Clinical knowledge modeling: An essential step in the digital transformation of healthcare

Eric Edelman,[1,2,*] Fabian Tijssen,[3,4,8] Popke Rein Munniksma,[4,8] Wim Bast,[4,8] Harold ten Bohmer,[4,8] Nicole van Eldik,[2,5] Marieke Spreeuwenberg,[1] Wolfgang Buhre,[6] and Frits van Merode[1,7]

[1]CAPHRI School for Public Health and Primary Care, Maastricht University, 6211 LM Maastricht, the Netherlands

[2]Healthcare Innovation Lab, Maastricht University Medical Center, 6229 HX Maastricht, the Netherlands

[3]Department of Anesthesiology and Pain Medicine, Maastricht University Medical Center, 6229 HX Maastricht, the Netherlands

[4]Open Walnoot, 6231 EB Meerssen, the Netherlands

[5]Department of Pediatrics, Maastricht University Medical Center, 6229 HX Maastricht, the Netherlands

[6]Department of Anesthesiology, Division of Vital Functions, University Medical Center Utrecht, 3584 CX Utrecht, the Netherlands

[7]Maastricht University Medical Center, 6229 HX Maastricht, the Netherlands

[8]These authors contributed equally

*Correspondence: e.edelman@maastrichtuniversity.nl

Received: June 5, 2024; Accepted: October 17, 2024; Published Online: October 19, 2024; https://doi.org/10.1016/j.xinn.2024.100718

Citation: Edelman E., Tijssen F., Munniksma P.R., et al., (2024). Clinical knowledge modeling: An essential step in the digital transformation of healthcare. The Innovation **5(6)**, 100718.

INTRODUCTION

The current healthcare system is labor intensive and relies heavily on the skills, knowledge, experience, and direct involvement of its individual human employees. It has thus acquired a reputation for being difficult to digitize and automate, making it challenging to achieve increases in productivity and efficiency. What makes automating or otherwise digitally supporting this work so difficult? The problem is that, to be implemented in automation, knowledge must be both structured and explicit, which is often not the case in the medical domain. This paper describes our method for the explicit structuring of medical knowledge in clinical knowledge models (CKMs) (Figure 1). These may be utilized as the building blocks for automation and decision support in medicine.

Much of human medical knowledge is currently locked in explicit but mostly unstructured documents, such as protocols, guidelines, and scientific publications. Compounding the problem, while these sources form the basis for good practice, they are continuously combined with the implicit knowledge of the individual healthcare professionals that apply them. Implicit knowledge is different than the additional knowledge professionals may have that simply has not been recorded, such as through on-the-job training. In contrast, implicit knowledge is knowledge that a holder is unable to state on their own. Polanyi's classic definition of "tacit knowledge" is closely related: we can know more than we can tell.[1] Examples of such implicit knowledge are knowledge gained through experience and personal development. This knowledge often varies between professionals. Despite being essential for providing responsible care, this implicit body of knowledge has so far been largely overlooked in attempts to provide automation and decision support to healthcare professionals.

A DECISION-DRIVEN APPROACH

Creating a CKM starts with a decision that is central to the knowledge being modeled. Such a decision might be whether a patient needs surgery, which medication to prescribe and in what dose, whether a patient may be discharged, etc. Which information is needed to take this decision? As stated above, only part of the required knowledge to determine this is explicitly present. The implicit knowledge that plays a role in caring for a group of patients needs to be elicited and combined with the formal, explicit knowledge. This usually requires multiple sessions with an expert group of relevant healthcare professionals and clinical modelers.

Applying such a decision-driven approach to elicit and organize knowledge offers significant advantages over the more common data- or process-driven approaches. (Note that these approaches do not cover probabilities based on partial information, such as described by Jaynes' information theory.[2])

1. In a data-driven approach, experts might be asked "what do you need to know about a patient to provide care?". From a technical point of view, this seems like a logical question, but it is a dangerous one. Asked by itself, it is an incredibly hard question to answer and leads to four risks:
   a. Experts might omit information by not recalling the need for it.
   b. Experts might omit information by subconsciously oversimplifying their work. This might happen when they are unaware of the complexity of their actions and their implicit knowledge.
   c. Experts might have trouble offering the information because they regard their work as more complex than it is. They may overestimate

