

AI in Medicine on its way from knowledge-intensive to data-intensive systems

Werner Horn ^a

^a*Department of Medical Cybernetics and Artificial Intelligence, University of Vienna
Freyung 6, A-1010 Vienna, Austria
and
Austrian Research Institute for Artificial Intelligence, Vienna
E-mail: werner@ai.univie.ac.at*

Abstract

The last 20 years of research and development in the field of artificial intelligence in medicine show a path from knowledge-intensive systems, which try to capture the essential knowledge of experts in a knowledge-based system, to data-intensive systems available today. Nowadays enormous amounts of information is accessible electronically. Large data sets are collected continuously monitoring physiological parameters of patients. Knowledge-based systems are needed to make use of all these data available and to help us to cope with the information explosion. In addition, temporal data analysis and intelligent information visualization can help us to get a summarized view of the change over time of clinical parameters. Integrating AIM modules into the daily-routine software environment of our care providers gives us a great chance for maintaining and improving quality of care.

1 AIM: a partial view of its scope and potential

This paper gives a personalized view of research and development in artificial intelligence in medicine (AIM) based on the work of my colleagues, myself, and cooperating scientists over the last 20 years.

My main point is that research in AIM has shifted during the last 20 years considerably with respect to the following aspects:

- a strong movement from knowledge-intensive systems to data-intensive systems, and
- static consultation systems which tried to capture the current situation of a patient are followed by dynamic systems, which have a close look at the change over time.

The second move was possible only because we succeeded in doing the first move. However, the first move towards intelligent data analysis does not imply there is no explicit knowledge inside such systems. On the contrary, the transfer of expert knowledge into data-intensive knowledge-based systems is a must for success.

It is worth noting that research in AIM has given us several major breakthroughs in AI technology in the past. Medicine is a domain very fruitful for exploring intelligent systems. This is because it is complex, not really well understood in all its aspects, and it requires adaptation for doing proper actions. This contrasts technical domains where one knows about correct functioning and optimal performance (at least theoretically). Medicine is a domain much more challenging. As a consequence I do expect most interesting breakthroughs in AI technology and application building in this domain.

There are many areas most interesting to explore in the future. Let me focus on a few of these, which are close to my personal interests:

- automation of more “intelligible” routine tasks in domains with sufficient on-line data like intensive care units,
- analysis and abstraction of temporal data, in order to support a comprehensive picture about what has happened in the recent past,
- graphical visualization tools to make visible what is hidden in the enormous amounts of data collected today,
- utilizing best practice guidelines for individual cases with the help of world-wide information technology.

In addition to the changes in methodology and application building we see a change in the form of system implementation. We started with monolithic consultation tools running on stand-alone hardware (like Lisp machines) which never found its way into practical use in a clinic or physician’s office. Today, integration is the main issue. AI modules become embedded into more general systems and devices. These modules are able to provide intelligent support and expert advice in a networked world of data collection, data storage, information extraction, interpretation, and medical action. We no longer sell these modules as distinguished “Artificial Intelligence” products. But they are essential to ease much of our daily work with computers.

Furthermore, AI technology is very essential to ensure quality standards in the health-care sector. Computerized patient records give us a chance today to have them examined by programs which know about the standards and rules of a clinic. AI modules should be developed and used in a much wider range to maintain and ensure such quality standards. Such modules have to be cooperative tools which assist the medical personnel in its daily routine tasks.

In the following I will explore a pathway of research in AIM starting with knowledge capturing activities and ending with the need for new AIM systems which

utilize knowledge to guarantee best practice in a medical world which is data-intensive, dynamically changing, and devoted to quality of care.

2 Capture the experts' knowledge

The initial motivation for research in AIM was the famous “knowledge is power” paradigm. The main goal was to capture as much knowledge as possible from experts, pack it into a computer system, add a little of inference, and finally have an excellent decision support system at hand.

My group started with work on an expert system in the field of primary medical care. Cooperating with a general practitioner we tried to capture the essential knowledge used in the daily practice. The first step was knowledge structuring, which resulted in a system for defining knowledge structures [4]. Today, one would use description logics [1], or concentrate on ontologies [2] and terminology [22]—fields of AI researchers are very active. Finding proper knowledge structures still remains a nearly unsolvable problem: we are looking for languages which are sufficiently expressive, formally well-defined and consistent (to support all kinds of verification and validation proofs—not to think about completeness). All these properties may make the knowledge expressed in such a system reusable. At the same time the knowledge structure should be intuitive to the user, i.e., even physicians should have a chance to read and comprehend the content of the knowledge base. Finally, inferences on such a structure should work efficiently in the large domain of medicine. A “nice” (unrestricted) formal structure offers all the advantages of expressiveness, consistency and ease of verification. Unfortunately, solving problems with the help of such structures tends to be NP-hard. All these conflicting requirements are good for having knowledge structuring activities and knowledge representation an open issue of research for the mid-term or even long future.

