INTERNATIONAL INSTITUTE FOR ANALYTICS™

RESEARCH BRIEF

RESEARCH & ADVISORY NETWORK

Generative AI in Health Care: Opportunities and Cautions

SAS EXPERTS: JASON DINOVI AND TOM WRIGGINS

IN DISCUSSION WITH BOB MORISON, SENIOR ADVISOR, IIA

JUNE 2024

Discussion Overview

Generative AI offers healthcare organizations – including physician practices, hospital systems, and health insurers – a wide variety of opportunities to make people more knowledgeable and processes more efficient. At the same time, generating new content introduces new challenges and concerns, while the technology landscape is changing fast. The use of Generative AI will affect all aspects of health care, from cost containment to patient care. Organizations must recognize both the potential and the limits of this new technology and be alert to ways it's being used to perpetrate fraud. To discuss the opportunities and the cautions of GenAI in health care, IIA spoke with experts from the SAS Fraud, Risk and Security Intelligence Practice: Tom Wriggins, Senior Manager and Global Lead for Health Care, Fraud and Compliance, and Jason DiNovi, Health Care, Fraud Analytics Consultant.

Generative AI tools, starting with ChatGPT, are all over the news. But that doesn't guarantee that they're well-understood. So let's start with a level set. What is generative AI? What are its new and distinctive capabilities?

Jason: GenAI is a subset of artificial intelligence, which is a very broad category. AI generally tries to handle information and perform tasks in ways similar to what people do. What makes GenAI different is that its output is new content. Rather than using data and algorithms to describe and analyze things, and even to make predictions, GenAI takes in sample data and then creates new data – text, audio, image, video. We could define GenAI simply as data analytics that can make new content for you.

Tom: So it's all in the name, "generative." It's going to build things. The efficiency of absorbing and summarizing large amounts of data is what's really revolutionary for many people and applications. And, as in traditional analytics, with proper supervision GenAI applications can "learn" over time and improve their accuracy and value. A second key feature is that, when we're working with textual data and language, GenAI can return information in a conversational way. It's not just the conventional output of data items, tables, and graphs. If I ask a question, GenAI can deliver information in sentences and paragraphs. If I iterate with GenAI, it's like having a conversation. The information system is finally talking to me.

What else is essential to know about GenAI?

Tom: We need to be clear about what GenAI is and maybe more importantly what it isn't. There are a lot of misconceptions, including attributing to GenAI more capabilities than it actually has. It's generating new results, but it's not really intelligent, not really thinking. It's accumulating and structuring information based on its algorithms and training data. It's using statistical methodology, for example, to predict the next best word to output. It's doing what it has been programmed by humans to do. And that includes the potential biases and unintended consequences of having humans do things. So we can't get complacent. If anything, we need to redouble our efforts around making sure that the models are unbiased and ethical, in two words, Trustworthy AI.

Jason: The public imagination may assume that we're going to be able to take a task and automate it entirely. So I can go from knowing very little about a subject to instantly being quite knowledgeable. But in fact, especially when dealing with a complex analysis or problem, with a variety of data and criteria, GenAI needs an iterative, conversational process to develop useful results. Iteration provides a kind of quality control.

Tom: We also have to keep in mind the rules and regulations around how patient information can and cannot be used. When we are generating new information, we have to be concerned not only with whether its source data is being properly used, but also with where and how the output, the newly generated information, is going to be seen and used. Regulations like HIPAA will continue to apply, even to generated content. For most applications, healthcare organizations want GenAI to be working with locally controlled data, as opposed to scraping the web for whatever it can find.

We also have to pay close attention to *who* is using GenAI. It really lowers the barrier, the level of technical proficiency, needed to engage with large data models. So it's a democratizing force. But it's also being used by malicious actors. They are using it to steal or misuse content, create deepfakes, and spread misinformation. And fraudsters are using it to generate fake medical records, images, and claims. They're using GenAI to create bad content more efficiently and harm the integrity of our healthcare programs.

How are the capabilities of GenAI especially important and valuable in the healthcare sector?

