

Artificial Intelligence in Medicine

Harjeev Singh¹, Sukhwinder Singh²

Student¹, Assistant Professor²

^{1,2}Electronics and electrical communications engineering, PEC University of Technology, Chandigarh, India

Abstract: This paper aims at characterizing the maturity and influence achieved in the field of Artificial Intelligence in Medicine (AIM) till date. It ranges from reasoning under uncertainty, clinical decision making, and knowledge representation to systems integration, translational bioinformatics, and cognitive issues. It also highlights current developments, as well as the challenges faced by A.I.M.

Keywords: Artificial intelligence, medicine, research, technology, development, knowledge.

I. 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 Dartmouth College conference in 1956 [24]). Early AI in medicine researchers had discovered the relevance of AI methods to life sciences, most visibly in the Dendral experiments [1] of the late 1960s and early 1970s, which brought together chemists (e.g., Carl Djerassi), computer scientists (e.g., Edward Feigenbaum), 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 a symbolic form. There was an immense interest in biomedical applications of AI during the 70s, catalyzed in part by the creation of the SUMEX-AIM Computing Resource [2] at Stanford University, and another at Rutgers University, which took advantage of the nascent ARPANET to make computing cycles available to a national (and international) community of researchers applying AI methods to problems in medicine and biology. Several early AIM systems including Internist-1 [3], CASNET [4], and MYCIN [5], 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. By 1978, a leading journal in the field had devoted a special issue [6] 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. The field was then in the midst of “AI winter” [23], 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 [7].

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 who believe that a system’s performance on tasks is more important 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 varied purposes in the AI in medicine community.

II. INFLUENCE OF A.I. ON MEDICINE

AI in medicine cannot be set off from the rest of biomedical informatics, or from the world of health planning and policy. For a realistic expectation of the field’s influence in health care and biomedical sciences we require drawing 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 like 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 application, and its value in biomedicine, that must drive the work, which means being liberal as well as oriented to policy and sociocultural realities as we are to the technical underpinnings of a medical AI application. 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 merger of knowledge-based tools with other applications. The notion of stand-alone consultation systems had been well disparaged by the late 1980s [8], and thus we

must be looking for ways to combine relegate AI notions with such ubiquitous systems as electronic medical records, provider order-entry systems, results reporting systems, e-prescribing systems, or 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. 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, is inherently information-management tasks – and must accordingly be tackled and supported as such. To this day it is 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.

III. DEVELOPMENT IN A. I.

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 [9]), 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.

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 [8], 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 [10], 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.

IV. KNOWLEDGE MANAGEMENT

As far as medicine is concerned, knowledge management (KM) is one of the most interesting AI fields [11]. 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 [12]. The main goals of HCOs 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, doesn’t only focus on improving the availability of instruments for improving communication, but also 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 [13]. From an organizational viewpoint [14,15]. Knowledge creation is the result of

a social interaction between two fundamental types of knowledge, tacit knowledge and explicit knowledge [16]. 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..

V. CHALLENGES

AIM research faces numerous interesting challenges:

A. Better data capture and handling

Much of the early AIM research focused on capturing the expertise of human experts in sophisticated computer programs. 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 misreport even basic facts such as the chief complaint, but 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.

B. Improved design, modeling and assistance for workflows

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. 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 [17]. 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 is 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 [18]. We need to make this part of routine practice, and to overcome impediments to its adoption and use.

C. Reliable methods for reassuring patients in their concerns for 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 [19]. 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 irrefutable 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 [20]. 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.

D. 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.

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.

CONCLUSION

In summary, the challenge for AI in the coming years will be to ground the current research scenario in its AI roots. As recognized, adequate representation of all kinds of knowledge and high-level system modeling are important topics for basic AIM research. Moreover, the 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 [21,22]. Since AI in medicine applications today span from molecular medicine to organizational modeling, the role of modeling human reasoning and cognitive science must be reevaluated. Modeling and reasoning plays 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. Finally, strong interdisciplinary education programs should be fostered, to improve the quality of researchers and practitioners and to help the dissemination of AI methods and principles in the biomedical informatics community.

