BANK

COMPLIANCE MEETS AI

A Banking Love Story

BY SAM HOLLE, JD, CRCM, AND KRIS STEWART, JD, CRCM

We decided to ask ChatGPT to help us write an introduction to this article.

Q: We're writing an article for a banking risk and compliance magazine on large language models. Tell us a fun fact about you that would encourage compliance officers to want to use you.

A: *Fun fact: As a large language model, I can significantly streamline and enhance the compliance process for banking professionals. Compliance officers can leverage my capabilities…*

Not what we hoped for, so, we tried again.

Q: Can you be a little more succinct and fun?

A: *Certainly! Fun fact: I'm your compliance co-pilot, navigating the complex terrain of banking regulations with lightning speed and precision. I'm the Swiss Army knife for compliance officers, simplifying searches, drafting documents, and keeping you on the straight and narrow in the exciting world of banking compliance!*

Better right? And we could not leave well enough alone – so –

Q: That's good. How about a catchy title for the article?

A: *"Compliance Revolution: How Large Language Models are Transforming Banking Regulation"*

Q: Too long and not fun. Do you have something shorter and pithier?

A: *"Compliance Meets AI: A Banking Love Story"*

Nailed it! We promise, from here on out, the writing is all human!

STOCK.ADOBE.COM / FRESHIDEA / RARINLADA / VLAD KOCHELAEVSKIY

Generative AI: What You Need to Know

Rarely does a day go by without our news feeds reporting on the rapidly evolving world of Generative AI (GAI)—artificial intelligence that can generate text, images, or other media in response to prompts. GAI models, and more specifically large language models (LLMs), a type of GAI, are being deployed across industries and financial institutions are thinking of ways they can do so, too—a challenge made all the greater by being highly regulated entities. Larger financial institutions are developing their own solutions, and of course vendors are rapidly figuring out ways to use this technology in their solutions.

So, what exactly are LLMs and why should you become familiar with them? LLM's, like ChatGPT and Bard, are models that are trained on vast amounts of textual data such as books, articles, white papers—anything you can find on the internet. The models recognize patterns and relationships in human language. They take that learning to generate text in response to your question (called a prompt in the AI world). Based on the vast data upon which they are trained; they answer your prompt based on predictive analytics—determining the most logical word or phrases in response to your question.

Risks Associated with GAI Technology

If you like to keep an eye on technology news, you have likely seen reports of some spectacular "fails" when it comes to the use of LLMs. Take for example, the attorney who submitted a brief to a court, citing cases, complete with quotes, citations, and judge's names. The only problem was that ChatGPT made them up. This type of inaccuracy is known in the AI world as "hallucination."

Hallucination is a tendency of a LLM to "make up facts," drawn from the many terabytes of data they have been trained with, to produce very convincing answers to prompts. Hallucination can result because LLMs are trained to be predictive, based on word patterns—with the goal of being conversational vs strictly knowledgeable.

The good news is that while hallucinations cannot be eliminated, they can be minimized and mitigated. LLM developers have been quick to respond to reported epic failures in their models to provide more safeguards in the results produced. OpenAI, the company behind ChatGPT, has initiated a new methodology that it refers to as "process supervision" that is intended to evaluate each logic step in the model's method for answering a prompt and, in doing so, to help detect and mitigate a model's logical mistakes. Also, controlling the data sources the model is trained on by using domain-specific, trusted data sources can provide better results. Ultimately, there is no substitute for expert (human) verification of the results.

Better Answers Through Better Questions

Another way to mitigate hallucination is to provide better context via prompt engineering. Prompt engineering refers to the practice of composing your question such that the system provides a more accurate answer. Using traditional search engines, we are used to asking very sparse questions and getting reasonable answers; those same queries do not work well as LLM prompts.

Researchers have identified best practices that produce much more accurate results. A couple of techniques that are useful are Roleplaying and Chain of Thought. Roleplaying refers to explicitly assigning a role to the chatbot: "You are an experienced compliance officer writing policies and procedures for a bank" or "analyze this case like Harvard Law Professor Cass Sunstein." The idea is to give the chatbot a lot of context. It is helpful to think of chatbots as new college interns: very smart but requiring guidance to produce quality work.

