Cognitive Informatics in Biomedicine and Healthcare

Series Editor
Vimla L. Patel, Ctr Cognitive Studies in Med & PH
New York Academy of Med, Suite 454, New York, NY, USA

Enormous advances in information technology have permeated essentially all facets of life. Although these technologies are transforming the workplace as well as leisure time, formidable challenges remain in fostering tools that enhance productivity, are sensitive to work practices, and are intuitive to learn and to use effectively. Informatics is a discipline concerned with applied and basic science of information, the practices involved in information processing, and the engineering of information systems.

Cognitive Informatics (CI), a term that has been adopted and applied particularly in the fields of biomedicine and health care, is the multidisciplinary study of cognition, information, and computational sciences. It investigates all facets of computer applications in biomedicine and health care, including system design and computer-mediated intelligent action. The basic scientific discipline of CI is strongly grounded in methods and theories derived from cognitive science. The discipline provides a framework for the analysis and modeling of complex human performance in technology-mediated settings and contributes to the design and development of better information systems for biomedicine and health care.

Despite the significant growth of this discipline, there have been few systematic published volumes for reference or instruction, intended for working professionals, scientists, or graduate students in cognitive science and biomedical informatics, beyond those published in this series. Although information technologies are now in widespread use globally for promoting increased self-reliance in patients, there is often a disparity between the scientific and technological knowledge underlying healthcare practices and the lay beliefs, mental models, and cognitive representations of illness and disease. The topics covered in this book series address the key research gaps in biomedical informatics related to the applicability of theories, models, and evaluation frameworks of HCI and human factors as they apply to clinicians as well as to the lay public.

Trevor A. Cohen • Vimla L. Patel
Edward H. Shortliffe
Editors

Intelligent Systems in Medicine and Health

The Role of AI

Springer

Editors
Trevor A. Cohen
University of Washington
Seattle, WA, USA

Vimla L. Patel
New York Academy of Medicine
New York, NY, USA

Edward H. Shortliffe
Columbia University
New York, NY, USA

ISSN 2662-7280 ISSN 2662-7299 (electronic)
Cognitive Informatics in Biomedicine and Healthcare
ISBN 978-3-031-09107-0 ISBN 978-3-031-09108-7 (eBook)
https://doi.org/10.1007/978-3-031-09108-7

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

We wish to dedicate this volume to three giants of science who greatly inspired and influenced us while helping to lay the groundwork for artificial intelligence in medicine, even though each of them made their primary contributions in other fields. Furthermore, each has been a personal friend of at least one of us and we look back fondly on their humanity and their contributions.

Herbert Simon (Carnegie Mellon University) won the Nobel Prize in Economics (1978) for his pioneering research into the decision-making process within economic organizations. He is also known as an early innovator in the field of artificial intelligence, whose long partnership with Allen Newell led to early work on the Logic Theory Machine and General Problem Solver. As a psychologist who studied human and machine problem solving, he pioneered notions of bounded rationality and satisficing.

Joshua Lederberg (Stanford and Rockefeller Universities) won the Nobel Prize in Physiology or Medicine (1958) for his discoveries of genetic transfer in bacteria. A computer scientist as well as a geneticist, he was instrumental in devising the notion of capturing expert knowledge in computers. His Dendral Project led to a body of work that dealt with the interpretation of mass spectral data to identify organic compounds, and he created the first national resource to support research on artificial intelligence in medicine (SUMEX-AIM).

Walter Kintsch (University of Colorado Boulder) is a psychologist and cognitive scientist who used both experimental and computational modeling techniques to study how people understand language. He reformulated discourse comprehension as a constraint satisfaction process. His work on latent semantic analysis demonstrated how machine learning could be used to construct a high-dimensional semantic space. His writing has influenced research on natural language understanding as well as the cognitive informatics focus reflected in this volume.

T.C., V.L.P., E.H.S.—March 2022

Herbert Simon
1916 - 2001

Joshua Lederberg
1925 - 2008

Walter Kintsch
1932 -

Photograph of Herbert Simon provided courtesy of Carnegie Mellon University.

Photograph of Joshua Lederberg provided courtesy of EH Shortliffe.

Photograph of Walter Kintsch provided by Professor Kintsch.

Foreword

Delivering quality, affordable health care to large populations is one of humankind's greatest challenges.