Returning to the domain of primary care much of the work we have done was devoted to represent nosological knowledge. There was no hope in the early eighties to reach a sufficient coverage of the basic knowledge units required. The progress here is clearly visible today. Electronic versions of, for example, the International Classification of Diseases [27] are available. Standardized nomenclatures and terminology servers provide a good starting point for representation tasks [23]. But we are still lacking medical knowledge bases with rich semantic relations. We find long lists of concepts, but the relations between these medical concepts are lacking. Most such projects (e.g., [24]) got stuck and even today the representation of general medical knowledge which is usable for tasks beyond statistics seems to be an effort too demanding.

In the next step we limited ourselves to a more restricted domain: rheumatology. The resulting expert system ESDAT [5] did its work properly in the limited space

LISP gave us those days. The reasoning capabilities were primarily a result of the hierarchical knowledge structure supporting efficient pruning, and a clever focusing mechanism which allowed to concentrate on a few disease hypotheses.

However, it became evident that the main problem of such an expert system is its lack of background knowledge. In rheumatology one has to know a lot about anatomy. Consequently we have deepened our knowledge base by adding an anatomical layer. This second layer contained all the basic knowledge about the musculo-skeletal system. All the concepts of general, systematic and functional anatomy were represented together with an extensive set of relations defining causal pathways between joints, bones, tendon insertion points, muscles, nerves and elementary movements. The reasoner combined generic disease descriptions about rheumatological disorders with this detailed knowledge of the underlying anatomy. This combination of causal (finding affected body structures and impaired functions) and associative reasoning (relating findings to disease hypotheses) resulted in a very deep but focussed exploration of the body regions [6,8]. This gave impressive results due to the complete and relation-rich representation of the anatomy of the musculo-skeletal system.

This project and several others (Village Health Worker [21], RADIO [7]) followed one major paradigm: acquire as much knowledge from the expert as possible and make it available for decision support to less experienced persons. This did work in the research setting, but why didn't it found its way to the daily routine? There are many possible answers. I would like to give a few:

- computers were rare and expensive. Most systems have been developed using specialized hardware and software like Lisp-machines. It was obvious that such machines would hardly find their way to a clinic or a physician's office,
- work concentrated on AI methodology and rarely on system integration aspects. Expert systems were designed as stand-alone applications. If they found their way to a clinical setting they were installed on a PC which was mainly used by other tasks, like word processing, data acquisition, and statistical analysis,
- use of the expert system consumed an enormous amount of time due to the need of entering all relevant data by hand. The lack of electronic patient records and the missing integration with data bases, if available, put the untolerated burden of extensive data acquisition to the user. As a result such systems were experienced as being too time-consuming.

On the positive side the largest benefit was achieved by the medical experts working in the projects. The knowledge structuring and the knowledge acquisition process during the expert system development forced them to (re-)organize their knowledge. They got new insights to their domain due to the requirement of a clear structuring. This was clearly seen as a large benefit by the experts.

AI research has done some initial steps in medical knowledge-based systems, but

I am convinced that we are still at infancy. We got some insights about the appropriateness of specific knowledge modeling techniques and about the advantages of deep knowledge structures [12], but extensive research is necessary to create medical knowledge-bases which are rich on semantic relations, efficiently usable, and reusable for future tasks and projects.

3 Data-intensive systems: Reduce the information overload

In the early years of AIM we concentrated on experts' knowledge capturing. Data about the patient had to be entered into consultation systems manually. Quite the opposite is the situation we envisage now in the modern health-care setting: data are collected electronically, patient data bases store enormous amounts of data, most patient records are electronic today. Especially the modern equipment of intensive care units (ICUs) collects large sets of data from patient monitors every second.

A severe problem in ICUs today are false alarms resulting from the rigid alarming ranges of monitors [15]. Too many false alarms result in insensitivity of the medical personnel. More severe, ranges are chosen which rarely give alarms—a situation far away from the initial intention of careful monitoring. This situation asks for new methods of intelligent alarming: alarms are to be produced taking into account the whole situation of the patient by combining different data streams, integrating data from several monitors, and looking how data values have evolved over time. This is a great challenge for the future. It will require monitoring equipment to become more intelligent. To my view this can be achieved only by integration of devices and by adding knowledge to the monitoring software.