Jason: When we think of the ways that health care is data-intensive, we find plenty of opportunities for GenAI to help. There's handling huge amounts of unstructured data, including medical images and audio-recorded notes as well as text from a variety of sources. Then there's dealing with the complexity created by all that data. Consider, for example, the doctor who accepts dozens of different insurance plans, each with its coverage terms and medical policies. And healthcare delivery has to generate enormous amounts of documentation both in service of patient care and in meeting the requirements of payers and regulators. In short, health care has a lot of data and documentation-intensive processes. Submitting and processing claims is a big one, but so is organizing data about patients so doctors can interact with them in effective and understandable ways. Anytime you need to look across a lot of data to address a clear question, that can be a use case for GenAI.

Tom: Take clinical trials as a specific example. I see big opportunity to use GenAI upstream in designing and locating trials, then downstream in evaluating and documenting the results. Design involves determining a representative participant population with the desired medical conditions and other characteristics, often starting with gender and ethnic mix. Locating trials involves pinpointing areas of the country where that representative population is available. Both are data-intensive analyses with a variety of data sources. Then once trials are under way, they generate extensive data. GenAI can summarize and document results as you go, helping to ensure that the trial is on course and meeting its protocol requirements and objectives.

What are some common use cases for physician practices?

Jason: Think about transcribing medical notes and conversations and then summarizing them for the medical records and patient aftercare instructions. These can bring in additional patient information, including medical history, to provide more clear and personalized information. That's something a GenAI assistant can be very good at, both taking notes and summarizing them as new content. Physician practices need these efficiencies so doctors can spend more time and attention on direct patient interactions.

Tom: GenAI can make things less error-prone. Too often physicians and their staffs make copy-paste errors, copying and updating the notes from the patient's last visit, but in the process possibly missing things that have changed. So the notes might say that blood pressure was taken when it wasn't. GenAI can also help with completeness of care. For example, for a diabetic patient, there are certain blood tests and other things that the physician wants to check and compare every visit. GenAI can give reminders, taking into account recent test results and trends, and so making sure that everything is covered. It's like having a smart and up-to-date checklist of what to do and discuss with the patient. "Have you changed your diet as we discussed last visit, and how has that worked out?" That's a simple example, but having a GenAI assistant listening in, taking notes, and giving reminders can potentially make a real difference in quality and consistency of care.

Jason: If you have a large physicians group, GenAI could also identify any differences in how physicians might be handling similar situations. Obviously, you still want to allow for a doctor to practice and make decisions individually, but spotting the differences can build a valuable feedback loop, a chance for physicians to learn from one another and further improve the quality of care.

Where is the potential for hospital systems?

Tom: What are some of the things that hospital systems are concerned with? Staffing levels and scheduling – are we covered? Occupancy – how many beds are full? Operating room and ICU utilization – are we close to our limits? Supplies management – are inventory levels optimal? I think we'll see a convergence of the predictive models of traditional AI with the natural language interactiveness of GenAI. And for the user, it doesn't really matter which is doing what. Predictive staffing models and patient acuity predictive algorithms we do today with traditional AI. Now we can incorporate GenAI reporting and prompting and custom graphics, so people don't have to study spreadsheet-type reports or drill-down dashboards unless they need more detail.

Jason: The capacity of GenAI in visual media is becoming a big application, with better and more accurate ways to assist radiologists and physicians generally in interpreting medical images. GenAI can look across vast libraries of images and help notice things that are nuanced and may contain telltale signs that humans commonly miss. Another example in basic patient care is translation. GenAI models are advancing the ability to translate between languages, making it easier to communicate with and serve a wide range of non-native speakers.

Tom: Also on the clinical side, GenAI can help with the standing objective of avoiding readmissions. Insurers and Medicaid/Medicare agencies have policies around whether to cover readmissions, especially within days of discharge when the patient may have been discharged too early. Hospitals use traditional AI to analyze readmissions and what drives them. But again, GenAI can enable them to look across broader data, perhaps demographic and social as well as clinical, and deliver analyses in more iterative and conversational fashion, enabling more what-if analysis. And large hospital systems can better compare the readmissions patterns of their individual hospitals.