There are so many facets to health care, however, that talking about "the problem" is a gross understatement. George Polya's heuristic (from Descartes) of dividing problems into subproblems[1] remains one of the best tools in our cognitive toolkit—human or artificial intelligence (AI).[2] This book explores many of those subproblems from the perspective of using computing—AI in particular—to meet them.

The idea of automating solutions to specific human problems is rooted in Greek mythology but actual implementations of automated problem solving with mechanical or electronic computers were initially restricted to single-purpose applications. For example, wartime uses included decoding encrypted messages and calculating trajectories of artillery shells; business uses included payroll accounting, information retrieval, and process automation. Medical uses included keeping medical records, recording physiological data, and calculating drug doses.

Modern medicine is grounded in data from observations as well as measurements. But the machines running initial applications of computers to medicine were so restricted in size and scope that demonstrations beyond single-purpose problems were severely limited.

Around 1950 Turing introduced the audacious idea that computers could be intelligent and could be used on problems that required reasoning, understanding,

[1] Polya G. *How to Solve It: A New Aspect of Mathematical Method*. Princeton, NJ: Princeton University Press, 2014 (originally published 1945).

[2] Two memorable examples of identifying key subproblems are cited in McCullough's book on building the Panama Canal [McCullough, D. *The Path between the Seas: The Creation of the Panama Canal, 1870-1914*. New York: Simon & Schuster, 1978]. First, after several years of failure, a new engineer put in charge of the project rethought the steps of digging a huge trench and recognized that the most rate-limiting step was getting rid of the dirt. Second, the same engineer recognized another rate-limiting step was keeping workers healthy. Neither was "the problem" per se; solving both subproblems made the difference.

and learning—to work smarter, not harder. Acquiring measurements is one thing, interpreting the data is another. Importantly, early AI research introduced a conceptual framework that includes two critical differences from mathematical computation: symbolic reasoning and heuristic reasoning. Reasoning with symbols, as opposed to numeric quantities, carries with it the essential link human beings made when they invented language between words and the world, i.e., semantics. Heuristic reasoning is even older than humankind, being bound up with the Darwinian notion of survival. A decision to fight or flee does not leave time to consider the consequences of all possible options,[3] and most decisions in everyday life have no algorithm that guarantees a correct answer in a short time (or often not in any amount of time).

The importance of human health quickly attracted early pioneers to the possibility of using computers to assist physicians, nurses, pharmacists, and other healthcare professionals.

The initial applications of AI limited their scope so they could be dealt with successfully in single demonstration projects on the small computers of the day. Cognitive psychologists recognized that clinicians *reason* about the available data in medicine just as they do in other fields like chess. The process of diagnostic reasoning became the focus of work on expert systems, with early programs becoming convincing examples that AI could replicate human reasoning in this area on narrowly defined problems.

An important part of this demonstration was with Bayesian programs that used statistics to reason in a mathematically rigorous way for clinical decision support. However, those initial applications of Bayes' Theorem to medicine were also limited until means were found to reduce their complexity with heuristic reasoning. Another important set of demonstrations encoded knowledge accumulated by human experts, along the lines suggested by cognitive science. One of the shortcomings of the expert systems approach, however, is the time and effort it takes to acquire the knowledge from experts in the first place, and to maintain a large knowledge base thereafter.

In practice, medical professionals are faced with several inconvenient truths, which further complicate efforts to use computers in health care. The chapters of this book address many of them. For example:

- The body of knowledge in, and relevant to, medicine is growing rapidly. E.g., diagnosis and treatment options for genetic diseases in the last few decades, and of viral infections in the last few years, have come into the mainstream.
- Complete paradigm shifts in medicine require rethinking whole areas of medicine. E.g., prion diseases have forced a reconceptualization of the mechanisms of pathogen transmission.

[3] Herb Simon, one of the founding fathers of AI, received a Nobel Prize in Economics in 1978 for elucidation of this concept in the world of decision making. His term for it was "Bounded Rationality." See https://www.cmu.edu/simon/what-is-simon/herbert-a-simon.html (accessed August 11, 2022).

