health outcomes while somewhat keeping in check the inexorable growth of health care spending. If, as most believe, it is true that

$$\text{Phenotype} = f(\text{Genotype}, \text{Environment})$$

and that our ability to exploit the "new biology" of high-throughput genetic measurements depends on an ability to match these to phenotypic data, then we must view the clinical record of "natural experiments" (diseases) as a most valuable source of data for biomedical research [35].

We also recognize that much of what ails health care is not innately technical at its roots. Many problems such as inequities in care, lack of insurance, unsupported practice variations, poor compliance with established guidelines, poor feedback on long-term outcomes of care, etc., require improvements in policy and management more than in technology. Nevertheless, technology, including AIM technology, can provide new options to help address these larger problems.

AIM research faces numerous interesting challenges, of which I will highlight just four: (1) better data capture and handling, (2) improved design, modeling and assistance for workflows, (3) reliable methods for reassuring patients in their concerns for confidentiality, and (4) better modeling techniques. These pose genuine basic research problems of the sort described in Shortliffe's earlier article, and therefore cannot be expected to yield short-term solutions to the problems of health care. They do, however, lay out a partial set of research goals that will, if successfully met, significantly improve health care.

5.1. Data

Much of the early AIM research focused on capturing the expertise of human experts in sophisticated computer programs. Today I joke with students that in those days we thought we knew a lot, but had little or no actual data. Today we are inundated with data, but have correspondingly devalued expertise. Yet despite the huge volume of data that are now routinely collected in health care, much of it remains incomplete or inaccurate in critical ways. Papers continue to document that notes of patient encounters sometimes misrecord even basic facts such as the chief complaint, and often get wrong details such as the patient's medical history or medications being taken. Lack of commonly accepted terminologies and ontologies makes exchange and interoperation of even well-recorded information difficult. Although we have moved beyond the days when lab instruments would print measurement results on paper and then discard the digital data, we still routinely see nurses and technicians transcribing data from one system to another because of standards for data exchange that are either lacking, poorly designed or poorly implemented. The vision of all instruments interoperating for seamless data exchange is an old one, but far from having been achieved. Whether through stricter standardization or more intelligent interfaces, this needs to be solved. Wireless and portable devices promise to support more convenient interactions, but will require good support for reliability and semantic reconciliation of conflicting records as well as great data exchange capabilities. Intelligent environments could combine speech understanding, computer vision systems, gesture tracking, comprehensive recording and models of how people interact to capture primary encounter data that is now often only recorded (incorrectly) from memory. Better natural language processing capabilities could help unlock the value now buried in narrative records whose content is opaque to traditional computer systems. Error models that take into account the typical sources of noise and corruption in data capture could help automatically "clean" data about clinical care to support both more robust assistance for the care process and better research data.

5.2. Workflow

Systems, whether based on AIM or other methods, must operate in conjunction with human practitioners. Therefore, they must model what those practitioners do, what information they need, and when the disruption caused by the system intervening

is more than offset by the value of its information. We read that many medical errors are due to omission rather than commission. This suggests that systems working in the background should be continuously monitoring care for every patient and checking to see if expectations are being met. For example, one could design a workflow system that requires inclusion, with every action, of a scheduled future step that verifies that the initially planned action was in fact performed and that its outcome was consistent with what was anticipated. Some systems already notify the doctor responsible for a patient's care of highly abnormal lab values, and then escalate the alert to others if they see no response [36]. Such a strategy should apply to all clinical actions, ranging from assuring that scheduled X-rays are actually taken to providing growingly insistent reminders that a child's check-ups or immunization schedule are not being met. Further, we know from Homer Warner's HELP system of 35 years ago that it is possible to incorporate decision support at every step of clinical care [37]. We need to make this part of routine practice, and to overcome impediments to its adoption and use.

5.3. Confidentiality

Much latent resistance to fully electronic tracking of health care arises from people's unfortunately correct beliefs that aggregation of vast amounts of sensitive health care data increase vulnerability to massive disclosures [38]. We need only read the daily newspapers to hear of institutional errors that release personal data on millions of people in a single incident. Thus far, most of these massive releases have threatened identity theft rather than medical disclosures, but those incidents have also occurred on a smaller scale and such vulnerabilities are widely recognized. To some extent, anxiety about such releases of information could be mitigated by universal guarantees of access to health care and non-discrimination in insurance based on patients' existing conditions. That would still leave embarrassment and a sense of violation of personal privacy as strong motivators for concern. Some technical advances that could help with these problems would be improved ways to establish identity, perhaps through distributed and local schemes that avoid the need for universal and irrepudiable identifiers. We need convenient and secure means of authentication, better than today's username/password combinations, whether by personal smart cards, biometrics, or some clever exploitation of already-existing technologies that can serve to identify people, such as their credit cards or cellular phones. We could also do a better job of decoupling

individuality (the ability of systems to determine that heterogeneous data all belong to the same person) from identity (who that person actually is). Such an approach could allow much of the quality and business analysis of health care to proceed and much of the research data to be used with much lower risk of divulging data about recognizable individuals [39]. A longer-term research challenge, perhaps unachievable, is to create data sets that naturally decay but without the need for cumbersome digital rights management infrastructures.

