Harnessing AI to Enhance Clinical Documentation Improvement

How Cerner is working to ensure the ethical and responsible use of advanced technology

Artificial intelligence — including machine learning, natural language processing, deep learning, and robotics — is being rapidly integrated into every aspect of healthcare IT. From behind-the-scenes applications in revenue cycle management to the high-profile prescriptive analytics powering precision medicine, these technologies are fundamentally changing the way clinicians practice, and driving new revenue opportunities and efficiencies for health systems.

The field of clinical documentation improvement (CDI) is an especially promising area for AI, because it stands at the critical intersection of data, diagnosis, treatment, and revenue. Although it may not garner as much attention as robotic surgeries or algorithmic predictions about cancer clusters, AI-injected CDI offers transformative opportunities to maximize patient outcomes while improving clinician satisfaction and ensuring the highest degree of transactional accuracy, according to Tanuj Gupta, MD, Vice President of Cerner Intelligence, who leads the product management, data science, and engineering teams for machine learning and artificial intelligence.

"Clinical documentation today is mostly human-driven, with the physician and care team spending a significant amount of their time manually entering data based on patient encounters," Gupta recently told Healthcare IT News. "On top of this, reimbursement and quality requirements heighten the need for robust documentation. Clinical documentation serves as a necessary communication tool, ensuring continuity of care for the patient between venues and informing financial reimbursement and care quality for reporting purposes."[1]

Several macro trends have highlighted the importance of CDI for healthcare organizations over the past two decades, including the introduction of value-based payment models, new capitation strategies, mergers and acquisitions, more stringent quality measures, and ever-increasing scrutiny from regulators and business partners. Hospitals and health systems have, with varying success, invested in human resources to keep up with the demand for greater accuracy and evidence in documentation, but it's becoming clear that AI technologies will be required to increase the percentage of concurrent chart reviews, to improve the quality of documentation during encounters, and to review 100% of the charts rather than settling for a representative sample.

"Even before the advent of AI — and at least since CMS's transition to Medicare-Severity Diagnosis-Related Groups — it had become increasingly evident that these reviews needed to be addressed while the patient was in-house," said Beth McCauley, Manager/Lead Clinical Terminologist for Cerner Intelligence. "Bringing AI into the

"Clinical documentation today is mostly human-driven, with the physician and care team spending a significant amount of their time manually entering data based on patient encounters."

TANUJ GUPTA, MD | VICE PRESIDENT | CERNER INTELLIGENCE

process makes it a faster, easier, and less cumbersome procedure leading to better tracking, better capture, and more complete reviews than using only a human staff, who can only do so much."

McCauley is leading a team of terminologists at Cerner who design, build, and support algorithms using machine learning, natural language processing, and ontology to assist physicians on documentation improvement. Ideally, she said, AI can be leveraged to reduce manual review of documentation and coding, minimize the risk of denied claims, and shift CDI staff from "reviewing all" to exception-based review of complex claims.

Ethical concerns

But the rise of AI has not arrived without practical and ethical concerns. Skeptics, for instance, have raised objections to "black box" deployments, in which AI tools generate findings and recommendations without allowing clinicians to see the evidence and reasoning behind them. Others have pointed to ways in which AI algorithms can undermine patient privacy norms that have been developed and refined over the course of many years.

In the lead editorial of the February 2019 issue of the AMA Journal of Ethics, Michael J. Rigby — a PhD candidate in molecular neuroscience at the University of Wisconsin School of Medicine — noted that AI "creates a novel set of ethical challenges that must be identified and mitigated since AI technology has tremendous capability to threaten patient preference, safety, and privacy. However, current policy and ethical guidelines for AI technology are lagging ... [and] the medical community remains ill-informed of the ethical complexities that budding AI technology can introduce."[2]

Writing for the Brooking Institute, W. Nicholson Price II, Law Professor at the University of Michigan Law School, agreed with Rigby's assessment of AI's promises and risks. "Artificial intelligence is rapidly entering healthcare and serving major roles, from automating drudgery and routine tasks in medical practice to managing patients and medical resources," Price argued in November 2019. But as "developers create AI systems to take on these tasks, several risks and challenges emerge, including the risk of injuries to patients from AI system errors, the risk to patient privacy of data acquisition and AI inference, and more."[3]

Addressing the challenges

But even as these concerns were being articulated, Gupta, McCauley, and their colleagues at Cerner Intelligence were already at work on an ethical approach to infusing AI into clinical documentation workflows.

"We were telling the industry that we could assist its CDI efforts with artificial intelligence," said Gupta, whose team is developing predictive and prescriptive intelligence, ambient voice solutions, natural language understanding, and semantic interoperability to make technology in healthcare more intuitive. "But if we were going to do that, then how do we assure the industry that we can do it responsibly? How could we mitigate AI's risks while taking advantage of its opportunities?"