- Medical knowledge is incomplete and there are no good treatment options, and sometimes no good diagnostic criteria, for some conditions. E.g., Parkinson's disease can be managed, but is still not curable after decades of research.
- Information about an individual patient is often erroneous and almost always incomplete. E.g., false positive and false negative test results are expected for almost every diagnostic test.
- Patients' medical problems exist within a larger emotional, social, economic, and cultural context. E.g., the most effective treatment options may be unaffordable or unacceptable to an individual.
- Professionals are expected to learn from their own, and others', experience (both positive and negative). E.g., continuing to recommend a failed treatment modality is reason for censure.
- Professionals at the level of recognized specialists are expected to deal with unique cases for which there are neither case studies nor established diagnostic or treatment wisdom. E.g., primary care providers refer recalcitrant cases to specialists for just these reasons.
- Communication between patients and professionals is imperfect. E.g., language is full of ambiguity, and we all have biases in what we want to hear or fear most.

Collectively, these issues are more than merely "inconvenient," but are humbling reminders that "the problem" of providing health care is overwhelming in the large. They also represent significant barriers to harnessing the presently available power of computers to actual healthcare delivery. The perspectives offered in this book summarize current approaches to these issues and highlight work that remains to be done. As such it is valuable as a textbook for biomedical informatics and as a roadmap for the possibilities of using AI for the benefit of humankind.

The book's emphasis on reasoning provides a central focus not found in other collections. The chapters here deal with transforming data about patients, once acquired, to actionable information and using that information in clinical contexts. With today's understanding, AI offers the means to augment human intelligence by making the accumulated knowledge available, suggesting possible options, and considering consequences. We are betting that computers can help to overcome human limitations of imperfect memory, reasoning biases, and sheer physical stamina. We are betting on the power of knowledge over the persistent forces of random mutations. Most of all we are betting that the synergy of human and computer intelligence will succeed in the noble quest of improving the quality of human lives.

University of Pittsburgh
Pittsburgh, PA, USA
August 2022

Bruce G. Buchanan

Preface

The State of AI in Medicine

Recent advances in computing power and the availability of large amounts of training data have spurred tremendous advances in the accuracy with which computers can perform tasks that were once considered the exclusive province of human intelligence. These include perceptual tasks related to medical diagnosis, where deep neural networks have attained expert-level performance for the diagnosis of diabetic retinopathy and the identification of biopsy-proven cases of skin cancer from dermatological images. These accomplishments reflect an increase in activity in Artificial Intelligence (AI), both in academia and industry. According to the 2021 AI Index report assembled and published at Stanford University, there was a close to twelvefold increase in the total number of annual peer-reviewed AI publications between 2000 and 2019, and a close to fivefold increase in annual private investment in funded AI companies between 2015 and 2020, with commercial applications of AI technologies such as speech-based digital assistants and personalized advertising and newsfeeds by now woven deeply into the threads of our everyday lives. These broad developments throughout society have led to a resurgence of public interest in the role of AI in medicine[1] (AIM), reviving long-standing debates about the nature of intelligence, the relative value of data-driven predictive models and human decision makers, and the potential for technology to enhance patient safety and to disseminate expertise broadly.

Consequently AIM is in the news, with frequent and often thoughtful accounts of the ways in which it might influence—and hopefully improve—the practice of medicine appearing in high visibility media venues such as *The New York Times* [1, 2], *The Atlantic* [3], *The New Yorker* [4], *The Economist* [5], and others [6, 7]. As biomedical informaticians with long-standing interests in AIM, we are encouraged to see this level of attention and investment in the area. However, we are also aware

[1] We consider the scope of AIM to include public health and clinically driven basic science research, as well as clinical practice.

that this is by no means the first time that the promise of AIM has emerged as a focus of media attention. For example, a 1977 *New York Times* article [8] describes the MYCIN system (discussed in Chap. 2), noting the ability of this system to make medical diagnoses, to request missing information, and to explain how it reaches conclusions. The article also mentions the extent of government funding for AI research at the time at $5 million a year, which when adjusted for inflation would be around $23 million annually. The question arises as to what has changed between then and now, and how these changes might affect the ongoing and future prospects for AI technologies to improve medical care. Some answers to this question are evident in this 1977 article, which is critical of the potential of the field to accomplish its goals of championship-level chess performance and machine translation ("neither has been accomplished successfully, and neither is likely to be any time soon") and face recognition ("they cannot begin to distinguish one face from another as babies can"). Today all of these tasks are well within the capabilities of contemporary AI systems, which is one indication of the methodological advances that have been made over the intervening four decades. However, translating methods with strong performance on tightly constrained tasks to applications with positive impact in health care is not (and has never been) a straightforward endeavor, on account of the inherent complexities of the healthcare system and the prominent role of uncertain and temporarily unavailable information in medical decision making, among other factors. While AI technologies do have the potential to transform the practice of medicine, computer programs demonstrating expert-level performance in diagnostic tasks have existed for decades, but significant challenges to realizing the potential value of AI in health care—such as how such AI systems might best be integrated into clinical practice—remain unresolved.

