What AI Can (and Can't) Do for Healthcare Security and Privacy Programs

Save to myBoK

By Robert Lord

The cybersecurity and privacy field is currently in the midst of an "AI moment," where every vendor, expert, and article about the future of cybersecurity describes how artificial intelligence (AI) is transforming operations and will continue to change the way cybersecurity and privacy professionals work. However, all too often, AI is spoken of in mythical, mysterious terms—a panacea that will save us from institutional confusion, myriad cyberthreats, and workforce shortages. However, industry professionals don't often talk about how AI is most useful in healthcare cybersecurity and privacy. Nor is it often discussed where AI is actually unhelpful and won't be replacing human effort anytime soon. This article covers both of these areas.

When AI Isn't the Cure in Privacy and Security

Professionals face a lot of challenges in healthcare security and compliance every day, and the unfortunate reality is that most of these challenges are ill-suited to be solved by AI. One important area where leaders in privacy and security rarely have time to focus is that of culture change. It's all about human behavior, and changing human behavior is incredibly hard; it requires creativity and judgment, and isn't going to be done entirely by AI in the near future. While AI might provide accountability or leverage that frees up time, culture change necessitates a human touch and intuitive knowledge of what the institution needs. In this realm, AI isn't going to help much with the long-term vision and culture modifications that characterize great systems.

AI also has significant limitations with respect to the insight and creativity required during investigations. AI should provide natural language descriptions of the facts, why those facts came together to present a suspicious event, and all the information one needs to decide whether or not something constitutes a breach. However, that final call still needs to rest with an actual person. The system may not know some critical fact, or lack human input that makes the difference between catastrophe and false positive. That's why every AI system has a "human on the trigger" when it comes to life-altering judgments—such as determining whether or not a cyberattack is a crime or an attack on an institution versus a more benign incident, which would determine whether or not law enforcement or regulatory agencies are contacted.

Finally, one must never forget that jobs in privacy and security are hard, filled with difficult choices that have real impact on the humans both violated by and responsible for violations. Without the uniquely human ethical judgment that each privacy and security officer brings to the team, a healthcare organization would just be following AI recipes, rather than using human judgment to dynamically weigh all considerations and reach the fairest possible conclusion. In some situations, a snooping incident or security infraction might constitute a fireable offense. However, in the context of a user's responsibilities, awareness of policies, and life events, a probationary period may be more appropriate. Good organizations are run by systems and rules, but they have the flexibility to exercise compassion and good judgment for the inevitable corner cases that arise when humans are involved. So, how can AI actually help healthcare?

When AI is Just What the Doctor Ordered

There are three important ways that AI can help with cybersecurity and privacy management—in challenges involving confusion, scale, and repetition.

One inherent complication of healthcare privacy is how confusing the work environments are. How can one differentiate between someone legitimately accessing a file for clinical reasons and someone accessing the data to commit identity theft? There's so much open access to health system information, with tens of thousands of users who can view patient data, it's nearly impossible to differentiate between appropriate and inappropriate access. Is that doctor from another department consulting for a colleague or inappropriately snooping on their ex-girlfriend?

Answering these questions requires deep clinical and administrative insight, and an ability to synthesize data from multiple sources. AI is a perfect tool for this as it can automatically combine and analyze these factors, complete with the context in which they're occurring. This allows systems to see how atypical an action might be and if there's an explanation. In this way, and in a completely automated manner, only validated and accurate alerts (whether potential privacy violations or potential network threats) are making it to privacy and security teams for review, rather than mostly false positives.

Scale is another challenge AI can help with. Privacy and security teams are continuously accountable for tens of millions of transactions and "touches" to patient data every day. No human can possibly look through all these logs—in fact, a team of thousands couldn't look through such voluminous logs. Current paradigms often rely on random audits or simple rules that help you cut through the clutter, which do indeed catch some issues, but only a small fraction of these events. This fact can lead to a profound sense of only seeing the tip of the iceberg. Privacy and security professionals often describe the feeling that they're drowning in data.

AI is a perfect solution for the challenge of scale—it's possible to leverage these technologies to review every single access, and place it in its appropriate context, as noted above. By continuously ingesting and reviewing these access points, privacy and security professionals can train AI programs to only send alerts when a human's judgment is needed. This allows for the best of both worlds. Comprehensive review of every access, as well as only having to review those needles that emerge from the haystack as true threats, are characteristics of AI at work.

Repetition is a huge challenge that privacy and security teams face, whether its analyzing months of past behavior, visiting and correlating physical locations, or simply requesting data and reports from multiple departments in order to piece together a case. There's a huge amount of repetitive, time-consuming work that goes into investigating cybersecurity or privacy incidents. Even once this is done, there are reports that need to be written and that need to be shared with different teams as well as regulatory agencies.

These tasks are perfect candidates for automation. Characterizing user behavior, summarizing facts of a case, and figuring out the details is all about gathering data, processing it, and presenting it for review. While coming to a conclusion about those factors is a human exercise, everything before and after can benefit from AI and bring together the sources to present the facts—nothing more and nothing less.

Doing an AI Privacy/Security Checkup

Privacy and security professionals should reflect on their teams' pain points and see how they match up to AI's challenges and opportunities. If most of the work involved relates to strategic culture change, consider low-cost solutions before thinking about the whiz-bang technologies. Conversely, if the main problems are high volumes of false positives and not having enough time to get your work done, an AI-based solution could be invaluable.

Remember this: AI is no panacea, but if applied correctly it can help users gain the confidence of knowing they're reviewing every access in a system appropriately. This removes a lot of manual work from users' plates, freeing them up for the more important and more enjoyable activities.

Robert Lord (Robert@protenus.com) is the president and cofounder of Protenus.

Article citation:
Lord, Robert. "What AI Can (and Can't) Do for Healthcare Security and Privacy Programs."
Journal of AHIMA 89, no.6 (June 2018): 36-37.

Driving the Power of Knowledge