POLICY & MEDICINE
COMPLIANCE UPDATE
Reprinted from
Volume 9.5 | May 2023

FEATURE

Embracing Generative Artificial Intelligence

The Race to Reshape the Life Sciences Industry

By Katie Jacyna and Schuyler Tasman[1]

Summary: Beginning in late 2022, generative AI has moved from a science fiction topic to a real tool with the potential to transform the life science industry. Therefore, life science compliance professionals must understand the technology and carefully consider its implications. This article provides a working background on generative AI and its potential uses within the industry.

Risk, risk mitigation, controls, and documentation are four universal concepts that dominate discussions about life science compliance operations. However, those discussions over the past several years involve a new fifth concept – artificial intelligence. For compliance functions struggling with limited resources and expansive responsibilities, artificial intelligence ("AI") provided hope that this new technology could increase the opportunity to standardize processes, enforce crucial controls, and provide greater accessibility to monitor operations. But until recently, AI's potential was illusory.

With the emergence of Generative AI, the latest advance in artificial intelligence swiftly being adopted across all industries worldwide, the unrealized potential has moved from illusion to reality. If the AI champions are correct, generative AI, used correctly, will transform the workplace and its compliance culture. However, like all new technologies, generative AI can be confusing and overwhelming for those compliance professionals struggling to understand what is happening inside this new "black box." Furthermore, with its complex coding structure and variable learning models, generative AI can appear overly technical and, thus, uncontrollable. This article explores the technical aspects of generative AI, current life sciences applications, and appropriately managing AI-associated risks.

Exploring Generative AI

OpenAI sent shockwaves in late 2022 when they released ChatGPT to the public as a free chatbot that generates an answer to any question asked.[2] Since its release, the chatbot has captured the attention of the business world, which is exploring ways to operationalize the technology.

In a compliance context, Helio has observed that ChatGPT can write a policy on healthcare professional ("HCP") meal limits, identify essential elements of recent Corporate Integrity Agreements ("CIAs"), and even list transparency regulations by country. Helio asked attendees how compliance professionals viewed these developments at a recent compliance innovations conference. The overwhelming response was that most compliance teams were taking a wait-and-see approach. However, most seemed hesitantly excited with a deep interest in learning more.

What is Generative AI?

Perhaps the simplest explanation of generative AI is that it is a better search engine. For example, currently answering everyday questions involves the following steps:

1. You enter a set of search terms or a question into Google.
2. Google provides a list of possible matches for your terms.
3. If unhappy with the results, you can modify your search criteria and search again.
4. You then look through multiple websites, articles, videos, and pictures to gain a broad understanding of the topic.

5. That information is now summarized, creating your understanding of the answer, and you are then free to apply that knowledge.

Although using generative AI involves similar steps, it reads more information and processes more data than humans have the time to do. Moreover, it utilizes different artificial intelligence techniques and models to generate unique content. It accomplishes this by using a few components and processes to provide a curated response. These components include:

1. **Data preprocessing:** One of the first steps in creating a generative AI model is data preprocessing. This step involves collecting and cleaning the input data to prepare it for training (i.e., learning). The data may need to be reformatted, normalized, or standardized to ensure the datasets are suitable for the model.

2. **Machine learning:** Generative AI uses machine learning algorithms to learn from large datasets and generate new samples. Specifically, it involves training a model to learn the input data's underlying patterns and statistical regularities. These patterns and statistical regularities can then be used to generate new data samples. Generative AI can employ several different machine learning algorithms in this process, such as generative adversarial networks ("GANs"), variational autoencoders ("VAEs"), and Boltzmann machines.

3. **Natural language processing ("NLP"):** Another facet of generative AI involves using NLP techniques to generate human-like text, such as chatbots, language models, and text generators. NLP utilizes statistical models and algorithms to analyze and generate natural language text. Achieving natural language responses requires training models to learn the statistical patterns and structures of language, such as grammar, syntax, and semantics.

4. **Deep learning:** Generative AI often involves deep learning techniques. These techniques are a subset of machine learning that involves training deep neural networks with multiple layers. Deep learning has proven particularly effective for specific generative AI tasks, such as image and speech generation because it can capture more complex patterns and relationships in the data. For example, GANs is a type of deep learning model that involves training a generator network to generate new samples that are evaluated, helping the generator improve its output.