It is our view that the discipline of Cognitive Informatics (CI) [9], which brings the perspective of the cognitive sciences to the study of medical decision making by human beings and machines, is uniquely positioned to address many of these challenges. Through its roots in the study of medical reasoning, CI provides a sound scientific basis from which to consider the relationship between current technologies and human intelligence. CI has extended its area of inquiry to include both human-computer interaction and the study of the effects of technology on the flow of work and information in clinical settings [10]. Accordingly, CI can inform the integration of AI systems into clinical practice. In recent years, patient safety has emerged as a research focus of the CI community, providing new insights into the ways in which technology might mitigate or, despite best intentions, facilitate medical error. Consequently, a volume describing approaches to AIM from a CI perspective seemed like an excellent fit for the Springer Nature Cognitive Informatics book series led by one of us (VLP). We fondly recall that this volume is a project that we first discussed several years ago over a bottle of Merlot in proximity to a conference in Washington, DC.

However, as our discussions developed over the course of subsequent meetings, it became apparent that there was a need for a more comprehensive account of the field. As educators, we considered the knowledge and skills that future researchers and practitioners in the field might need in order to realize the transformative potential of AI methods for the practice of health care. We were aware of books on the subject, such as cardiologist Eric Topol's cogent account of the implications of contemporary machine learning approaches for the practice of medicine [11], and the development of pediatric cardiologist Anthony C. Chang's excellent clinician-oriented introduction to AI methods with a focus on their practical application to problems in health care and medicine [12].[2] Books drawing together the perspectives of multiple authors were also available, including a compendium of chapters in which authors focus on their research interests within the field [13], and an account of the organizational implications of "big data" and predictive models that can be derived from such large collections of information [14]. However, none of the books we encountered was developed with the focused intention to provide a basis for curricular development in AIM, and we were hearing an increasing demand for undergraduate, graduate, and postgraduate education from our own trainees—students and aspiring physician-informaticians.

Consequently, we pivoted from our original goal of a volume primarily concerned with highlighting the role of CI in AIM, to the goal of developing the first comprehensive coauthored textbook in the AIM field, still with a CI emphasis. We reached out to our friends and colleagues, prominent researchers with deep expertise in the application of AI methods to clinical problems, several of whom have been engaged with AIM since the inception of the field. This was a deliberate decision on our part, as we felt that in addition to lacking a cognitive perspective (despite the emergence of the term "cognitive computing" as a catch-all for AI methods), much work we encountered was presented without apparent consideration of the history of the field. Our concern with this disconnect was not only a matter of the academic impropriety of failure to acknowledge prior work. We were also concerned that work conducted from this perspective would not be informed by the many lessons learned from decades of work and careful consideration of the issues involved in implementing AI at the point of care. Thus, in developing our ideas for the structure of this volume, and in our selection of chapter authors, we endeavored to make sure that presentations of current methods and applications were contextualized in relation to the history of the field, and informed by a CI perspective.

Introducing Intelligent Systems in Medicine and Health: The Role of AI

The result of these efforts is the current volume, a comprehensive textbook that takes stock of the current state of the art of AIM, places current developments in their historical context, and identifies cognitive and systemic factors that may help or hinder their ability to improve patient care. It is our intention that a reader of this

[2] A trained data scientist, Dr. Chang also contributed a chapter on the future of medical AI and data science in this volume.

volume will attain an accurate picture of the strengths and limitations of these emerging technologies, emphasizing how they relate to the AI systems that preceded them, to the intelligence of human decision makers in medicine, and to the needs and expectations of those who use the resulting tools and systems. This will lay a foundation for an informed discussion of the potential of such technology to enhance patient care, the obstacles that must be overcome for this to take place, and the ways in which emerging and as-yet-undeveloped technologies may transform the practice of health care.