Looking more broadly, GenAI can prove invaluable as hospital systems continue the transition away from fee-for-service and into population health and managed care. State Medicaid agencies in particular have made the transition to a managed care model. A central challenge is absorbing and making sense of all of the health demographics data on the populations being served. And then calibrating staffing, facilities, services, and supplies to meet those needs.

How about some use cases for insurers?

Jason: For starters, GenAI can enable insurers to better handle the hundreds or maybe thousands of medical policy bulletins that they consult to determine coverages for the whole range of conditions and procedures. That includes both the bulletins they develop internally and those issued by CMS for Medicare and Medicaid coverage. Insurers should be able to identify in plain language what policies apply to a given procedure, as well as what discrepancies or conflicts might occur among them. In addition to giving clarity to beneficiaries and providers, this also benefits insurers from a payment integrity perspective.

Tom: That knowledge can then enable policy teams to reconcile, clarify, and improve the policies around coverages. GenAI can also assist in up-front policy development by pulling together all the information potentially relevant to a given medical policy, including population health information and the latest clinical trials for medications and devices. Insurers have to be able to make sense of these large volumes of data in order to justify their decisions on coverages – what treatments are covered, for how long, and at what cost.

Jason: Another use case is providing patients better access to in-network providers, especially those who take Medicaid patients. People need a straightforward and friendly way to identify and connect with primary care physicians and specialists within manageable distance. With GenAI, such a locator function can delve into information about the physician's practice and experience. Provided information about the patient's specific needs and preferences, it can then make "good match" recommendations. A state's Medicaid website should certainly have a smart locator function.

Tom: That example reminds me that we shouldn't let all the attention to AI distract us from the fact that the core information systems of insurers and hospital systems need significant improvement, especially in terms of information exchange. People's medical data and history are typically spread across different providers' electronic medical records and different insurers' databases. Each time someone relocates or changes insurers, often with a job change, their data doesn't follow them. Rather, it gets more fragmented, and the new providers and insurers make less informed decisions. The general public may assume that all of these electronic systems talk to each other, but they don't, in part because of how patient information must be safeguarded. Carefully orchestrated GenAI may have a role to play in remedying this situation, but it's not a silver bullet.

Let's discuss the possible downsides of GenAI. How can things go wrong, especially in health care? What precautions are therefore needed?

Jason: As with analytics generally, there are three categories of problems. The first is implicit and unrecognized bias, in this case mainly attributable to a GenAI model's training data. This pitfall is much discussed – and a big reason to scrutinize GenAI outputs and choose solutions that incorporate Trustworthy AI into their models. The second is intentional harm through misuse, for example in healthcare claims fraud. As we said, GenAI has made the fraudsters' tasks easier as well. The third is unintentional harm when GenAI output is inaccurate but taken at face value. Some outputs – text, audio, or image – are noticeably, even comically, incorrect. The term "AI hallucination" has become popular to describe this. But most inaccuracies are harder to spot, and inaccuracies in clinical information can have large consequences. This is another reason to do your due diligence, scrutinizing outputs and conversing with models until the results look right.

Tom: That second problem is of enormous concern to payers – private insurers, Medicare, and Medicaid – who are already dealing with waves of claims fraud and abuse/error. GenAI lowers the bar to get into the game. It used to require more medical knowledge to create fake medical records and claims, and creating a slew of them was labor-intensive. With GenAI, it's easier to iterate to produce realistic-looking records, diagnoses, and documentation including medical images. "Create a dental X-ray for a middle-aged adult with a cavity on tooth 22 and a few old fillings elsewhere." With a template to follow, GenAI might create 50 variations on a fraudulent claim in seconds. So from the payers' standpoint, fraudulent claims can become more sophisticated, and their volume can increase dramatically. Cost containment is not just about fraudulent activity, but the overuse of services and errors in billing. GenAI can also help to stem this tide by examining treatment patterns and alerting providers when what they are doing is not following policy or normal treatment modalities.