5. **Reinforcement learning:** Generative AI can also utilize reinforcement learning, which involves training the AI to interact with an environment and learn from the results. Reinforcement learning can be used for game AI, where the system learns to generate new game scenarios based on previous interactions within the environment. For example, in War Games, the 1983 movie, which was well ahead of its time; the computer "played" multiple nuclear war game scenarios to ultimately "learn" that the only winning move was not to play.

6. **Model Training:** This component involves training the neural network using large datasets of examples. As the training progresses, the weights and biases[3] of the network are adjusted to minimize any disconnect between the predicted output and the actual output. These adjustments are typically done using gradient descent optimization algorithms.

7. **Evaluation:** Evaluating the performance of a generative AI model can be challenging because the output is not always clearly defined. Nevertheless, evaluating the model's performance generally involves the use of metrics. For example, perplexity measures how well the model predicts the next word in a sentence. Alternatively, the Fréchet Inception Distance ("FID") score can measure the similarity between AI-generated and real images.

Combined within a generative AI framework, these components can create a chatbot that can comprehend, learn, and respond to any query (Figure 1).

The generative AI model, in this case, ChatGPT, makes predictions based on specific patterns it is "trained" on. Depending on the model's training, it can produce a wide range of written materials, including creating new content, such as images and audio, summarizing complex information, and even providing detailed answers to questions.[5]

FIGURE 1: *Components within the ChatGPT framework*[4]