In the ICU physicians are confronted with dozens of values from each patient. Selecting appropriate patient management activities requires extensive interpretation of these large data sets. Intelligent data analysis (IDA) research [14] focuses on AI methods able to support this task of handling large amounts of heterogenous data. Data abstraction tries to find a high level description of the many elementary data items of a patient.

We used data abstraction to support a critical task: ventilator management in the neonatal ICU. VIE-VENT [18] combined information from different monitors to get a comprehensive (abstract) picture of the respiratory situation of the patient. This resulted in recommendations for changes in the ventilator settings if the situation was unsatisfying. This is quite a complex task which needs a lot a experience before becoming expert. Clinical expertise together with basic textbook knowledge about ventilation was built into VIE-VENT to support this patient management task.

One of the main advantages of AI systems, like VIE-VENT, are their ability to

maintain a continuous picture of the development of parameters over time. Most of the time clinicians just see snapshots of the overall situation of a patient by looking at the monitors. Patient data management systems (PDMS) provide more information, but usually in spreadsheet form with low resolution (e.g., one value of each parameter every 15 minutes). As a result the PDMS is an archive for a sequence of snapshots of the data situation of a patient. But what is needed for supporting decisions about proper therapeutic actions is a summary about the development of parameters in the recent past. Temporal data abstraction is able to provide essential methods for abstracting the development of parameters over time. It produces abstract descriptions of how certain parameters evolve over time. In VIE-VENT we used four different trends to abstract the time course—from a very short term trend (one minute) to a long term trend (three hours) [17]. These trends were utilized to control the recommendation of ventilator setting changes.

The temporal parameters used may vary depending on the requirements of the application. It may include the observation of stability, periods of continuous increase, the detection of repetitions and cyclic patterns. Temporal reasoning [13] has a lot of flexibility in adapting to such requirements—both in high-frequency domains like ICU monitoring where data are collected every second and in low-frequency domains where we look at the change over years. The more we are going to collect and store data, the more we need methods which support us in comprehending the complex data picture, both by combining the many different parameters and by abstracting parameters over time. Ultimately, we want our AI system to give us, e.g., a short summary on what has happened in the last 60 minutes.

An alternative method for easy comprehension of complex data is data visualization. Metaphor graphics is a method which uses graphical objects for visualizing different aspects of complex data sets. In VIE-VISU [11] we used shapes and color to depict 13 different parameters within one metaphorical graphical object. By showing multiples (sequences of these objects) we were able to visualize stability and change of the situation of a patient over 24 hours. Such visualization systems will get increased importance the more data we are going to collect.

A further step in intelligent data interpretation is the automatic control of medical devices. Automated FiO₂ supply controlled by continuous SpO₂ readings from pulseoximetry is such an example we are currently working at [25]. The main obstacle for such closed-loop systems is the large number of artifacts we receive from the monitoring equipment. Sophisticated data validation and data abstraction methods [9,20] form an essential prerequisite for fully automated data interpretation. However, we see quite substantial improvements in the robustness of monitoring systems which gives new opportunities for biomedical technology utilizing intelligent data analysis.

4 Ways to improve quality of care

Quality of care was always a key issue. But today it is discussed more intensively. We do expect to have excellent standards and best practice guidelines everywhere. Knowledge-based modules are superb tools able to ensure such quality of care if we have them integrated into the daily routine. Unfortunately, as argued in section 2, most of the knowledge-based systems did not find their way to daily routine use.

One successful example we have developed is an AI module for calculating the daily nutrition needs of newborn infants with low birth weight. VIE-PNN [10,16] use the nutrition standards of a neonatal intensive care unit to compose the components of a parenteral nutrition solution based on the daily requirements and the prescription of the previous day. It encodes the clinical practice rules of expert neonatologists together with textbook knowledge about nutritional needs. Being in daily routine use now for several years at two neonatal ICUs in Vienna VIE-PNN ensures and maintains the high quality of parenteral nutrition at these ICUs.

Essential for the success of this knowledge-based module are a few facts which we should keep in mind on our way to deliver AIM applications:

- the system is time-saving to the users. Physicians save approximately one hour per day if they use VIE-PNN. In addition the error rate has reduced considerably,
- VIE-PNN is a specific module integrated into the patient data management system of the clinic. Its use is part of the daily routine use of computer software for patient management,
- it encodes clinical knowledge and it is able to explain its recommendations based on these clinical rules. Interestingly, these explanations are used by novice users during the first days only. After finding the system trustworthy and having learned how the rules of prescription are applied, explanations are rarely needed,
- the system maintains continuity in the prescription of nutrition components with the ability of dynamic adaptation to the daily individual needs of a patient.