Jason: The same applies to other sectors like banking. Fraudsters don't necessarily have to come up with novel frauds to perpetrate. They can refine and automate common frauds and inundate institutions with false transactions. That includes phishing attacks. Compromised credentials account for the lion's share of security breaches, and GenAI is allowing for more personalized and sophisticated attacks. Imagine criminals gathering information about an individual through their professional social media posts, identifying what kind of work that person does and who they associate with frequently, and then generating a convincing phishing attempt designed specifically for that individual.

Tom: If criminals can have GenAI create more and better phishing emails, the chance of that one person being fooled is greatly increased. We have seen the impact of these cyber and malware attacks in health care recently in the news.

How can organizations defend against higher volumes of fraud?

Jason: It's a race just to keep pace. As fraudsters get more efficient, those of us combatting them have to as well. That means using conventional analytics and AI to scan growing volumes of claims, more accurately spot anomalies, and triage and prioritize cases for investigation.

GenAI can help accelerate the investigative processes and provide efficiencies so that strained staffing levels are able to accomplish more with augmented decisioning and heightened focus on real threats. Applied to investigative case management systems, it can identify intelligence gaps in notes, summarize what tasks have (or haven't) been taken throughout a lengthy investigation, and even suggest next steps. Medical record review is another opportunity. Confirming clinical notes to substantiate services billed is a highly manual process. Using document vision to read hundreds or thousands of pages and applying AI to parse through it all to identify relevant sections for a reviewer to focus on would be a tremendous efficiency gain.

Tom: Payers need better "radar-detector detectors," ways to identify those who are trying to game and cheat the system. GenAI outputs don't come labeled, so we need to be able to recognize what's suspicious in documents and images. These "tells" can be as basic as finding British spellings in claims nominally originating in Kansas or finding that a provider's listed address is not a medical office building but a vacant lot. Not that long ago, a fraud investigator had to make a site visit and photograph the address. Now we can fold in Google Earth data and find out instantly. With greater volume of claims, it becomes increasingly important to examine them not just individually but in groups to find patterns, commonality, and timing, such as spikes in submissions, that might indicate mass-production of false claims. What we're really doing here is paying more attention to the behaviors of the fraudsters.

How do GenAI projects and their technical underpinnings differ from other analytics-based initiatives?

Jason: The differences stem from the fundamental distinctions we started with. GenAI doesn't just describe patterns and make predictions; it creates new and consumable data. We need to be very careful about the data GenAI is trained on and accesses, what data is used and how. Most importantly, the people who interrogate or converse with GenAI applications need training and understanding, including about the limits of their tools. They need to sense when to take their GenAI output with a grain of salt and reframe the question.

Tom: For the reasons Jason just mentioned, even though the technology is advancing fast, GenAI initiatives should be approached with an extra degree of deliberation. Move a little slower and develop some clarity. What do we want GenAI to do for us that we can't do with conventional AI and analytics? What do we want to accomplish – and learn in the process? What data sources do we want to incorporate? And given the data guardrails in health care, what data must we exclude, and how will that shape the outputs? With GenAI there are still many unknowns. Experiment with the technology for sure, but as with all serious decisions, don't jump in with two feet only to discover how deep the water really is.

Jason: On the underlying technology side, recognize that a lot is going on behind the scenes. An enormous amount of processing power is required to run large language models, and for some applications the training data sets are huge and varied.

Even for what may appear to be a straightforward prompt and output, it's not that simple. There are all kinds of steps including security and API calls, preprocessing and filtering data, and constructing and returning output. SAS and Intel have partnered to handle this processing load and complexity. So all SAS analytics and AI capabilities can take full advantage of the processing performance of Intel chips.

To wrap up, what are your key pieces of advice to healthcare sector leaders who want to make the most of GenAI?

