Risk Matters: Maximizing the Benefits and Recognizing the Risks of AI in Healthcare

By Justin Joy, JD, CIPP

Various applications utilizing some form[1] of artificial intelligence ("AI") have been in place in healthcare for decades. The utility and utilization of AI have increased dramatically as the technology continues to develop and improve. A 2019 Harvard Business Review study estimated that AI applications for back-office activity save the industry $18 billion annually, noting that "activities that have nothing to do with patient care consume over half (51%) of a nurse's workload and nearly a fifth (16%) of physician activities." [2] With the dramatic improvement and availability of AI applications over the past half decade, such cost-saving estimates are likely more generous today.

Recent improvements in, and increased adoption of, generative artificial intelligence ("Gen AI") have reinvigorated imaginations on how AI can be leveraged to improve healthcare on

a wide scale. For example, prior to Gen AI, voice-to-text technology could automatically and instantly transcribe notes dictated by a provider. With the introduction of Gen AI, those voice recordings from a patient visit can be taken from unstructured text and adapted into a structured office visit note with conversational language.[3]

While categorically dividing AI healthcare applications into two likely-overbroad categories of clinical and administrative may be helpful for discussion and comprehension purposes, clinicians utilizing any AI application must be aware of the risk no matter how the technology is used. The risks of employing AI in any clinical application, such as assisting with diagnoses, should be evident on the surface. The challenge-which is well beyond the scope of this article-is to mitigate that risk in a meaningful way while not significantly diminishing the recognized efficiencies and other benefits of utilizing the technology for clinical purposes.

While not as obvious, the risks for administrative tasks should not be underestimated. Particularly with the increasing use of Gen AI, the benefit of the technology in quickly providing customized material unique to each patient, such as after-visit summaries specifically addressing points discussed during the visit, must be weighed against the possibility of error in that output upon which a patient may rely. Similarly, with the increased availability of AI technology for improving remote monitoring,[4] unchecked reliance on the technology could potentially lead to adverse results. Privacy and security concerns must also be addressed. For example, ambient clinical intelligence is a technology that "listens to" a conversation between a provider and patient and then automatically creates a clinical note based on the encounter. Those using the technology must understand if and whether any audio recording is maintained, and the security of the information collected during the visit utilized to generate the note.

Risk management resources continue to develop with the continued commercial proliferation and adoption of AI systems and applications. The National Institute of Standards and Technology (NIST) released the first version of the AI Risk Management Framework in January 2023 with the goal "to improve the ability to incorporate trustworthiness considerations into the design, development, use, and evaluation of AI products, services, and systems."[5] For organizations seeking adoption of a management system standard to structure how they address "the unique challenges AI poses, such as ethical considerations, transparency, and continuous learning," the International Organization for Standardization released ISO/IEC 42001:2023 in December 2023 to provide "a structured way to manage risks and opportunities associated with AI, [while] balancing innovation with governance."[6] While comprehensive and innovative, recognized organizational standards and frameworks consider the entire AI development lifecycle. A considerably more straightforward and focused approach can be followed for a medical practice seeking to find a starting place to address its AI risk.

Medical practices must first understand and identify where AI is used anywhere within the organization, including software or systems provided by outside vendors. Next, groups should identify any output created by AI or, relatedly, any data derived from AI processing.

Given that AI may be embedded into applications and not always apparent at the surface, IT staff or others familiar with the practice's software and systems should be involved in this identification process. Once AI applications and systems have been identified, the risk posed by the output or AI data should be assessed, with any application assisting in rendering medical diagnosis or judgments generally weighted as a potentially higher risk than AI processes geared toward administrative tasks. The assessment process should challenge how AI outputs are validated and how automation bias is mitigated. In other words, just because the process seems correct nine times in a row, that alone does not justify a presumption, without some check or control, that it will be correct the tenth time. Similarly, quality checks comparable to human-generated output should be utilized for administrative tasks. For example, just like human-transcribed dictation should be proofread for errors, text generated by an AI application should be subject to the same review process.

While AI in healthcare offers tremendous potential for improving patient care and operational efficiency, healthcare organizations must recognize and proactively manage the associated risks.

[1]. For an overview of certain types of AI-machine learning, natural language processing, rule- based expert systems, physical robots, and robotic process automation-relevant to healthcare applications, see Thomas Davenport and Ravi Kalakota, "The potential for artificial intelligence in healthcare," Future Healthcare Journal, Vol. 6, No. 2:94-98 (2019). Of course, as the technology continues to develop, new variations and combinations of AI are introduced.

[2]. Brian Kalis, Matt Collier, and Richard Fu, "10 Promising AI Applications in Health Care," Harvard Business Review (May 10, 2018).

[3]. Shashank Bhasker, Damien Bruce, Jessica Lamb, and George Stein, "Tackling healthcare's biggest burdens with generative AI," McKinsey & Company (July 10, 2023), available at ***********.mckinsey.com/industries/healthcare/our-insights/tackling-healthcares-biggest- burdens-with-generative-ai.

[4]. Shannon Flynn, "10 top artificial intelligence (AI) applications in healthcare," VentureBeat (Sept. 30, 2022), available at ********venturebeat.com/ai/10-top-artificial-intelligence-ai- applications-in-healthcare/.

[5]. ***********.nist.gov/itl/ai-risk-management-framework

[6]. *********** .iso.org/standard/81230.html