VIE-PNN is an example of a successful system in a narrow domain. It ensures the quality of care by applying clinical standards for daily patient management actions. We need many more such systems which actively maintain, ensure and control best practice standards, tailored to the need of the individual patient. A promising path is the work on automated support to guideline-based care [3,26]. Clinical guidelines are a set of schematic plans [19] for the management of patients showing a particular clinical condition. Computerized guidelines represent clinical expertise about patient data required, data interpretation, standard therapeutic actions and methods for revising a therapeutic plan if necessary. Again we see the need for extensive work in knowledge representation, even with a quite formal flavor. But in contrast to early approaches of decision support, procedural and temporal aspects form an

essential part. The resulting plans represent guidelines open for dynamic adaptation to the needs of an individual patient.

There is a strong need for methods improving and maintaining quality of care. AI modules embedded into the daily working environment of physicians and nurses have a great potential to contribute to these quality standards. They have the chance to widely disseminate best practice guidelines. But these AIM modules have to be implemented, verified and validated very carefully. If approved, they must be integrated smoothly into the daily-use environment of our care providers.

5 Conclusion

The last 20 years of AIM research and development show a path from knowledge-intensive systems, which try to capture the essential knowledge of experts in a KBS, to data-intensive systems available today. The enormous amount of information available in our networked world and the enormous amount of data collected daily with our modern medical equipment leads us to new tasks to cope with the information overload. We need knowledge to be able to see the essentials in this large bunch of information and data. Knowledge-based modules are essential tools which should help us to find out what is truly relevant and to give us a comprehensive picture of the situation of the patient. Most clearly, time is a prominent feature of our life. Temporal reasoning is a fruitful method to react to changes which appear over time. This gives our AIM systems the chance of dynamic behavior.

The examples given in this paper show some steps on this path. We are now in a situation where data-intensive systems have entered the health care domain. But this does not mean that data-intensive systems have replaced knowledge-intensive systems. On the contrary, knowledge-based systems are needed to make use of all these data available and to help us to cope with the information explosion.

Acknowledgements

We greatly appreciate the support given to the Austrian Research Institute of Artificial Intelligence (ÖFAI) by the Austrian Federal Ministry of Education, Science and Culture.

References

- [1] F. Baader, Logic-Based Knowledge Representation, in M.J. Wooldridge and M. Veloso, eds., *Artificial Intelligence Today* (Springer, Berlin, 1999) 13–42.
- [2] B. Chandrasekaran, J.R. Josephson, V.R. Benjamins, What Are Ontologies, and Why Do We Need Them?, *IEEE Intelligent Systems* 14 (1999) 20–26.
- [3] J. Fox, N. Johns, A. Rahmanzadeh, Disseminating Medical Knowledge: The PROforma Approach, *Artificial Intelligence in Medicine* 14 (1998) 157–182.
- [4] W. Horn, W. Buchstaller, R. Trappl, Knowledge Structure Definition for an Expert System in Primary Medical Care, in *Proceedings of the 7th International Joint Conference on Artificial Intelligence (IJCAI-81)*, (University of British Columbia, Vancouver, Canada, 1981) 850–852.
- [5] W. Horn, ESDAT – Decision Support for Primary Medical Care, in J.H. van Bommel, et al., eds., *MEDINFO 83* (North-Holland, Amsterdam, 1983) 484–487.
- [6] W. Horn, K. Ammer, Causal and Anatomical Reasoning in a Two-Level Medical Expert System, in T. O’Shea, ed., *Proc. of the 6th European Conference on Artificial Intelligence (ECAI-84)* (Elsevier, Amsterdam, 1984) 61–64.
- [7] W. Horn, H. Imhof, B. Pfahringer, E. Salomonowitz, A Radiological Expert System for the PC – Design and Implementation Issues, in J. Fox, et al., eds., *AIME 87 – European Conference on Artificial Intelligence in Medicine* (Springer, Berlin, 1987) 169–176.
- [8] W. Horn, Utilizing Detailed Anatomical Knowledge for Hypothesis Formation and Hypothesis Testing in Rheumatological Decision Support, *Artificial Intelligence in Medicine* 3 (1991) 21–39.
- [9] W. Horn, S. Miksch, G. Egghart, C. Popow, F. Paky, Effective Data Validation of High-Frequency Data: Time-Point-, Time-Interval-, and Trend-Based Methods, *Computers in Biology and Medicine* 27 (1997) 389–409.
- [10] W. Horn, C. Popow, S. Miksch, A. Seyfang: Quicker, More Accurate Nutrition Plans for Newborn Infants, *IEEE Intelligent Systems* 13 (1998) 65–69.
- [11] W. Horn, C. Popow, L. Unterasinger, Metaphor Graphics to Visualize ICU Data over Time, in R. Bellazzi and B. Zupan, eds., *Intelligent Data Analysis in Medicine and Pharmacology (IDAMAP-98)*, Workshop Notes of the ECAI-98 Workshop, (Brighton, UK, 25 August 1998) 76–81.
- [12] E. Keravnou, ed., *Deep Models for Medical Knowledge Engineering* (Elsevier, Amsterdam, 1992).
- [13] E. Keravnou, Temporal Reasoning in Medicine, *Special Issue on Temporal Reasoning in Medicine, Artificial Intelligence in Medicine* 8 (1996) 187–192.
- [14] N. Lavrac, E. Keravnou, B. Zupan, Intelligent Data Analysis in Medicine, in A. Kent et al., eds., *Encyclopedia of Computer Science and Technology*, Vol.42 (Dekker, New York, 2000) 113–157.