5.4. Modeling

I have noted the dramatically increased availability of large collections of data, even in routine clinical settings. New measurement techniques such as microarrays that simultaneously determine hundreds of thousands of DNA, RNA and protein levels and methods that determine a half million SNPs or, soon, an individual's entire genetic sequence, cannot be treated as simply a huge number of additional "findings" in traditional diagnostic or therapeutic reasoning systems. Simply to make sense of such volumes of data will require advanced AI methods that can automate their analysis. As a community, we have already adopted traditional statistical and more novel data mining and machine learning approaches to deal with this wealth of data. Unfortunately, these techniques tend to discover relatively simple relationships in data and have not yet demonstrated the ability to discover complex causal chains of relationships that underlie our human understanding of everything from molecular biology to the complex multi-organism and environmental factors in the epidemiology of diseases such as malaria. Human expertise, developed over centuries of experience and experimentation, cannot be discarded in the hope that it will all be re-discovered (more accurately) by analyzing data. For example, I do not know of any automated methods that would be able, from terabytes of recorded intensive care unit monitoring data, to discover even elementary facts such as that blood circulates because it is pumped by the heart. Therefore, I think it is a great challenge to build better modeling tools that permit the integration of human expertise (recognizing its fallibility) with machine learning methods that exploit a huge variety of available data to formulate and test hypotheses about how the human organism "works" in health and illness.

Challenges for AIM remain vital and exciting. However, we recognize that our crisis in health care demands an ever-broader set of disciplines to create integrated solutions. AI in general has come closer over the years to statistics and operations research,

linguistics, communications engineering, theoretical computer science, computer systems architecture, brain and cognitive science, etc. Fundamental research progress in medicine depends on biochemistry, molecular biology, physiology and a host of medical specialties. Improvements in health care demand coordination with economics, management, industrial engineering and policy. These trends demand that we educate our students more broadly and that we continue the laudable tradition of interdisciplinary projects in AIM.

6. Comments by Michael Berthold

Before investigating the progress and ongoing challenges of AI approaches in medicine, it may be helpful to categorize the type of science going on in this research area.

An often used categorization of scientific research concentrates on three phases:

(1) *Collection*: The initial effort relates to gathering of data about the problems at hand. No clear knowledge about underlying regularities or systems is available nor do researchers know much about the domains of the data of interest.
(2) *Systematization*: The collected data is organized better and models are being built to predict certain properties—most of these models, however, are build without a clear knowledge about the underlying system. The system that has generated the original data still is very much a black box.
(3) *Formalization*: A better understanding of the underlying system has been achieved and theories can be formed and validated through targeted, systematic experimentation.

In sharp contrast to many other scientific disciplines, research in medical domains is still very much stuck in the early phases. Some isolated knowledge fragments are available about medical systems but no fine-grained, global model exists. One could argue that some of this research has reached phase 2, Systematization. However, especially in pharmaceutical drug development, experiments often end up creating data without a clear idea about its use. In fact, much of this data will hardly ever be read again. In these areas, research still mostly focuses on data collection with the sometimes rather vague hope to stumble across discoveries which will ultimately lead to new medications. One of the key problems in these areas is the increasing ability to generate the data and the much slower advent

in methods to deal with the resulting, gigantic data repositories. Converting these heaps of data into information and ultimately knowledge is still one of the most pressing needs in biomedical research.

The interesting question is: do current AI methods support this type of research scenario? Most applications of current AI methods are either focussing on unsupervised approaches which try to identify structure in data by clustering or similar approaches or by more or less complex means to present visualizations or summaries of the data. Supervised approaches on the other hand, focus on either finding patterns of very particular, predefined type (e.g., association rules, subgroups) or build predictive models. These models can be black boxes (e.g., artificial neural networks) or interpretable models (e.g., decision trees or rules). No matter which of these techniques is used, the underlying model families or similarity metric push a strong bias into the analytic process. Hence current applications of AI methods mainly focus on answering rather well-posed questions. One could argue that this type of problem solving approach was appropriate a decade ago when data resources were considerably smaller and one could hope to make sense out of them using such restricted approaches. However, in recent years data has far outgrown our ability to analyse them and new, more powerful and versatile methods are needed. One could even say that the increasing amount of data keeps pushing this area of scientific research back towards phase 1, the sheer collection of new data!