Jason: First, find good opportunities to deploy GenAI, starting with situations where you need natural language processing, where you have lots of data and documents to deal with, and where the clinical or business process calls for people to create narratives. Then consider your desired outcomes and the possible consequences of using GenAI. Is the newly generated information, whether in provider or payer organizations, going to affect the course of care for patients? To avoid unintended consequences, think of GenAI not as a source of airtight answers, but as a time-saver and assistant for people in their work.

Tom: As I mentioned, be deliberate with GenAI. Tread lightly while you determine what works well. Be cognizant of the familiar pitfalls of AI, such as introducing implicit bias, as well as the new ones associated with generating new content. GenAI has many potential uses in the healthcare space, alongside many effective applications of traditional analytics and AI. Some of the best use cases are in the convergence – where GenAI makes other outputs more consumable and conversable, while making established processes more efficient.

An Intel Perspective on GenAI in Health Care

The healthcare industry is dealing with a supply and demand problem. Across the globe, people are getting older, dealing with a greater number of health issues, and the cost of care is increasing. At the same time, the industry is faced with a significant caregiver shortage and clinician burnout is a top-of-mind concern. While it isn't a panacea, GenAI can help address this situation by streamlining manual tasks, enabling clinicians to spend a higher percentage of their time with patients. Generating drafts of answers to patient questions, drafts of patient care summaries, or drafts notes derived from patient encounters are all examples where GenAI is streamlining clinical workflows today.

While the new capabilities enabled by GenAI are exciting, it is important to consider GenAI as a decision support capability that requires trained professionals to consider the generated content and how/if to apply it in the current clinical context. It is critical that these capabilities be developed and used in a safe, effective, and ethical manner.

—Chris Gough, Head of Strategy, Intel Health & Life Sciences

Jason DiNovi

Jason DiNovi has been part of the insurance industry for 15 years with over a decade dedicated to medical fraud analytics. Over several years of designing analytic solutions for commercial health payers, DiNovi has developed extensive expertise in utilizing billing, adjudication, and external data sources to identify anomalous behaviors. DiNovi is a certified professional medical auditor accredited by the American Academy of Professional Coders and an NHCAA Accredited Healthcare Fraud Investigator.

Chris Gough

Chris Gough is head of strategy for the Health & Life Sciences team at Intel. In this role, he works across Intel and partners to integrate advanced technology into solutions that address top industry business challenges. With over 20 years of industry experience working closely with health care, biopharma and government organizations, Gough brings deep health care and life sciences expertise along with a foundation in technology and its application to multiple industries. He's been with Intel for over 25 years and has previously served as worldwide General Manager, Chief Healthcare Solution Architect, and both product and security architect in Intel's Digital Health Division. He holds a Bachelor of Science in computer science from the University of California at San Diego.

Tom Wriggins

Tom Wriggins, with over 30 years of healthcare experience, is considered a thought leader within the commercial and government healthcare and data space. Calling upon his practitioner-level clinical knowledge and experience, Wriggins is responsible for business and analytic design, data management consultation, and application and interpretation of comprehensive program/payment integrity and data analysis solutions.

Wriggins has shared his knowledge through publications and conferences—and has been the invited speaker for several industry association meetings.

Prior to joining SAS, Wriggins spent over 10 years as the practice leader for one of the world's largest technology companies leading global multidisciplinary teams that delivered large and complex data solutions for government healthcare entities, as well as creating training programs associated with fraud and abuse investigative solutions.

Wriggins holds a Bachelor of Science degree in health policy and administration from The Pennsylvania State University. He resides with his family in Southwest Florida and Midcoast Maine.

Robert Morison

Robert Morison serves as senior advisor to IIA. An accomplished business researcher, writer, discussion leader, and management consultant, he has been leading breakthrough research at the intersection of business, technology, and human asset management for more than 25 years. He is co-author of *Analytics At Work: Smarter Decisions, Better Results* (Harvard Business Press, 2010), *Workforce Crisis: How to Beat the Coming Shortage of Skills And Talent* (Harvard Business Press, 2006), and three Harvard Business Review articles, one of which received a McKinsey Award as best article of 2004. He holds an A.B. from Dartmouth College and an M.A. from Boston University.