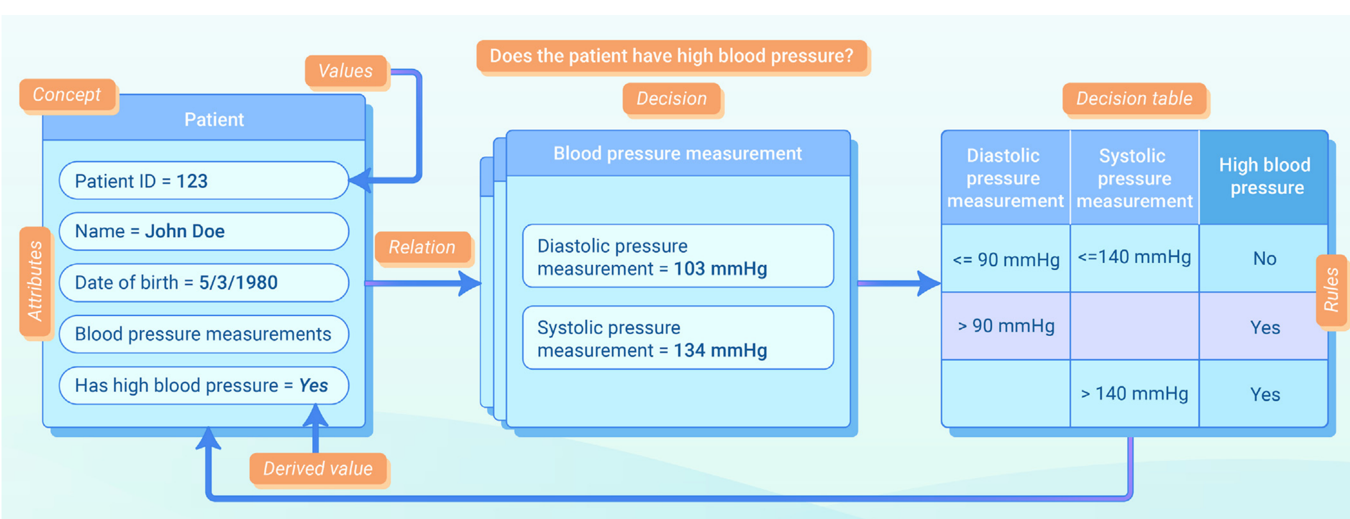

Figure 1. Conceptual modeling explained The most important terminology in conceptual modeling is explained using a simple example. This example demonstrates how a derived value (whether the patient suffers from high blood pressure) is determined using attribute values and decision rules.

the role of their unique experience and the variation between patients.

d. Finally, the question may also be interpreted as "what do you *want* to know about a patient?", leading down a path of including superfluous information that does not actually influence any later decisions. This may be specifically tested by asking an expert "what does that information affect down the line? Think of a patient where this information caused the plan to change."

Focusing on the decisions the experts need to take and only then listing the data required for these decisions provides a framework for the discussion. Using this approach allowed us to develop a prototype model on the treatment of esophageal cancer, specifying the data that would be needed to be exchanged between centers of care. Three days of expert sessions led to 90% of the data specifications that had previously taken 3 years of discussion in the region to agree on.

2. In a process-driven approach, data might be connected to a process step. Consider two clinical pathways that both make use of the same blood test to determine the next step in their processes but at different moments in time. In current pathway modeling (and real-life execution of these pathways), it is a common occurrence that the two independent pathway teams never realize this duplication. In contrast, decision-driven modeling runs on a continuously current representation of the state of the patient. If the result of a test is known, this is immediately available to each CKM that makes use of it.

   The aim is to integrate all of the knowledge relevant to a patient. This breaks the barriers that now often exist between independently developed clinical pathways and could prove especially valuable in the management of comorbidity and subsequent complications. Juggling the knowledge from multiple cross-influencing guidelines across multiple medical specialties is extremely difficult to do accurately, forming an Achilles' heel of many clinical decision support systems.[3] This may lead to suboptimal treatment and, in some cases, an increased risk of adverse incidents.
3. Consequently, in CKMs, decisions are not considered steps in a process, like they are in procedural modeling. Instead, all CKMs are completely declarative representations of knowledge, in which decisions form coherent networks of related knowledge that can easily be integrated. Each decision table (containing a specification of the required input data, the decision rules, and the possible outcomes) can be evaluated independent of each other. If the needed variables are set and the logic of the table is complete, a decision can be made. This reduces complexity because a decision can be modeled separately from all of the rest of the model.