Therefore new methods are needed which allow to uncover the unexpected, allow the user to interactively form new, initially often confusing hypotheses and assist them in discovering truly new insights—ultimately leading to an understanding of the underlying system. One could describe such a system as an "external AI", assisting the user in what she can do best: quickly sorting out the useless aspects from the currently interesting information pieces, probing and discarding potentially interesting connections and associations and narrowing down on the gems hidden in the vast amounts of available data. Such a system should not attempt to do the discovery job for users—instead it needs to support them by giving associative, intuitive access to everything the system has access to: unstructured and semi-structured data all the way to humanly annotated pieces of expert knowledge. Hence we need to be developing *discovery-support systems* rather than automated discovery systems [40].

7. Concluding remarks by Riccardo Bellazzi and Ameen Abu-Hanna (AIME 2007 Program Chairs)

Over the last few years medicine's identity as a data-rich quantitative field has become much more appreciated—especially with the use of electronic data capture and data management systems for both clinical care and biomedical research. The abundance of data is strongly accelerating the process of transforming medicine from art to science and is providing new ways to carry on biomedical research. Data-driven studies are more and more frequent, looking at the discovery of new, unexpected knowledge as the "holy grail" buried in the data. Image-based and molecular-based diagnoses are becoming standard ways to assess a patient's disease more precisely. Guidelines and protocols are disseminated to standardize a patient's treatment. Finally, health care organizations are now considered complex companies, which may be studied from a business perspective. It is against this background that the panelists of AIME 2007 offered their thoughts on the "coming age of AI in medicine". The coming of age of a person is the transition from adolescence to adulthood. AIM is approaching 40 years of age, but for scientific disciplines it is hard to discern whether and when such a transition takes place, partly due to the lack of standard criteria to establish this transition. For example when AIM was about 25 years old Coiera argued that AIM was not yet being successful—if success is judged as making an impact on the practice of medicine [41]. Haux is of the opinion that the field of medical informatics as a whole is still relatively young but that it has had an impact on the quality and efficiency of health care and on biomedical research [42]. Regardless of the specific criteria one chooses to use to mark transitions on the maturity scale, the authors of this paper are of the opinion that:

- AIM draws upon many disciplines. Computer science, the background perhaps characterizing most AIM researchers, is only one such discipline—albeit an important one. AIM research is continuously widening its scope and there is a need for more people with background in the disciplines at the intersection defining AIM and its parent field of biomedical informatics.
- AIM methods are becoming more and more integrated within other applications. Paradoxically, this diminishing of explicit visibility is a sign of the success of the AIM program.
- We have come a long way in creating and/or utilizing the information and communication infrastructures needed for the AIM applications, but there are still challenges and barriers such as defining communication and data sharing standards, having access to data which are complete and coded according to agreed upon terminological systems.
- There is a move from "does the system work?" to "does the system also help?" This implies implementing and testing AIM-based solutions within the environment of clinical practice. Sophisticated evaluation designs are being used to assess impact on both process and patient outcomes.
- The staggering amounts of data generated and collected in the biomedical field gave impetus to research on (statistical) machine learning that tries to make sense of these data. There is still a long way to go in order to find causal relations in the data, but an equally useful purpose is to create tools that act as discovery-support systems facilitating the work of the human interpreter.
- Evidence-based medicine has fostered the implementation of guidelines and protocols; AI approaches have been demonstrated to be useful for building and checking them, and workflow systems appear to be the proper way to apply guidelines in dynamic environments.
- There is, however, a strong need to apply AI tools and methods besides data and guidelines. Scientists working in a "data-driven world" are recognizing the strong risk of concentrating on data gathering and analysis alone. Poor systematization and poor formalization of knowledge may result in accumulating data without knowledge extraction and/or without knowledge exploitation. On the other hand, a "guideline-based world" may strongly suffer from a lack of flexibility; dogmatic guidelines may constrain efforts to deal effectively with tailored decision-making and may overlook the importance of research on complex planning, decision-making under uncertainty, and individual risk management.