With increasing interest and investment in AIM technologies will come the recognition of the need for professionals with the prerequisite expertise to see these technologies through to the point of positive impact. Progress toward this goal will require both advancement of the state of the science through scholarly research and measurably successful deployment of AIM systems in clinical practice settings. As the first comprehensive coauthored textbook in the field of AIM, this volume aims to define and aggregate the knowledge that researchers and practitioners in the field will require to advance it. As such, it draws together a range of expert perspectives to provide a holistic picture of the current state of the field, to identify opportunities for further research and development, and to provide guidance to inform the successful integration of AIM into clinical and public health practice.

We intend to provide a sound basis for a seminar series or a university level course on AIM. To this end, authors have been made aware of the context of their chapters within the logical flow of the entire volume. We have sought to assure coordination among authors to facilitate cross-references between chapters, and to minimize either coverage gaps or redundancies. Furthermore, all chapters have followed the same basic organizational structure, which includes explicit learning objectives, questions for self-study, and annotated suggestions for further reading. Chapters have been written for an intended audience of students in biomedical informatics, AI, machine learning, cognitive engineering, and clinical decision support. We also offer the book to established researchers and practitioners of these disciplines, as well as those in medicine, public health, and other health professions, who would like to learn more about the potential for these emerging technologies to transform their fields.

Structure and Content

The book is divided into four parts. They are designed to emphasize pertinent *concepts* rather than technical detail. There are other excellent sources for exploring the technical details of the topics we introduce. The **Introduction** provides readers with an overview of the field. Chapter 1 provides an introduction to the fields of Artificial Intelligence and Cognitive Informatics and describes how they relate to one another. Chapter 2 provides a historical perspective, drawing attention to recurring themes, issues, and approaches that emerged during the course of the development of early AI systems, most of which remain highly relevant today. Chapter 3 provides an overview of the landscape of biomedical data and information, to familiarize readers with the resources AIM systems can draw upon for training and evaluation, and as sources of structured knowledge. This chapter also serves as a bridge to subsequent parts in the book, introducing some of the methods, applications, and issues that will be covered in greater detail later.

The **Approaches** part covers methods used by AIM systems. Chapter 4 commences this part with a discussion of knowledge-based approaches, spanning from knowledge modeling efforts, used in expert systems at the inception of AIM as a field, to contemporary efforts to infuse machine learning models with structured biomedical knowledge. Chapter 5 considers the idea of AIM from a cognitive perspective, beginning with an account of the parallel development of the two fields, and advancing an argument for the development of complementary human/AI systems with superior clinical utility when compared with either of these components alone. Chapter 6 introduces the machine learning approaches that have come to predominate in the current wave of AIM systems, with the intention of providing readers with a conceptual understanding of some key algorithms, and issues that may arise during their application to modeling medical data. Chapter 7 considers Natural Language Processing (NLP), which has been intertwined with AI at least since Turing's proposal of the ability to conduct a passably human conversation as an observable surrogate for "thinking" [15]. The chapter focuses on biomedical NLP, giving an account of the main problems to be addressed in this area, and the methods that predominate at present, from rule-based methods through to deep learning approaches. Chapter 8 considers approaches to the explanation of decisions recommended by AIM models, with an emphasis on its importance in clinical settings in which "black box" predictions may not—and arguably should not—be taken at face value without some justification to earn trust and to permit detection of predictions made on tenuous grounds. Chapter 9 revisits the issue of language, with a focus on methods that support human-like conversational interactions with the goal of supporting health care.

Having introduced the fundamental methods of the field, the book focuses in the third part on **Applications**. Chapter 10 describes those in which AIM methods are applied to support decision making in clinical care, through integration with existing platforms and workflows. Chapter 11 focuses on the prediction of clinical outcomes, such as near-term readmission, with an emphasis on dynamic models of variables and outcomes that change over time. Chapter 12 describes what is arguably the most successful area of application of contemporary AIM methods to date—the interpretation of medical images. Chapter 13 shifts from a focus on methods to support the care of individual patients to the population level, with a discussion of the emerging role of AI in the field of public health. Chapter 14 then describes AI applications at the molecular level: methods that yield clinically actionable insights from -omics data. Chapter 15 discusses AIM applications that are of administrative importance to a healthcare system, such as approaches to optimize the use of resources. This chapter also includes a discussion of the practical concerns that may arise when attempting to implement AI solutions in the context of an operational healthcare system. Finally, Chap. 16 focuses on AI applications in medical education, including the roles of cognitive and learning sciences in informing how clinicians-in-training should be educated about AIM, and how AI might support the education of clinicians in their chosen fields.