Once a CKM is complete (a simple example is displayed in Figure 1), it can be used to generate usable software tailored to its intended purpose. The resultant software can be thought of as a service: waiting to receive input, continuously evaluating all decision rules at once, and returning output as soon as the conditions for a decision rule are met. The service may be integrated with existing medical software, but in many cases, the CKM service will not be able to collect all necessary data by itself. The user must input the information needed as it becomes available; for example, while talking to a patient. Preferably, this would take place in the existing user interface (UI) of the software to which the user is accustomed, such as an electronic health record system. A prerequisite for this scenario is that the existing software supports a dynamic UI. Each change of the content of a field in the UI (such as a patient's medication, history, a measurement, etc.) calls the CKM service. The service can take any decision for which all information is available and return the result. Currently, medical software often still presents input fields for each piece of information that might need to be entered sometime in the care process. This is part of the reason why healthcare professionals have experienced an increase in administrative burden while caring for patients, interfering with their face-to-face time with patients and contributing to clinician burnout. In one study, nearly 87% of clinicians mentioned excessive data entry requirements as a concern about their electronic health record's design.[4] It is often unclear whether certain information is required and what consequence its value will have later on in the process.

MANAGEABILITY AND EXPLAINABILITY

Medicine was an early testing ground for electronic expert systems. Despite having advantages for the explainability of results, it soon became clear that neither formal logic (standard rules of inference) nor production rules (conditional sentences with an "if" part and a "then" part) would be suitable models for capturing and managing the complexity of medical knowledge, which often include many-to-many relationships between concepts. In contrast, many modern AI technologies are not knowledge driven like expert systems but could be considered to produce "data-driven knowledge."[5] In simpler forms of machine learning, this results in models that are fundamentally based on statistics, allowing human interpretation. But in more complex deep learning, such as neural networks, interpretation is far less intuitive. This has serious consequences for the way users view AI systems regarding reliability and trustworthiness. In medicine especially, the explainability of AI advice is considered a key factor to gaining acceptance for human-machine collaboration.

Declarative modeling as employed in CKMs offers a more manageable way of organizing knowledge than procedural logic, even when faced with vast and complex bodies of knowledge, including many-to-many relationships between concepts. CKMs may be unraveled to present only the relevant elements for that moment. This offers three advantages. First, during the development and maintenance of models, it enables developers to focus on specific areas at a time, showing the possible interdependence with other parts of the model but hiding anything else. Second, it enables the modeling tools to check the model for missing rules and inconsistencies. Third, the final CKM service can leverage this unraveling by not only delivering a result but also an explanation of how it got there.

CONCLUSION

In summary, CKMs hold the potential to enable further digitalization of healthcare. Structuring and explicitly storing knowledge that has previously remained out of view makes it manageable and shareable and enables its use for automation.

Automation and decision support are regarded as important building blocks in the digital transformation of healthcare. The wish for this transformation is partially driven by a fear of scarcity, sparked by the growing divide between demand and the capability to deliver in our healthcare system. Although it is expected that digital transformation will have a significant positive impact on the capacity of healthcare workers and the system as a whole, we would be underselling its benefits by only seeing it in the light of forced change. We consider well-executed automation of "rational intelligence" (taking the obviously correct decision based on available knowledge and information) the next step in increasing medicine's quality and safety and increasing its capacity for providing care. Healthcare professionals should be unburdened from the mundane decisions they are currently forced to take daily. When all information and knowledge available lead to a relevant decision being able to be taken rationally, the resulting advice should be available automatically and not cost any human cognitive capacity. That precious cognitive capacity is then fully available for the parts of medicine that will never be digitalized nor automated: thinking about cases that do not fall within the bounds of current knowledge, employing emotional intelligence to understand a patient's condition in a way a machine never can, and building meaningful relations with patients.

REFERENCES

1. Polanyi, M. (1967). The Tacit Dimension (Anchor Books).
2. Jaynes, E.T. (1957). Information Theory and Statistical Mechanics. Phys. Rev. ***106***(4): 620–630. https://doi.org/10.1103/PhysRev.106.620.
3. Fraccaro, P., Casteleiro, M.A., Ainsworth, J., et al. (2015). Adoption of Clinical Decision Support in Multimorbidity: A Systematic Review. JMIR Med. Inform. ***3***(1): e3503. https://doi.org/10.2196/medinform.3503.
4. Kroth, P.J., Morioka-Douglas, N., Veres, S., et al. (2019). Association of Electronic Health Record Design and Use Factors With Clinician Stress and Burnout. JAMA Netw. Open ***2***(8): e199609. https://doi.org/10.1001/jamanetworkopen.2019.9609.
5. Dubois, D., Hájek, P., and Prade, H. (2000). Knowledge-Driven versus Data-Driven Logics. J. Log. Lang. Inf. ***9***(1): 65–89. https://doi.org/10.1023/A:1008370109997.

DECLARATION OF INTERESTS

F.T., P.R.M., W. Bast, and H.t.B. are shareholders of Open Walnoot, a company that specializes in providing the techniques and solutions described in this paper to various clients within the healthcare sector.