The challenge for AI in the next years will be to ground the current research scenario in its AI roots. As recognized by all panelists, the representation of all kinds of knowledge and high-level systems modeling are important topics for basic AI in medicine research. Moreover, the effective exploitation of knowledge in building decision-making tools and in extracting information from the data is also very important. The field of intelligent data analysis seems relevant in this regard [43,44]. Since AI in medicine applications today span from molecular medicine to organizational modeling, the role of modeling human reasoning and cognitive science must be re-evaluated. Modeling and reasoning will play a significant role as

we strive to build successful systems and to deal with their impact on how people, from research groups to healthcare teams, perform their work. Last but not least, strong interdisciplinary education programs should be further fostered, to improve the quality of researchers and practitioners and to help the dissemination of AI methods and principles in the biomedical informatics community.

The AI in medicine leaders participating in the AIME 2007 panel have argued that AI in medicine is coming of age as a discipline. An assessment of its current status has been helpful as we seek to propose future directions to improve not only biomedical informatics but also biomedical research more generally

Acknowledgement

References

[1] History of artificial intelligence. http://en.wikipedia.org/wiki/History_of_artificial_intelligence. (Accessed June 1, 2008).

[2] Lindsay RK, Buchanan BG, Feigenbaum EA, Lederberg J. Applications of artificial intelligence for organic chemistry: the DENDRAL Project. New York: McGraw-Hill; 1980.

[3] Freiherr G. The seeds of artificial intelligence: SUMEX-AIM (1980). U.S. G.P.O; DHEW publication no.(NIH) 80-2071. Washington, D.C.; U.S. Dept. of Health, Education, and Welfare, Public Health Service, National Institutes of Health; 1980.

[4] Miller RA, Pople HE, Myers JD. Internist-1: an experimental computer-based diagnostic consultant for general internal medicine. New England Journal of Medicine 1982;307(8): 468—76.

[5] Weiss SM, Kulikowski CA, Amarel S, Safir A. A model-based method for computer-aided medical decision making. Artificial Intelligence 1978;11:145—72.

[6] Shortliffe EH. Computer-based medical consultations: MYCIN. New York: Elsevier; 1976.

[7] Sridharan NS. Guest editorial. Artificial Intelligence 1978:11 (1—2);1—4.

[8] Shortliffe EH. The adolescence of AI in medicine: will the field come of age in the 90s? Artificial Intelligence in Medicine 1993;5:93—106.

[9] Miller RA, Maserie F. The demise of the Greek oracle model for medical diagnosis systems. Methods of Information in Medicine 1990;29:1—2.

[10] Noy NF, Crubezy M, Fergerson RW, Knublauch H, Tu SW, Vendetti J, et al. Protégé-2000: an open-source ontology-development and knowledge-acquisition environment. In: Musen MA, Friedman CP, Teich JM, editors. Proceedings of the 27th annual symposium of the American Medical Informatics Association AMIA 2003: biomedical and health informatics: from foundations to applications. Bethesda: American Medical Informatics Association; 2003. p. 953.

[11] Patel VL, Groen GJ. Real versus artificial expertise: the development of cognitive models of clinical reasoning. In: Stefanelli M, Hasman A, Fieschi M, Talmon J, editors. Proceedings of the 3rd conference on artificial intelligence in medicine AIME 91. Berlin: Springer-Verlag; 1991. p. 25—37.

[12] Patel VL, Ramoni M. Cognitive models of directional inference in expert medical reasoning. In: Ford K, Feltovich P, Hoffman R, editors. Human & machine, cognition. Hillsdale, NJ: Lawrence Erlbaum Associates; 1997. p. 67—99.

[13] Horsky J, Kuperman GJ, Patel VL. Comprehensive analysis of a medication dosing error related to CPOE: a case report. Journal of the American Medical Informatics Association 2005;12:377—82.

[14] Patel VL, Arocha JF, Zhang J. Thinking and reasoning in medicine. In: Holyoak K, Morrison RG, editors. The Cambridge handbook of thinking and reasoning. Cambridge University Press: Cambridge, UK; 2005. p. 727—50.

[15] Patel VL, Zhang J, Yoskowitz NA, Green R, Sayan OR. Translational cognition for decision support in critical care environments: a review. Journal of Biomedical Informatics 2008;41:413—31.

[16] King RD, Whelan KE, Jones FM, Reiser PG, Bryant CH, Muggleton SH, et al. Functional genomic hypothesis generation and experimentation by a robot scientist. Nature 2004;427(6971):247—52.

[17] Soldatova LN, Clare A, Sparkes A, King RD. An ontology for a robot scientist. Bioinformatics 2006;22(14):e464—71.