> **Roleplaying and chain-of-thought promptings are just two prompt engineering tips. There are dozens of other good ones just a quick web search away.**

Another useful technique is called chain-of-thought prompting. All you need to do is add the phrase "think step by step" to the end of a prompt. When you do so, the bot will explain the steps in its reasoning, which is essential for double-checking the work. It also leads to better results, which will not be any surprise to math teachers who have been admonishing students to show their work for centuries.

Roleplaying and chain-of-thought promptings are just two prompt engineering tips. There are dozens of other good ones just a quick web search away. While AI is rapidly improving at discerning what the questioner is looking for, there is still value in experimenting with learning how to write better prompts.

Opportunities Abound

Most financial institutions are at preliminary stages in evaluating opportunities to use Generative AI in their operations. Some of the areas where we are seeing the anticipated use of LLMs are in customer services. Large language models can interact with a bank's customers in very natural conversations. Depending on the data that the bank trains the LLM on, the chat bots can answer questions about customer accounts and even provide recommended product offerings and investment advice. Several large banks are working with internal LLM models to capture call center notes, organize information for investment advisors and organize other product data for customer service reps, with plans to roll out to more customer-facing uses as extensive testing addresses potential risks.

Banks are also assessing opportunities to improve internal operations. Generative AI capabilities enable new ways to analyze data. One practical use case for most organizations is to train LLMs on all the pockets of organizational information that employees need to access to do their jobs. Imagine spending less time looking for things and more time solving problems, with the information you need, quickly presented to you via an interactive chat with your friendly LLM bot.

LLMs also lend themselves to reading large regulatory documents and providing structured responses to aid a compliance or risk professional in digesting complex regulations and locating practical implementation guidance. The recent publication of the Small Business Lending rule by the CFPB illustrates these challenges. Between the regulation (complete with preamble and commentary) and the supplementary documents, there are over 1,200 pages of material to read and digest. A large language model, with structuring guidance from a legal expert, can digest all this material and provide useful responses to staff charged with understanding the regulation and producing policies, procedures, and implementation plans.

Yes, There's Risk

The continuing, rapid evolution of technology means that the opportunities for new and creative uses will continue to grow. According to a study from McKinsey, across the banking industry, GAI could deliver value equal to an additional $200 billion to $340 billion annually if the use cases were fully implemented.[1] With that kind of economic potential in the industry, we compliance and risk professionals need to start now to learn more about this technology and how to help manage risk for our organizations.

Regulators are also raising concerns that are important to address. Concerns around privacy of customer data, bias in algorithms and other UDAP/UDAAP concerns are on their minds. CFPB Director Chopra has reiterated that there are a wide range of potential threats to consumers' civil rights, and he reminds us that "there is no exemption in our nation's civil rights laws for new technologies that engage in unlawful discrimination. Companies must take responsibility for their use of these tools."[2]

Other institutional risks can result from employees using public versions of LLMs. When ChatGPT burst on the scene in late 2022, many of us logged on and were amazed at how it could author poems, suggest birthday party ideas for a nine-year-old, and devise creative holiday cocktail names. Then we started thinking about how it could help us with work. Helping draft emails, write a performance review, or suggest some creative marketing ideas for our yet-to-be-launched product. While it might seem harmless on the surface, you are exposing trade secrets and creating privacy concerns for your organization. It is important to read the fine print. Many existing chatbots have terms of service that allow the company to reuse user data to "develop and improve" their services.

Generative AI capabilities enable new ways to analyze data. One practical use case for most organizations is to train LLMs on all the pockets of organizational information that employees need to access to do their jobs.

STOCK.ADOBE.COM / N ON NE ON / FRESHIDEA / RARINLADA

A critical early step is to develop and implement an acceptable use policy. While many of your existing policies should apply to employee use of public LLMs, with the rapid evolution of this technology and the attraction it generates, it is better to be direct on what actions are permissible. It may not be obvious to someone using a public chat bot that they are exposing confidential information. So, an explicit policy that clearly states what your staff can and cannot do with LLMs is a prudent early step. Having employees acknowledge the policy and taking training courses on it are possible next steps.