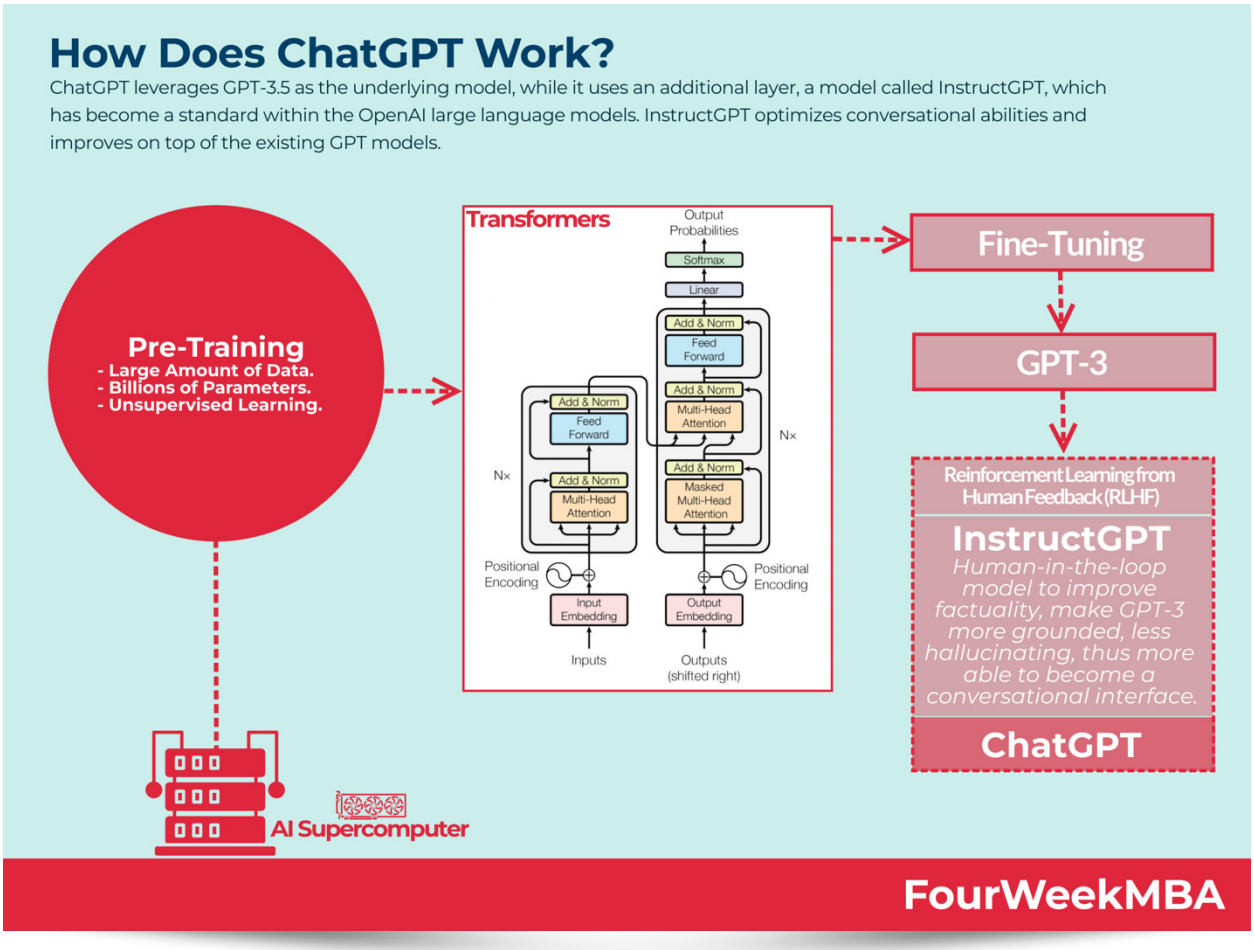

Moreover, developers are exploring new ways that generative AI can assist with process workflows. These potential improvements include:

- Automating the manual process of writing content,
- Reducing the effort of responding to emails,
- Improving the responses to specific technical queries,
- Creating realistic representations of people,
- Summarizing complex information into a coherent narrative, and
- Simplifying the process of creating content in a particular style.

However, these improvements necessitate reliance on massive amounts of information and constant machine learning for refinement and learning.

Using Generative AI Within Life Sciences

Life sciences have used generative AI in a limited capacity, typically for R&D and clinical research, to assist with aggregating and accessing large sets of research data. However, the launch of ChatGPT opened the world to accessing publicly available knowledge, bringing with it potential new ways to utilize AI within life sciences.

R&D

Life science research and development has employed generative AI in drug discovery, personalized medicine, and medical imaging.[6] So far, drug discovery seems to be the primary application for this technology.

For example, companies can feed thousands of rows of clinical and research data into a generative AI model to generate new optimized molecules. After adding the data to the model, the development team can select specific properties, such as potency, selectivity, or solubility.

Used in this manner, generative AI modeling, in essence, is the advanced successor to earlier high-throughput target screening.

Even at this early phase of the technology, AI has shown the capability of generating quality drug candidates quickly, with a much higher probability of a successful transition to clinical development.[7] Consequently, by employing AI, companies can develop new drugs and devices faster and enter the market sooner, substantially increasing revenues and decreasing costs.

Commercial Operations & Medical Affairs

Given the use of generative AI in R&D, coupled with the successful launch of ChatGPT, companies are poised to employ generative AI in commercial operations and medical affairs.[8] However, while it is too early to tell how sales, marketing, or medical science liaisons will use this new technology, many believe it is inevitable.

There are many possible uses for generative AI-supported initiatives within life science commercial operations and medical affairs. For example, generative AI can assist companies in accessing large quantities of unstructured data within their systems. MIT researchers estimate that unstructured data account for 80% to 90% of the data collected by businesses, with much of it coming from commercial teams.[9]

Accessing this data can provide significant insights into drug and device customers resulting in improved targeted advertising and sales approaches. Thus, generative AI can help sales teams provide more personalized experiences for HCPs based on available data, such as prescribing information, social media use, and publication access.

Generative AI can also assist marketing teams by analyzing market segments using publicly available data such as patient populations and geographic location. With these insights, marketing can better target customer segments with focused messaging and outreach.[10]

Compliance

Juxtaposed against significant compliance risks resulting from the use of generative AI in commercial operations or medical affairs, the same technology has the potential to assist compliance in mitigating risk, enforcing controls, and embedding a culture of ethics and compliance. For example, compliance professionals have already identified that generative AI can help them answer employee compliance questions.

Even where policies, procedures, industry codes, and other compliance guardrails are centralized, indexed, and highly accessible, subtleties and complexities abound that require the assistance of experienced human compliance professionals. Using generative AI to assist with employee compliance questions can provide timely compliance support around-the-clock, reducing the risks associated with delayed or missed responses.

Consider the following hypothetical exchange between a company sale representative and a generative AI chatbot:

> **Sam Sales Rep:** ComplianceBot, I am having dinner tonight with some physicians here in Minneapolis. What are the meal limits for HCPs?
>
> **ComplianceBot:** The general company policy limits the cost of dinner for each HCP to $120 (U.S.), including drinks, taxes, gratuities, and delivery fees. However, for certain high-cost cities, that amount can be increased to $150 (U.S.). Minneapolis is not a high-cost city.
>
> **Sam Sales Rep:** Thank you. My HCPs are from Texas and Vermont. Also, they are planning on bringing their spouses. Can I buy their spouses' meals?
>
> **ComplianceBot:** HCPs with an active Vermont license cannot be provided with a meal during your interaction with them. Moreover, meals may not be provided to spouses of Vermont or Texas HCPs. See section 3 of the HCP Meals and Entertainment Policy for additional details and restrictions.