[18] Abidi SS. Knowledge management in healthcare: towards 'knowledge-driven' decision-support services. International Journal of Medical Informatics 2001;63:5—18.

[19] Stefanelli M. The socio-organizational age of artificial intelligence in medicine. Artificial Intelligence in Medicine 2001;23(1):25—47.

[20] Brooking A. Intellectual capital. London: International Thomson Business Press; 1996.

[21] Choo CW. The knowing organization. New York: Oxford University Press; 1998.

[22] Nonaka I, Takeuchi H. The knowledge-creating company. Oxford, UK: University Press; 1995.

[23] Polanyi M. The tacit dimension. London: Routledge & Kegan Paul; 1966.

[24] Bellazzi R, Zupan B. Predictive data mining in clinical medicine: current issues and guidelines. International Journal of Medical Informatics 2008;77(2):81—97.

[25] Quaglini S, Dazzi L, Gatti L, Stefanelli M, Fassino C, Tondini C. Supporting tools for guideline development and dissemination. Artificial Intelligence in Medicine 1998;14:119—37.

[26] Fox J, Johns N, Rahmanzadeh A. Disseminating medical knowledge: the PROforma approach. Artificial Intelligence in Medicine 1998;14:157—81.

[27] Argyris C, Schön D. Organizational learning II. London: Addison Wesley; 1996.

[28] Quaglini S, Stefanelli M, Cavallini A, Micieli G, Fassino C, Mossa C. Guideline-based careflow systems. Artificial Intelligence in Medicine 2000;20(1):5—22.

[29] Panzarasa S, Maddè S, Quaglini S, Pistarini C, Stefanelli M. Evidence-based careflow management systems: the case of post-stroke rehabilitation. Journal of Biomedical Informatics 2002;35(2):123—39.

[30] Leonardi G, Panzarasa S, Quaglini S, Stefanelli M, Van der Aalst WMP. Interacting agents through a web-based health serviceflow management system. Journal of Biomedical Informatics 2007;40(5):486—99.

[31] Panzarasa S, Quaglini S, Micieli G, Marcheselli S, Pessina M, Pernice C, et al. Improving compliance to guidelines through workflow technology: implementation and results in a stroke unit. In: Kuhn KA, Warren JR, Leong T-Y, editors. Proceedings of the 12th world congress on health (medical) informatics MEDINFO 2007. Amsterdam: IOS Press; 2007. p. 834—9.

[32] Wenger E, Snyder W. Communities of practice: the organizational frontier. Harvard Business Review 2000;January—February:139—45.

[33] Palladini G, Kyle RA, Larson DR, Therneau TM, Merlini G, Gertz MA. Multicentre versus single centre approach to rare diseases: the model of systemic light chain amyloidosis. Amyloid 2005;12(2):120—6.

[34] Kurzweil R. The singularity is near. New York: Viking Press; 2005.

[35] Butte AJ, Kohane IS. Creation and implications of a phenome—genome network. Nature Biotechnology 2006;24 (1):55—62.

[36] Rind DM, Safran C, Phillips RS, Wang Q, Calkins DR, Delbanco TL, et al. Effect of computer-based alerts on the treatment and outcomes of hospitalized patients. Archives of Internal Medicine 1994;154(13):1511—7.

[37] Warner HR. Computer assisted medical decision-making. New York: Academic Press, Inc.; 1979.

[38] Institute of Medicine. For the record: protecting electronic health information. Washington, DC: National Academy Press; 1997.

[39] Trepetin S. Privacy in context: the costs and benefits of a new deidentification method. PhD dissertation (computer science). Cambridge, MA: Massachsetts Institute of Technology; 2006.

[40] Berthold MR, Dill F, Koetter T, Thiel K. Supporting creativity: towards associative discovery of new insights. In: Washio T, Suzuki E, Ting KM, Inokuchi A, editors. Advances in knowledge discovery and data mining: proceedings of the 12th Pacific—Asia conference on knowledge discovery and data mining, PAKDD 2008; 2008. p. 14—25.

[41] Coiera EW. Artificial intelligence in medicine: the challenges ahead. Journal of the American Medical Informatics Association 1996;3(6):363—6.

[42] Haux R. Preparing for change: medical informatics international initiatives for health care and biomedical research. Computer Methods and Programs in Biomedicine 2007;88: 191—6.

[43] Zupan B, Holmes JH, Bellazzi R. Knowledge-based data analysis and interpretation. Artificial Intelligence in Medicine 2006;37(3):163—5.

[44] Holmes JH, Peek N. Intelligent data analysis in biomedicine. Journal of Biomedical Informatics 2007;40(6):605—8.