The final part in the volume considers **the road ahead** for AIM. It addresses issues that are likely to be of importance for the successful progression of the field, including two potential stumbling blocks: inadequate evaluation, and failure to consider the ethical issues that may accompany the deployment of AI systems in healthcare settings. Accordingly, the first chapter in this part, Chap. 17, is focused on evaluation of AIM systems, including enabling capabilities such as usability and the need to move beyond "in situ" evaluations of accuracy toward demonstrations of acceptance and clinical utility in the natural world. Chapter 18 concerns the need for a robust ethical framework to address issues proactively such as algorithmic bias, and exacerbation of healthcare inequities due to limited portability of methods and algorithms. Chapter 19 projects from the trajectory of current trends to anticipate the future of AI in medicine with an emphasis on data science, and how broader deployment of AI systems may affect the practice of medicine. Finally, Chap. 20 provides a summary and synthesis of the volume, including the editors' perspectives on the prospects and challenges for the field. The final part is then followed by a detailed glossary that provides definitions of all terms displayed in bold throughout the body of the book (with an indication of the chapter(s) in which each term was used). The book closes with a subject index for the entire volume.

Guide to Use of This Book

This book is written as a textbook, such that it can be used in formal courses or seminars. For this purpose, we would anticipate curricular design to follow the overall structure of the book, with a logical progression from introduction through approaches to applications and projections for the future. For example, this structure could support an undergraduate or graduate level course in a Computer Science, Biomedical Informatics, or Cognitive Engineering program that aims to provide students with a comprehensive survey of current applications and concerns in the field. At the graduate level, this could be coupled to a student-led research project. Alternatively, one might imagine an MS level course that aims to provide clinical practitioners seeking additional training in clinical informatics with the knowledge they will need to be informed users of AIM systems, in which case content could be drawn from the book selectively with an emphasis on introductory content, and clinical applications and issues that relate to them directly (Chaps. 1–3, 10–12, 15, and 17–19). Of course, the book may be used for self-study and reference, and readers may wish to explore particular topics in greater detail—starting with a particular chapter (say, machine learning methods) and then exploring the cross-references in this chapter to find out more about how this topic features in the context of particular applications, or issues that are anticipated to emerge as the field progresses.

This is an exciting time to be working in the field of AIM, and an ideal time to enter it. There is increasing support for AIM work, both through federal funding initiatives such as the NIH-wide Bridge to Artificial Intelligence program[3] and in light of an acceleration in private investment in digital health technologies. Such support has been stimulated in part by the field's demonstrated utility and acceptance as a way to diagnose disease, to deliver care, and to support public health efforts during the COVID-19 pandemic. On account of the pervasiveness of AI technologies across industries outside of health care, skepticism about the ability of these technologies to deliver meaningful improvements is balanced by enthusiasm for their potential to improve the practice of medicine. It is our goal that readers of this volume will emerge equipped with the knowledge needed to realize this potential and to proceed to lead the advancement of health care through AIM.