The hypothetical illustrates how a generative AI chatbot can improve compliance. First, the chatbot not only answers the initial question but continues to understand the conversation in the context of the first question. For example, the chatbot highlighted the restriction on meals to Vermont HCPs even though the question concerned providing meals to spouses.

Second, it provided a comprehensive response on the standard limits and those for high-cost cities. It also understood and highlighted that Minneapolis was not a

high-cost city. Finally, the sales representative received an immediate response because the chatbot was trained on the company's compliance policy.

The above hypothetical exchange is just one use for a generative AI chatbot in compliance. Other possible uses include ensuring sales representative communications with HCPs remain "on-label" or providing Medical Science Liaisons ("MSLs") with quick access to detailed medical information.

Regardless of the use, a compliance team can use a generative AI platform to reduce the time and effort of answering questions to reduce the risk of minor violations. It also allows compliance teams to focus on the more challenging questions and high-risk areas within an organization while simultaneously improving transparency and accessibility. Thus, it can help compliance meet the typical business challenge of doing more with less.

Generative AI's powerful ability to consume and understand large amounts of data also provides opportunities to monitor HCP transactions efficiently and trigger automatic escalations to compliance as needed. Using generative AI in this capacity allows organizations to remove manual compliance data analytics review and filtering currently used by many organizations.

Finally, generative AI can assist with ensuring operational compliance at its core. It does so by ensuring all operations and transactions go through established flows and providing real-time notifications to the compliance team if any deviations from predetermined processes are detected.[11]

Generative AI Risk Mitigation

ChatGPT is built on utilizing a public model, which provides access to an expansive knowledge base, but that may come with organizational risk. The key to managing risk within AI is ensuring controlled change and document management. In addition, utilizing a local generative AI ensures that strict control can be maintained by controlling what the AI consumes (i.e., feeding it a compliant diet).

As noted at the beginning of this article, risk and risk mitigation are universal compliance concepts. Within life sciences, there is no shortage of risks that need mitigating, and generative AI is no exception. Moreover, given the speed of its adoption, compliance professionals, at a minimum, need to understand the capabilities, limitations, and potential risks posed by generative AI. As companies are racing to adopt and use generative AI to improve the industry, there are a few critical factors to consider when assessing how generative AI is used in a life science context.

The technical approach and configuration are crucial decision points that will impact data security and accessibility. Generative AI systems have varying levels and terms for security and how "private" collected data will be from the rest of the data collected.

Ensuring proprietary and confidential information is stored and accessible only to the source organization is highly recommended. Some AI models, such as BERT and T5, are typically trained and use only specific datasets provided to them. Therefore, BERT and T5 are ideal models for compliance because the information cannot be altered or supplemented by other organizations or information.

However, GPT-3, on which ChatGPT is built, utilizes a more general-purpose approach. Thus, GPT-3 may supplement or include information from other sources. Therefore, GPT-3 is likely a better model for commercial teams.

Training the generative AI system is also essential to ensuring it reacts and provides the expected responses. Initially, training identifies knowledge gaps or ambiguous information, which can be corrected and fed back into the AI system. It also can help remove bias, which can be mitigated by careful selection and curation of training data, as well as continuous monitoring of the AI.

Likewise, citations and sources can be added to the AI system responses to provide traceability and auditability. Doing so allows the user to clarify and potentially detect inaccurate responses. One inherent weakness of generative AI is its inability to admit what it does not know. Therefore, setting up the system to highlight when it cannot correctly answer a question is critical to promoting reliability and trust.

Another important consideration is that machine learning feedback is a key component of generative AI, and it is essential to fine-tune the AI. For example, if the model is trained too closely to any specific dataset, including compliance information, it may be

"overfitted," which means it may perform poorly when encountering new data. Poor performance translates into poor decision-making or inaccurate predictability. Machine learning can help mitigate this by ensuring the model is evolving and taking new information in to train the model.

Human escalation is another crucial aspect of the machine-learning process. In this scenario, questions that the AI cannot answer are directed to a human to review and respond manually. Once the response has been provided, machine learning will absorb that information and initial question feedback into the system building a more robust knowledge base for future questions.

