Mining Medical Records: A Case for Artificial Intelligence in Health Systems

By Kapila Monga, MS

According to the Oxford Living Dictionary, artificial intelligence (AI) is "the theory and development of computer systems able to perform tasks that normally require human intelligence, such as visual perception, speech recognition, decision-making and translation between languages." There is a lot of debate underway about whether AI can ever match or precisely mimic the intelligence a human being possesses.[1] But the chances that this technology will eventually be perfected—and soon, at that—are growing by the day.

AI has been around in some form since the 1940s, when the first digitally programmable computer was invented based on principles abstracted from mathematical reasoning. The field of AI research was founded in 1956 at a workshop in Dartmouth College. Investment and interest in AI boomed in the first decades of the 21st century, when machine learning was successfully applied to many problems in academia and industry.[2] So artificial intelligence is not new. However, with tech companies like Facebook, Google, and Amazon investing significantly in AI, and around 3,108 start-ups currently operating in the AI space, this seems like an idea whose time has come.[3]

The Case for a Medical Records Mining AI Agent Named 'Aida'

Taking Oxford's definition of AI as a basis, let's assume that health IT developers were able to develop an engine that runs continuously in the background of an electronic health record (EHR) system and guides physicians, nurse practitioners, and health information management (HIM) professionals on several clinical and operational aspects.

Let's call this program "Aida." Aida would have a "brain" of its own in the form of a cognitively intelligent text, audio, and video processing engine. In this case, cognitively intelligent means having the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, and learn quickly from experience. This brain will enable Aida to give recommendations to physicians, nurse practitioners, fraud, waste, and abuse inspectors, and several other stakeholders. The limit on the kind of insights and recommendations generated by Aida is only set by the maturity of the cognitively intelligent engine. In an ideal world one could at least expect Aida to give recommendations on the following:

1. **Medical documentation completeness:** Aida would advise physicians and HIM professionals on any possible missing ICD-10 codes, procedure codes, and HCPCS codes that should be added to the record. The recommendations for a particular patient can be generated when the provider is accessing the record. This way there will be an opportunity to both validate the recommendation and act on it (if determined necessary) without an additional overhead on the users.
2. **Medical necessity determination:** Before a procedure is administered, Aida could scan all available data and pre-existing conditions for a particular patient and help determine if the procedure does or does not seem medically necessary for any reason.
3. **Care management recommendation:** At the time of patient discharge from an inpatient setting, Aida would alert a provider on the likelihood of a patient being readmitted within 30 days and possible reasons, giving the opportunity to prevent 30-day hospital readmissions. Aida would also be able to predict if the patient is likely to stop taking recommended therapies post-discharge from the hospital, giving an opportunity to the provider to prevent that from happening.
4. **Treatment options:** As of now this seems to be a slightly distant dream, but in an ideal world Aida would be able to scan patients' data from an EHR and medical research journals and recommend possible treatment options to the healthcare provider.

Healthcare is a complex field and it seems unlikely that any computer system will be able to fully comprehend those complexities with existing technologies. Any AI engine should at best serve as an aid to providers and HIM professionals. It can't replace them—it can only speed up identification of any possible gaps. The final decision on whether the recommendations by an AI engine are worth considering will still be in the hands of the healthcare professional. However, as the AI engine matures and becomes more accurate, the reliance on healthcare professionals' validation can be minimized.

Prototype for Aida's Records Mining Engine

An AI engine has four different parts that need to work together to form a real AI application. The author's point of view on what those different parts are for Aida, and how they work, includes :

1. **A cognitively intelligent processing engine:** This is the core engine that would do all the data crunching, run the machine learning algorithms, and do the recommendation generation. Technology is the easy part for this piece since there are several machine learning AI platforms available from Microsoft, IBM, and Python, etc. What's difficult to get is the required clinical knowledge and experience integrated in algorithm development. There is a huge gap between the statistical and medical acumen of clinical experts and data scientists, respectively.
2. **Human interface:** Since Aida would be interacting with physicians, nurse practitioners, and HIM professionals, the conversation maps need to be appropriately designed so that Aida can explain and understand humans in both clinical and non-clinical terms. This is easier said than done. Human conversations typically can go in any direction—starting from simple "What is" questions to more complicated "What if" and "So what" types of questions. Given the complexity of healthcare, there can't be the same answers to the same questions all the time—advances in medical technology, changes in regulations, and Centers for Medicare and Medicaid Services (CMS) guidelines all would have an impact on the "truthfulness" of an answer at any given time. For Aida to be able to successfully interact with humans, conversation maps need to be designed using principles from design thinking and in consultation with anthropologists, UI designers, software programmers, and clinical experts all in one room.
3. **Self-learning engine:** Complete self-learning is still a distant dream, and even further away in healthcare because of the complexities of the field. Assisted self-learning is something that AI systems can aim for today. Through appropriately designed learning databases used to capture feedback on the accuracy of Aida, Aida can become intelligent and accurate as time passes. The learning database is like a performance score card that will enable the collection of feedback on the accuracy of results of Aida's performance (learning) from HIM professionals (assisted learning) in a structured format. The scientifically designed structure of the learning database will also enable incorporation of the feedback back into a machine learning algorithm in an automated manner, thereby increasing the accuracy of the algorithm.
4. **Triggering/invocation mechanism:** The Aida engine will keep on running at the back-end of the EHR system. In other words, it will be like a crawler program that runs 24/7. As soon as Aida encounters an anomaly it will display an alert on the screen of the HIM professional who is logged in and viewing the EHR. Alternatively, in case a HIM professional encounters a record they find suspicious, Aida can be invoked on demand through a click of a button in the EHR application or through the voice prompt "Aida."

Healthcare has a direct impact on human lives, so putting AI to action calls for an interdisciplinary and collaborative effort.

Notes

[1] Marcus, Gary. "Artificial Intelligence Is Stuck. Here's How to Move It Forward." *New York Times*. July 29, 2017.

[2] Moor, James. "The Dartmouth College Artificial Intelligence Conference: The Next Fifty Years." *AI Magazine* 27, no. 4 (2016): 87-89.

[3] AngelList. "Artificial Intelligence Startups."

Kapila Monga (kapila.monga@gmail.com) is a healthcare analytics professional specializing in advanced analytics, data science, and consulting for healthcare and life sciences customers. She currently works with Cognizant Technology Solutions in their healthcare analytics practice. The opinions expressed in this article are the author's own

and do not reflect the views of the author's employer or of any other corporate entity. The name Aida was used as a theoretical example namesake for an AI entity.

Article citation:
Monga, Kapila. "Mining Medical Records: A Case for Artificial Intelligence in Health Systems" *Journal of AHIMA* 88, no.10 (October 2017): 54-56.