- [15] S.T. Lawless, Crying Wolf: False Alarms in a Pediatric Intensive Care Unit, *Critical Care Medicine* 22 (1994) 981–985.
- [16] S. Miksch, M. Dobner, W. Horn, C. Popow, VIE-PNN: An Expert System for Parenteral Nutrition of Neonates, in *Proceedings of the Ninth Conference on Artificial Intelligence for Applications (CAIA-93)* (IEEE Press, New York, 1993) 285–291.
- [17] S. Miksch, W. Horn, C. Popow, F. Paky, Therapy Planning Using Qualitative Trend Descriptions, in P. Barahona, et al., eds., *Artificial Intelligence in Medicine, Proc. AIME-95* (Springer, Berlin, 1995) 197–208.
- [18] S. Miksch, W. Horn, C. Popow, F. Paky, Utilizing Temporal Data Abstraction for Data Validation and Therapy Planning for Artificially Ventilated Newborn Infants, *Artificial Intelligence in Medicine* 8 (1996) 543–576.
- [19] S. Miksch, Plan Management in the Medical Domain, *AI Communications* 12 (1999) 209–235.
- [20] S. Miksch, A. Seyfang, W. Horn, C. Popow, Abstracting Steady Qualitative Descriptions over Time from Noisy, High-Frequency Data, in W. Horn, et al., eds., *Artificial Intelligence in Medicine, Proc. AIMDM-99* (Springer, Berlin, 1999) 281–290.
- [21] G. Porenta, B. Pfahringer, M. Hoberstorfer, R. Trappl, A Decision Support System for Village Health Workers in Developing Countries, *Applied Artificial Intelligence* 2 (1988) 47–63.
- [22] A.L. Rector, J.E. Rogers, P. Pole, The GALEN High Level Ontology, in *Proc. Medical Informatics in Europe (MIE)* (Copenhagen, Denmark, 1996).
- [23] A.L. Rector, Terminology and Concept Representation Languages: Where are We? *Artificial Intelligence in Medicine* 15 (1999) 1–4.
- [24] O. Senyk, R.S. Patil, F.A. Sonnenberg, Systematic Knowledge Base Design for Medical Diagnosis, *Special Issue on Causal Modeling, Applied Artificial Intelligence* 3 (1989) 249–274.
- [25] A. Seyfang, S. Miksch, W. Horn, M.S. Urschitz, C. Popow, C.F. Poets, Using Time-Oriented Data Abstraction Methods to Optimize Oxygen Supply for Neonates, Austrian Research Institute for Artificial Intelligence, Technical Report TR-2001-03, Vienna, Austria, 2001.
- [26] Y. Shahar, S. Miksch, P. Johnson, The Asgaard Project: A Task-Specific Framework for the Application and Critiquing of Time-Oriented Clinical Guidelines, *Artificial Intelligence in Medicine* 14 (1998) 29–52.
- [27] World Health Organization, *The International Statistical Classification of Diseases and Related Health Problems (ICD-10)* (WHO, Geneva, 1992-1995), available at www.who.int/whosis/icd10/.