REFERENCES

- [1]. Lindsay RK, Buchanan BG, Feigenbaum EA, Lederberg J. Applications of Artificial intelligence for Organic Chemistry: The DENDRAL Project. New York: McGraw-Hill; 1980.
- [2]. Freiherr G. US Dept of Health, Education, and Welfare, Public Health Service. National Institutes of Health; Washington, D.C.; U.S. G.P.O: 1980. The Seeds of Artificial intelligence: SUMEX-AIM; pp. 80-2071. DHEW publication no. (NIH)
- [3]. 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):468-476. [PubMed]
- [4]. Weiss SM, Kulikowski CA, Amarel S, Safir A. A model-based method for computer-aided medical decision making. Artificial Intelligence. 1978;11:145-172.
- [5]. Shortliffe EH. Computer-Based Medical Consultations: MYCIN. New York: Elsevier; 1976.
- [6]. Sridharan NS. Guest editorial. Artificial Intelligence. 1978;11(1-2):1-4.
- [7]. Shortliffe EH. The adolescence of AI in Medicine: Will the field come of age in the '90s? Artif Intell Med. 1993;5:93-106. [PubMed]
- [8]. Miller RA, Maserie F. The demise of the Greek oracle model for medical diagnosis systems. Methods of Information in Medicine. 1990;29:1-2.

- [9]. Noy NF, Crubezy M, Ferguson RW, Knublauch H, Tu SW, Vendetti J, Musen MA. In: Musen MA, Friedman CP, Teich JM, editors. Protégé-2000: an open-source ontology-development and knowledge-acquisition environment; 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.
- [10]. 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. pp. 67–99.
- [11]. Abidi SS. Knowledge management in healthcare: Towards 'knowledge-driven' decision-support services. *International J of Medical Informatics*. 2001;63:5–18.
- [12]. Stefanelli M. The socio-organizational age of artificial intelligence in medicine. *Artificial Intelligence in Medicine*. 2001;23(1):25–47.
- [13]. Brooking A. *Intellectual Capital*. London: International Thomson Business Press; 1996.
- [14]. Choo CW. *The Knowing Organization*. New York: Oxford University Press; 1998.
- [15]. Nonaka I, Takeuchi H. *The Knowledge-Creating Company*. Oxford, UK: University Press; 1995.
- [16]. Polanyi M. *The Tacit Dimension*. London: Routledge & Kegan Paul; 1966.
- [17]. Rind DM, Safran C, Phillips RS, Wang Q, Calkins DR, Delbanco TL, Bleich HL, Slack WV. Effect of computer-based alerts on the treatment and outcomes of hospitalized patients. *Arch Intern Med*. 1994;154(13):1511–1517
- [18]. Warner HR. *Computer Assisted Medical Decision-Making*. New York: Academic Press, Inc; 1979.
- [19]. Institute of Medicine. *For the Record: Protecting Electronic Health Information*. Washington, DC: National Academy Press; 1997.
- [20]. Trepetin S. PhD Dissertation (Computer Science) Massachusetts Institute of Technology; Cambridge, MA: 2006. Privacy in Context: The Costs and Benefits of a New Deidentification Method.
- [21]. Zupan B, Holmes JH, Bellazzi R. Knowledge-based data analysis and interpretation. *Artif Intell Med*. 2006;37(3):163–5.
- [22]. Holmes JH, Peek N. Intelligent data analysis in biomedicine. *J Biomed Inform*. 2007;40(6):605–8
- [23]. The Coming of Age of Artificial Intelligence in Medicine, *Artif Intell Med*. May 2009; 46(1):5-17.
- [24]. History of artificial intelligence. http://en.wikipedia.org/wiki/History_of_artificial_intelligence.