Regulation on the Horizon

Nearly 200 AI-related bills were introduced in 31 different states this year according to the Business Software Alliance.[3] It is expected that 2024 will see even more activity. These bills respond to concerns about misuse of AI—including its ability to create fake images and audio "deepfakes"—and seek to define and put safeguards around AI's role in employment and credit decisioning. California has taken the lead with two agencies, the California Civil Rights Council and the California Privacy Protection Agency, in seeking to regulate automated decision-making.[4] Reviewing AI through the lens of privacy may seem surprising, but it is part of a long-running trend of technological change prompting expanded conceptions of privacy. There was not a need for a "right of publicity" until photography was invented. Nobody conceived of a "right to be forgotten" until the cost of documenting and archiving every youthful indiscretion became essentially free.

At the federal level, Senator Charles Schumer in September 2023, held the first of nine A.I. Insight Forums. These are closed-door events where tech leaders, such as Sam Altman, Bill Gates and Elon Musk, together with tech ethicists such as Tristan Harris and Dr. Rumman Chowdhury, addressed the assembled senators to explain the technology and their concerns for the future.[5] Notably, when Sen. Schumer asked the panel whether the government should regulate A.I., every hand was raised.[6]

While there may be broad agreement that regulation is needed, the federal government's recent track record for regulating technology does not inspire confidence that D.C. will get the job done. "We've yet to pass a meaningful bill, for example, protecting people's privacy on social media," said Sen. John Neely Kennedy (R-La.). "If we can't do it for an important … but less difficult issue like protecting privacy on the internet, I think artificial intelligence will be tougher."[7]

And finally, on October 30, 2023, President Biden issued an executive order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. The order lays out eight guiding principles and priorities, and tasks federal regulatory agencies to use these to advance and govern AI development. The overarching goal is to encourage responsible AI development while still creating a governing environment to mitigate the substantial risk. Financial services, as part of the critical infrastructure as defined by the Patriot Act, will be impacted by the regulatory activities that will emerge from this order.

Conclusion

How do you wrap up a banking love story? (You remember the title, right?) Generative AI will change so much about how we function in the world in the months and years to come. Like many great love stories, only time will tell how the relationship between Compliance and AI evolves. However, there is no time like the present to learn about Generative AI and the opportunities and risks that will need to be managed. Enjoy the journey. ■

ABOUT THE AUTHORS

SAM HOLLE, JD, CRCM, is an Associate Director with the Regulatory Compliance Analysis team at Wolters Kluwer. With over 15 years of industry experience, Holle helps financial institutions leverage advanced technologies to make the complex simple. Reach him at **linkedin.com/in/sam-holle-76178614/**.

KRIS STEWART, JD, CRCM, is a Senior Director in the Decision Science team at Wolters Kluwer. With over 30 years of industry experience, Stewart works with customers to find ways to apply advanced technologies to solve complex problems. Reach her at **linkedin.com/in/krisstewartjd/**.

Endnotes

1. McKinsey | The economic potential of generative AI: The next productivity frontier
2. https://www.consumerfinance.gov/about-us/newsroom/director-chopra-prepared-remarks-on-interagency-enforcement-policy-statement-artificial-intelligence/
3. Software industry alliance BSA says states are regulating AI faster than DC (axios.com)
4. Bill Text: CA AB331 | 2023-2024 | Regular Session | Amended | LegiScan; https://cppa.ca.gov/regulations/pre_rulemaking_activities_pr_02-2023.html
5. https://techpolicy.press/us-senate-ai-insight-forum-tracker/
6. https://www.democrats.senate.gov/news/press-releases/majority-leader-schumer-floor-remarks-on-the-success-of-the-inaugural-ai-insight-forum
7. https://www.washingtonpost.com/politics/2023/09/14/3-takeaways-senators-private-huddle-with-tech-execs-ai/

ABA MEMBER RESOURCES

ABA Webinar: Tech Innovation and AI
aba.com/training-events/online-training/tech-innovation-and-ai-managing-risk-in-the-rapidly-evolving-landscape

ABA Advocacy: Executive order on artificial intelligence
aba.com/advocacy/policy-analysis/eo-artificial-intelligence