The result of this reflection and analysis is what the Cerner Intelligence team calls the "Six Principles of Responsibly Applying AI to Documentation and Coding." The principles seek to articulate the basic considerations that developers, vendors, healthcare providers, and users should insist on when applying AI to any processes or workflows within healthcare.

"This is something that hasn't necessarily been defined before in the industry," Gupta said. "But I think and hope that these six principles end up being reviewed, considered reasonable, and adopted by other companies, by other people in healthcare. It's an approach to AI not exclusive to clinical documentation, but AI in other fields as well."

The Six Principles

1. No black box deployments

It's not enough for an AI-supported application to recommend what it believes is an undocumented diagnosis. According to the first principle, AI developers should be transparent about the algorithm and the factors it uses when generating recommendations, and presenting the evidence in the same context [Figure 1].

Figure 1. No black box deployments

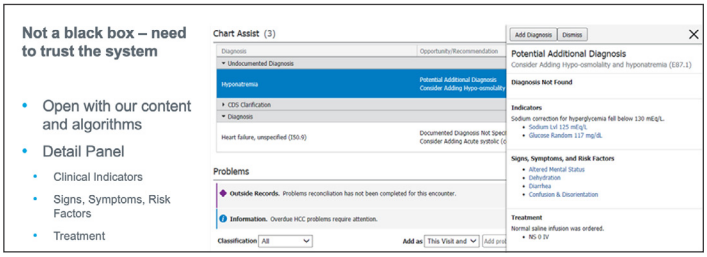

"Clinicians need to be able to trust the system they use. Having been in the field for a while, it was evident when we first started this project that there needed to be transparency, particularly with content. If providers don't know what that content is or how it was arrived at, it forces them to guess or make assumptions. And that's unethical and a real disconnect."

BETH MCCAULEY | MANAGER/LEAD CLINICAL TERMINOLOGIST | CERNER INTELLIGENCE

"Clinicians need to be able to trust the system they use," McCauley explained. "Having been in the field for a while, it was evident when we first started this project that there needed to be transparency, particularly with content. If providers don't know what that content is or how it was arrived at, it forces them to guess or make assumptions. And that's unethical and a real disconnect."

2. Leave medical decisions to the clinician

Responsible AI ensures that physicians retain decision-making power. It enables providers to consider the information presented, its rationale, and allows them to determine if it's acceptable or not. It permits human decision-makers to make an additional diagnosis based on available information, and if it makes sense to do so, gives them the freedom to include it in the patient encounter and to the notes.

"We often are asked by clients — when they see AI-based software like this — if we are telling the clinicians what to do," said Gupta. "Are we leading them? Who's accountable for the choices? We say it shouldn't be the software. Otherwise, it's the software that's practicing medicine. So, that principle becomes important. Once you state that the medical decision-making belongs to the professional, it defines your entire approach."

3. Incorporate clinical validation

Acknowledging the burden of proof is the foundation for the third principle of responsible AI. Given the potential fallibility of algorithms, clinicians should be able to validate diagnoses that the system may not have generated and dismiss system-generated findings that they feel are not warranted.

"Let's say a diagnosis for renal failure was found, but not all the clinical indicators and evidence of treatment were there," McCauley said. "Maybe there's supportive evidence that was not located in the encounter. Maybe there were signs for treatment, but it didn't necessarily meet all of the clinical indicators. Perhaps documented lab values don't suggest it's acute renal failure, and you want to call it to the attention of the physician or clinical documentation specialist so it can be reviewed. Either way, clinical information can be validated — and it's advantageous to validate the information while the patient is still in-house."

4. Make AI inherent in the workflow

Early AI applications weren't integrated into the daily workflows and main applications of clinicians, which required them to log into different tools, constantly switch windows, and review charts retrospectively. The fourth principle of responsible AI recommends that solutions be so tightly integrated that it feels a natural part of physicians' normal, daily workflows.

"Context switching is frustrating for clinicians," McCauley said. "AI recommendations should not feel like an added additional burden, but one that just fits into their workflows so that they really don't even notice that there's something new that they're looking at. This enables more attention on the patient, better decision-making, and reduced cognitive load."

5. Customization allowances

The fifth principle stipulates that there must be some room for content customization with any AI application. Subtle variations in practice are not necessarily unacceptable, especially when context matters.

"We must recognize that not all facilities are the same," said McCauley. " wA facility in one region may have slightly different standards compared to another facility in a completely different region. Some organizations develop their own internal guidelines or best practices, based on national standards, and they want different levels of lab values, for instance, to trigger different treatment. It's not a one-size-fits-all, and that's okay."