References

1. Metz C. A.I. shows promise assisting physicians. The New York Times. 2019. [cited 2021 Apr 15]. https://www.nytimes.com/2019/02/11/health/artificial-intelligence-medical-diagnosis.html.
2. O'Connor A. How artificial intelligence could transform medicine. The New York Times. 2019. [cited 2021 Apr 15]. https://www.nytimes.com/2019/03/11/well/live/how-artificial-intelligence-could-transform-medicine.html.
3. Cohn J. The robot will see you now. The Atlantic. 2013. [cited 2021 Jun 14]. https://www.theatlantic.com/magazine/archive/2013/03/the-robot-will-see-you-now/309216/.
4. Mukherjee S. A.I. versus M.D.. The New Yorker. [cited 2021 Apr 15]. https://www.newyorker.com/magazine/2017/04/03/ai-versus-md.
5. Artificial intelligence will improve medical treatments. The Economist. 2018. [cited 2021 Jun 16]. https://www.economist.com/science-and-technology/2018/06/07/artificial-intelligence-will-improve-medical-treatments.
6. Aaronovitch D. DeepMind, artificial intelligence and the future of the NHS. [cited 2021 Jun 16]. https://www.thetimes.co.uk/article/deepmind-artificial-intelligence-and-the-future-of-the-nhs-r8c28v3j6.
7. Artificial intelligence has come to medicine. Are patients being put at risk?. Los Angeles Times. 2020. [cited 2021 Jun 16]. https://www.latimes.com/business/story/2020-01-03/artificial-intelligence-healthcare.
8. Experts argue whether computers could reason, and if they should. The New York Times. [cited 2021 Apr 15]. https://www.nytimes.com/1977/05/08/archives/experts-argue-whether-computers-could-reason-and-if-they-should.html.
9. Patel VL, Kaufman DR, Cohen T, editors. Cognitive informatics in health and biomedicine: case studies on critical care, complexity and errors. London:

[3] https://commonfund.nih.gov/bridge2ai (accessed August 11, 2022).

Springer; 2014. [cited 2021 Jun 16]. (Health Informatics). https://www.springer.com/gp/book/9781447154891.
10. Patel VL, Kaufman DR. Cognitive science and biomedical informatics. In: Shortliffe EH, Cimino JJ, editors. Biomedical informatics: computer applications in health care and biomedicine. 5th ed. New York: Springer; 2021. p. 133–85.
11. Topol E. Deep medicine. How artificial intelligence can make healthcare human again. Hachette UK; 2019.
12. Chang AC. Intelligence-based medicine: artificial intelligence and human cognition in clinical medicine and healthcare. Academic Press; 2020.
13. Agah A, editor. Medical applications of artificial intelligence. Boca Raton: CRC Press; 2013. 526 p.
14. Natarajan P, Frenzel JC, Smaltz DH. Demystifying big data and machine learning for healthcare. Boca Raton: CRC Press; 2017. 210 p.
15. Turing AM. Computing machinery and intelligence. Mind : a quarterly review of psychology and philosophy. 1950;LIX:433.

Seattle, WA, USA — Trevor A. Cohen
New York, NY, USA — Vimla L. Patel
New York, NY, USA — Edward H. Shortliffe
August 2022

Acknowledgments

While my coeditors are well known for their prescience in anticipating (and influencing) the evolution of AIM, I think it is safe to say that none of us envisioned developing a textbook together when we first met at the annual retreat of Columbia University's Department of Medical Informatics in the Catskills in 2002. At the time, I had just joined the program as an incoming graduate student, attracted in part by what I had learned of Ted's work in medical AI and Vimla's in medical cognition. As these topics have remained core components of my subsequent research, it is especially encouraging to see the upsurge in interest in the field. The expanding pool of talented graduate students and physician-informaticians excited about the potential of AIM was a key motivator for our development of this volume, as was our recognition of the need for a textbook in the field that represented historical and cognitive perspectives, in addition to recent methodological developments. However, given the breadth of methods and biomedical applications that by now fall under the AIM umbrella, it was apparent to us that we would need to draw upon the expertise of leaders in relevant fields to develop a multiauthored textbook. Our efforts to weave the perspectives of these authors into a coherent textbook benefited considerably from Ted's experience as lead editor of a multiauthored textbook in biomedical informatics (fondly known as the "blue bible" of BMI, and now in its fifth edition). We modeled both our approach to encouraging the integration of ideas across chapters and the key structural elements of each chapter on the example of that book. We were aided in this endeavor by our authors, who were highly responsive to our suggestions for points of reference between chapters, as well as our recommendations to reduce overlap. We especially appreciated our correspondence with authors during the editing process, which broadened our own perspectives on AIM, and influenced the themes we focused on when developing the final chapter (Reflections and Projections). We owe additional gratitude to Grant Weston, Executive Editor of Springer's Medicine and Life Sciences division, for his steadfast support and guidance throughout the development of this volume. We would also like to acknowledge our Production Editor Rakesh Jotheeswaran, Project Manager Hashwini Vytheswaran, and Editorial Assistant Leo Johnson for their assistance in keeping the project on track. Special thanks also go to Bruce Buchanan (an AI luminary, key mentor for Ted's work in the 1970s, and his long-term collaborator and colleague) for his willingness to craft the foreword for this volume. The development of this volume coincided with a global pandemic. Many involved in the project were affected by this in a professional capacity through their practice of medicine, their role in an informatics response at their institutions, or the position of their professional home within a healthcare system under siege. We trust the development of their chapters provided a welcome respite for these authors, as ours did for us, and hope that readers of this volume will be inspired to develop AIM solutions that equip us to anticipate and manage global health crises more effectively in the future.