Finally, cost considerations always come with new technology until it has scaled to common use. For example, open-source generative AI models can cost ten times or more than utilizing a more local model. Additionally, the cost of uptraining human resources to build, manage, and implement generative AI may be a costly endeavor. However, third-party vendors may be able to reduce costs if the generative AI is purchased as a SaaS solution, given that model is purchased and can be spread across several organizations defraying the cost.

Conclusion

Love it, like it, or hate it, the use of generative AI is likely here to stay. Therefore, regardless of adoption, compliance professionals must understand and recognize that generative AI requires a thoughtful and concerted approach.

Moreover, generative AI is not human. Consequently, when considering adoption, it is essential not to discount the human implications surrounding generative AI (e.g., culture, change management, and basic human emotions). Compliance's seat at the table is crucial when deciding what enterprise content is appropriate for generative AI and the best ways to consolidate, use, and structure that information. However, through understanding and careful implementation, generative AI could become one of the most useful tools in the compliance toolbox.

References

1 Ms. Jacyna is the Vice President of Product Management, and Ms. Tasman is an Analyst with Helio Health Group.

2 *See* Fionna Agomuoh and Luke Larsen, *ChatGPT: How to use the AI chatbot that's changing everything*, DIGITALTRENDS (Apr. 23, 2023), https://www.digitaltrends.com/computing/how-to-use-openai-chatgpt-text-generation-chatbot/.

3 AI bias refers to the tendency of artificial intelligence (AI) algorithms to make decisions or predictions that are systematically and unfairly skewed toward certain groups or outcomes.

4 *See* Gennaro Cuofano, *How Does ChatGPT Work?*, FOURWEEKMBA (Feb. 24, 2023), https://fourweekmba.com/how-does-chatgpt-work/.

5 *See* MCKINSEY & COMPANY, WHAT IS GENERATIVE AI?, https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-generative-ai (Jan. 19, 2023); *see also* Kaushik Pal, *How Can an AI Model Create Music*, TECHOPEDIA (Apr. 28, 2023), https://www.techopedia.com/how-can-an-ai-model-create-music.

6 *See, e.g.*, Linda Rosencrance, *9 Uses of Generative AI in Healthcare*, TECHNOPEDA (last updated Apr. 17, 2023), https://www.techopedia.com/9-uses-of-generative-ai-in-healthcare#:~:text=9%20Uses%20of%20Generative%20AI%20in%20Healthcare%201,8%20Developing%20New%20Research%20Ideas%20...%20More%20items.

7 *See, e.g.*, MCKINSEY & COMPANY, AI IN BIOPHARMA RESEARCH: A TIME TO FOCUS AND SCALE, https://www.mckinsey.com/industries/life-sciences/our-insights/ai-in-biopharma-research-a-time-to-focus-and-scale (Oct. 10, 2022).

8 *See, e.g.*, Manas Das, *The Power of Generative AI in [the] Pharmaceutical Industry*, LINKEDIN (Apr. 22, 2023), https://www.linkedin.com/pulse/power-generative-ai-pharmaceutical-industry-manas-das/; Daniel Fallmann, *Making the Most Out of Generative AI In Your Enterprise*, FORBES (Apr. 28, 2023), https://www.forbes.com/sites/forbestechcouncil/2023/04/28/making-the-most-out-of-generative-ai-in-your-enterprise/?sh=440019fb445e.

9 *See* Tam Harbert, *Tapping the power of unstructured data*, MIT SLOAN SCHOOL (Feb. 1, 2021), https://mitsloan.mit.edu/ideas-made-to-matter/tapping-power-unstructured-data.

10 *See* Anupam Nandwana, *How Technology Will Impact Pharmaceutical Sales in 2023*, FORBES (Feb. 27, 2023), https://www.forbes.com/sites/forbesbusinesscouncil/2023/02/27/how-technology-will-impact-pharmaceutical-sales-in-2023/?sh=1654b72511e8.

11 *See* Sumeet Singh, *Generative vs. Deterministic Artificial Intelligence in Compliance Workflows*, THE COMPLIANCE & ETHICS BLOG (Apr. 10, 2023), https://www.complianceandethics.org/generative-vs-deterministic-artificial-intelligence-in-compliance-workflows/.

www.heliohealthgroup.com

David Herr
Account Executive | Helio Health Group
dherr@heliogrp.com
605.390.6737