ahia
Assoc. of Healthcare Internal Auditors

NEW PERSPECTIVES

on Healthcare Risk Management, Control and Governance

www.AHIA.org Journal of the Association of Healthcare Internal Auditors Vol. 43, Number 3, 2024

Next-Gen(AI) Fraud

Identify the threats, responses and governance of Generative Artificial Intelligence
page 7

Patient Monitoring in the Digital Age

Evaluate the risks and benefits of hands-off patient care
page 12

Navigating Hospital-Physician Subsidy Arrangements

Examine hospital-based specialist compensation
page 17

From Risk to Resilience

Navigate workplace violence internal audits
page 28

Next-Gen(AI) Fraud

Identify the threats, responses and governance of Generative Artificial Intelligence

By Victor Hartman, JD, CPA, CFF, CFE

Generative Artificial Intelligence (GenAI) burst onto the world scene in November of 2022 with the introduction of ChatGPT, and it is now radically changing business processes. The internal audit function for organizations is also being impacted as staff wrestles with the threats and benefits of this disruptive technology.

To appreciate the threats Generative Artificial Intelligence (GenAI) brings and how organizations can use it to respond, you need to understand how this fascinating technology works. Regulators, creators of GenAI, and implementors of this technology are also debating whether and to what extent governance protocols should be adopted to control GenAI so that it is beneficial and not harmful to humanity.

What is GenAI?

GenAI is a type of artificial intelligence that can create original content. This content may be answers to complex questions or it could be images, lyrics, videos, computer code, and other media. There are already many different GenAI models, and they work in a variety of ways. Many GenAI models try to mimic the human brain using a neural network architecture.

GenAI models

The base feature of most GenAI is a large language model that uses transformer technology. A transformer works well with language by simultaneously interpreting each word in a sentence, paragraph, or document. It will then rank the importance of each word in proximity to the other words, enabling it to comprehend human language.

Once a prompt is understood, the transformer model can resolve the task within its pre-trained data network and search the internet for new information. The final step is to provide the requested output in various mediums, which could include a verbal response.

Another GenAI model is known as an autoregression model. Autoregression models are time-series models that predict future values based on past observations in a sequence. The model predicts the next data point in the sequence

The need for vast data encourages GenAI implementors to acquire and use the data illegally, unethically, and in ways not imagined even a few years ago.

The courts and legislatures are playing catch up to address novel legal issues.

using its previous data points as inputs. In other words, it models the relationship between a data point and the data points that came before it. Autoregression models are widely used in time-series forecasting, financial market predictions, and other sequential data analysis tasks.

One of the more intriguing GenAI models is the generative adversarial network (GAN). GANs are a class of deep learning models consisting of two neural networks, the generator and the discriminator, which are trained adversarially. The generator tries to produce synthetic data similar to real data, while the discriminator aims to differentiate between real and generated data. GANs have shown remarkable success in generating lifelike images, creating new art, and generating synthetic data for various applications.

Another example of a GenAI model is the variational autoencoder (VAE). VAEs use an encoder and a decoder to generate content. The encoder takes the input data, such as images or text, and simplifies it into a more compact form known as *latent space*.

Consider a cat as an example. The encoder will learn by being given, say, 1000 cat images. The VAE will measure each cat's color, eyes, size, dimensions, etc. From this data, the encoder will take the mean and a standard deviation of all the data and place it in the latent space. When prompted, the decoder takes this encoded data and creates an entirely new image of a cat.

GenAI threat picture

GenAI presents threats to organizations. These threats may include existential threats to an organization or an entire industry. They also include fraud, espionage, and past threats that have been repurposed with this technology. Internal auditors' skillsets uniquely enable them to assist organizations in preventing, detecting, and investigating threats posed by GenAI.

Data threats

Fundamentally, GenAI needs lots of training data to be effective. In the healthcare setting, GenAI innovators need healthcare data including patient files, medical procedures, outcomes, and billing records to train their models. The Centers for Medicare & Medicaid Services (CMS) houses the largest repository of health records and is using GenAI to improve healthcare outcomes for the public.

GenAI terminology

Input/Output: The input is the media (language, text, image, data, etc.) placed in the GenAI prompt, and the output is the response that it provides.

Prompt: The prompt is the user's input or request to a GenAI model seeking an answer or an output. It is usually a verbal or written request but could include various mediums. In a hospital context, it could include doctor dictations and the results of an MRI, X-ray, and echocardiogram.

GenAI model: The model is the algorithm that drives the software being used. While the number of models constantly grows, some standard models include the transformer, generative adversarial networks, and autoencoders.