Seattle, WA, USA Trevor A. Cohen
August 2022

Contents

Contributors

Emily Alsentzer Massachusetts Institute of Technology, Cambridge, MA, USA

Riccardo Bellazzi Department of Electrical, Computer and Biomedical Engineering, University of Pavia, Pavia, Italy

Laboratory of Informatics and Systems Engineering for Clinical Research, Istituti Clinici Scientifici Maugeri, Pavia, Italy

Timothy Bickmore Khoury College of Computer Sciences, Northeastern University, Boston, MA, USA

Bethene D. Britt UW Medicine, University of Washington, Seattle, WA, USA

David L. Buckeridge McGill University, Montreal, QC, Canada

Anthony C. Chang Medical Intelligence and Innovation Institute (MI3), Children's Health of Orange County, Orange, CA, USA

Jake Y. Chen Informatics Institute, University of Alabama at Birmingham, Birmingham, AL, USA

Trevor A. Cohen University of Washington, Seattle, WA, USA

Gregory F. Cooper Department of Biomedical Informatics, University of Pittsburgh, Pittsburgh, PA, USA

Arianna Dagliati Department of Electrical, Computer and Biomedical Engineering, University of Pavia, Pavia, Italy

Dev Dash Stanford University School of Medicine, Stanford, CA, USA

Parvati Dev SimTabs LLC, Los Altos Hills, CA, USA

Kenneth W. Goodman Institute for Bioethics and Health Policy, University of Miami, Miami, FL, USA

Michael B. Gotway Mayo Clinic, Scottsdale, AZ, USA

Tina Hernandez-Boussard Stanford University School of Medicine, Stanford, CA, USA

Andrew J. King Department of Critical Care Medicine, University of Pittsburgh, Pittsburgh, PA, USA

Diane M. Korngiebel Department of Biomedical Informatics and Medical Education, University of Washington, Seattle, WA, USA

Ron C. Li Stanford University School of Medicine, Stanford, CA, USA

Jianming Liang Arizona State University, Phoenix, AZ, USA

Naveen Muthu University of Pennsylvania School of Medicine, Philadelphia, PA, USA

Children's Hospital of Philadelphia, Philadelphia, PA, USA

Thanh M. Nguyen Informatics Institute, University of Alabama at Birmingham, Birmingham, AL, USA

Giovanna Nicora Department of Electrical, Computer and Biomedical Engineering, University of Pavia, Pavia, Italy

Vimla L. Patel The New York Academy of Medicine, New York, NY, USA

Kirk Roberts The University of Texas Health Science Center at Houston, Houston, TX, USA

Martìn-Josè Sepùlveda Florida International University, Miami, FL, USA

Nigam H. Shah Stanford University School of Medicine, Stanford, CA, USA

Edward H. Shortliffe Vagelos College of Physicians and Surgeons, Columbia University, New York, NY, USA

Ida Sim University of California San Francisco, San Francisco, CA, USA

Marina Sirota University of California San Francisco, San Francisco, CA, USA

Anthony Solomonides Research Institute, NorthShore University HealthSystem, Evanston, IL, USA

Devika Subramanian Rice University, Houston, TX, USA

Peter Szolovits Massachusetts Institute of Technology, Cambridge, MA, USA

Shyam Visweswaran Department of Biomedical Informatics, University of Pittsburgh, Pittsburgh, PA, USA

Byron Wallace Khoury College of Computer Sciences, Northeastern University, Boston, MA, USA

Adam B. Wilcox Washington University in St. Louis School of Medicine, St. Louis, MO, USA

Hua Xu The University of Texas Health Science Center at Houston, Houston, TX, USA

Zongwei Zhou Johns Hopkins University, Baltimore, MD, USA