6. Enable a silent mode or evaluation period

Finally, given the uncertainty that clinicians and documentation specialists may harbor about AI, incorporating the ability to evaluate it using real data before rolling it into production is crucially important.

"If you have any doubt about how the AI component works and you want more proof, you should be able to try it out," said McCauley. "Run it in silent mode for an evaluation period. Be able to play around with it in the background so it's not disrupting anyone's workflow in your own environment, and see if it meets your needs."

Validating the basis

All six principles strive to validate the basis of AI findings and recommendations, Gupta says, and make it more helpful to clinicians and other users. Earning that trust is fundamental. "We have to be able to show our work and explain why the software is recommending what it does," he said. "And we would have to do that quickly because we never want it to be cumbersome for users."

Furthermore, he said, "these principles guide every design choice we make, from the language that we use in recommendations to the colors we apply to alerts, because we should never be directive or creating some bias toward action. They truly affect the design of the software from the bottom up."

In the end, McCauley adds, the responsible use of AI in coding and documentation is not to replace or supplant human decision-making, but simply a means of empowering users to make better decisions based on more accurate information. That, in turn, will achieve the goals of ensuring improved patient outcomes and the most efficient use of scarce resources.

Cerner's Six Principles

Responsibly Applying AI to Documentation and Coding

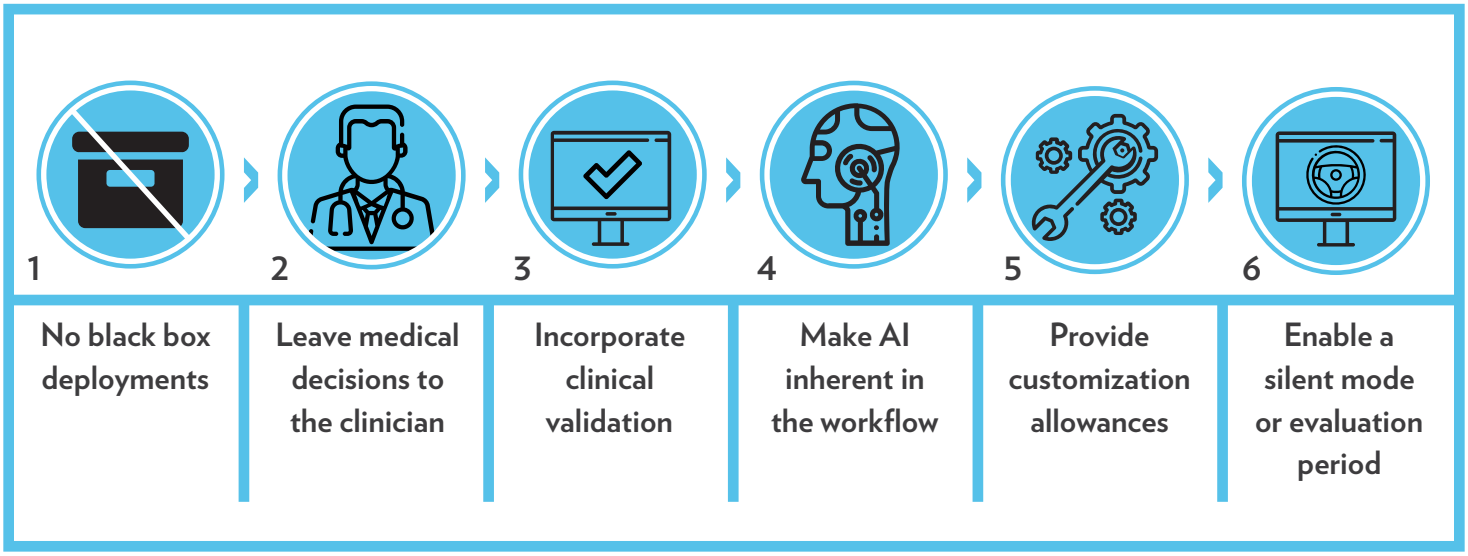

References

[1] Bill Siwicki, "At HIMSS20, Cerner will be talking AI-powered voice technology," Healthcare IT News, February 06, 2020, https://www.healthcareitnews.com/news/himss20-cerner-will-be-talking-ai-powered-voice- technology.

[2] Michael J. Rigby, "Ethical dimensions of using artificial intelligence in health care," AMA J Ethics, 2019;21(2):E121-124, https://journalofethics.ama-assn.org/article/ethical-dimensions-using-artificial-intelligence-health-care/2019-02.

[3] W. Nicholson Price II, "Risks and remedies for artificial intelligence in health care," The Brookings Institution, Nov. 14, 2019, https://www.brookings.edu/research/risks-and-remedies-for-artificial-intelligence-in-health-care/.