Training data: The GenAI models need copious amounts of data to train the algorithm. If the training data is flawed, skewed, or incomplete, the GenAI output could be low quality, biased, or inaccurate.

Hallucination: A *hallucination* is a GenAI output that is completely made up or fabricated.

The need for vast data encourages GenAI implementors to acquire and use the data illegally, unethically, and in ways not imagined even a few years ago. Data is being scraped from the internet, including corporate websites, X (Twitter) feeds, Wikipedia, open-source software from GitHub, and many other sources.

Intellectual property owners have had their original content taken and used to train GenAI models as part of this insatiable need for data. These models are then producing derivative content from legally protected intellectual property. The courts and legislatures are playing catch up to address this and other novel legal issues.

An organization's data presents a variety of issues. Corporate employees may place sensitive company data into a publicly available GenAI chatbot, e.g., ChatGPT. That information may be used to train the GenAI model; consequently, corporate secrets can end up in the public domain. Malicious actors may intentionally cause foreign adversaries' or corporate competitors' GenAI to be trained on erroneous, false, or biased data so that their GenAI produces flawed results in a scheme known as adversarial machine learning.

An employee of a publicly traded company could ask the company's GenAI to describe the new products the company is developing, when the employee could not otherwise obtain this nonpublic information. With the GenAI-created information, the employee could engage in insider trading.

Theft of intellectual property is already a significant problem for organizations. GenAI may make it easier for a departing employee to take the company's secrets before working for a competitor.

Deepfakes

GenAI is exceptionally adroit at cloning images and voices. This will undoubtedly provide the entertainment industry with new creative tools. It will also enable fraudsters to use deepfakes to commit fraud. There are already many examples of a grandchild's voice being cloned to entice a grandparent to send money to a fraudster as part of some falsely portrayed harm to a grandchild such as an arrest, accident, or stolen wallet.

For the past decade, the most virulent fraud that has been victimizing organizations worldwide is the business email compromise. In this scheme, the fraudster convinces a corporate employee to pay a hefty invoice to a bank account controlled by the fraudster. These fraudsters have enhanced their tradecraft for years, resulting in billions of dollars in losses annually. GenAI can make this fraud all the easier.

Fraudsters are now using GenAI cloning technology to defraud victim companies by creating a voice clone from a recording of a senior executive's voice obtained from social media or a pretext conversation. The voice clone can be used to leave voicemail instructions or even speak in real-time with staff. Due to the convincing nature of hearing a known voice from a senior executive, a payment official will confidently make an errant payment to the fraudster for an invoice, real estate purchase, or corporate merger. This scheme has already resulted in substantial losses across the globe.

Victim targeting

Frauds and scams that have long victimized individuals and organizations are now being enhanced with GenAI. In the past, fraudsters were limited in how effectively they could perpetrate typical schemes such as spear phishing, romance scams, and the Nigerian 419 letter fraud. Of the world's eight billion people, only 1.5 billion speak English. Consequently, fraudsters from a foreign country may not be able to speak the victim's language, understand local customs, or be conversant in a particular business jargon. GenAI changes all of this. Fraudsters can now draft emails and use voice clones to communicate with anyone, anywhere.

Other malicious uses

GenAI can perform the tasks of skilled humans at lightning-fast speeds. This will both enhance and threaten professions as it is implemented. Further, this same technology can also be used to carry out malicious tasks with enormous efficiency.

The ability to create fake news or generate widely spread false narratives can now be easily accomplished. GenAI can constantly populate and update an exceedingly large number of malicious websites with false narratives in pursuit of political, social or economic goals.

Similarly, malicious actors with limited or no software coding skills can create malware. The creation of child pornography will frustrate law enforcement efforts because the current database of known images will not be of benefit. Cyberbullying and sextortion schemes now include taking a victim's likeness and superimposing it on pornographic images or videos.

While plagiarism is nothing new, GenAI makes it easier, and more difficult to detect. This will not only be a challenge for

Fraudsters can now draft emails and use voice clones to communicate with anyone, anywhere.

Malicious actors with limited or no software coding skills can create malware.

colleges and universities but for research institutes and the business world as well. Researchers may be tempted to claim GenAI work product as their original work. GenAI's ability to create synthetic data has already been used in several frauds. In one case, an attorney was referred for disciplinary action by a judge for his use of a GenAI-created legal brief that contained hallucinated citations to nonexistent cases.

GenAI fraud detection

GenAI can assist in the prevention and detection of historical frauds as well as new frauds created by this technology. This can be accomplished by various business models used to implement GenAI. These models range from a pay-as-you-go subscription model to an organization bringing the entire operation in-house, where it owns the technology and trains the algorithms with its data.

Using GenAI in the prevention and detection of fraud is a new field. GenAI can be programmed to assess fraud indicators based on human behavior, sentiment analysis, internal control breaches, executive actions, financial incentives, and company performance. After an organization navigates various legal and corporate cultural issues, GenAI could be trained on an organization's entire dataset. This data would include an organization's structured data, e.g., quantifiable data such as spreadsheets and financial statements, and unstructured data, such as emails, voicemails, and calendar notes.

GenAI implementation teams are expensive and include software engineers, data scientists, compliance professionals, lawyers, and subject matter experts. The subject matter experts may include internal auditors, management, and certified fraud examiners, and they are critical for both creating fraud prevention and detection algorithms and identifying relevant and quality datasets to train the GenAI model. Internal auditors are a natural choice to be on this type of implementation team because they know primary fraud methodologies, such as the fraud triangle, and the organization's fraud risks.

GenAI governance

Those most knowledgeable about GenAI strongly believe there should be guardrails around GenAI technology out of fear that it could be an existential threat to humanity. Governments around the world have begun debating and passing GenAI legislation. China has already developed a robust regulatory framework as part of its goal of becoming the world leader in GenAI by 2030. The European Union, which has historically been a world leader in technology regulations, is in the process of passing comprehensive legislation. The United States has been slow to develop a GenAI regulatory framework. Arguably, this is because of the United States' philosophy to innovate first and then regulate. A more sinister view is that Congress is inherently divided, and the passage of any legislation is just too daunting.

Corporations are now striving to develop GenAI governance models in the current legislative void that will serve their interests. Thought leaders in this space generally recognize several GenAI governance principles

- Data Privacy
- Trust
- Accountability
- Explainability
- Fair and unbiased
- Transparency
- Human control
- Reliability
- Safety
- Security

When implementing GenAI, organizations will need strict policies and procedures with robust oversight from those in charge of governance. The implementation team should have built-in audit procedures where deficiencies are quickly reported to management. For example, a primary factor in the successful implementation of GenAI is the acquiring, safeguarding, and using various datasets during each iteration of training the GenAI model. Just as important, the organization must document each facet of data integration so that someone is accountable for an output that produces flawed, biased, or harmful results. This is a prime example of the GenAI governance principles regarding data privacy, transparency, human control, and accountability.

Using GenAI in the prevention and detection of fraud is a new field.

Conclusion

GenAI presents new challenges for organizations. They must learn to implement this technology to create greater efficiency and effectiveness. Conversely, GenAI presents new threats that could not be anticipated just a few years ago. Fortunately, GenAI can also be used to moderate these threats while conforming to accepted governance principles. Internal auditors are well placed to anticipate, implement, and utilize this emerging and exciting technology. **NP**

Victor Hartman, JD, CPA, CFF, CFE, is Principal of The Hartman Firm, LLC, specializing in forensic accounting and internal investigations. He was an FBI Special Agent. He is now an Adjunct Professor at Georgia State University and serves on Georgia Southern University's Forensic Accounting Advisory Board. Vic can be reached at 404-369-0616 and Vic@HartmanFirm.com.

"The computer is tired of you taking all the credit and it's demanding half of your paycheck."

About *New Perspectives*

New Perspectives (NP) is a refereed and peer-reviewed journal that focuses on up-to-date information, trends and issues in the healthcare industry and the internal auditing profession. Practical guidance is provided on risks and controls that can be applied by internal audit professionals in their jobs.

NP is published quarterly in an electronic format. Issues are accessed by online viewing or through download. See the NP archives at https://ahia.org/new-perspectives-archive/ (login required).

For author guidelines or to submit an article, please contact Jen Conley at 801-803-2361 or Jen.ahia.np@gmail.com.

Yearly subscription rates, including postage, are $100, payable in U.S. funds. No refunds or cancellations. Send publication and subscription inquiries, address changes and other inquiries to: Association of Healthcare Internal Auditors, Inc., 19 Mantua Road, Mount Royal, NJ 08061 USA. Phone 856-554-1083 or email, info@ahia.org.

New Perspectives, its editors and the Association of Healthcare Internal Auditors, Inc. are not responsible for the opinions and statements of its contributors and advertisers. The authors do not necessarily reflect the official policies of AHIA nor does AHIA endorse any products. AHIA does not attest to the originality of the author's content. Reprints of any portion of *New Perspectives* may be used for educational or instructional purposes only, provided the following statement appears on each reprint: "Reprinted with permission from *New Perspectives*, Journal of the Association of Healthcare Internal Auditors, Inc. Volume/Number."