

BANKING ON

Risk, readiness and the next

BY KRIS STEWART AND AMY BAISDEN, JD



Artificial intelligence (AI) Adoption is Surging.

AI IS NO LONGER A DISTANT PROMISE IN BANKING

— it's reshaping everything from fraud detection to underwriting.

According to a recent report on AI use in financial services, up to 91% of financial services companies are in the process of adopting or already utilizing AI for a wide variety of operational activities — from portfolio optimization, fraud detection, and risk management to marketing and customer engagement.¹

Yet, while AI adoption is surging, risk and compliance professionals face the daunting challenge of keeping abreast of rapidly changing technology, a patchwork of emerging regulatory expectations, and the need to help guide their institutions to remain innovative and accountable.

What is AI? Demystifying the basics

At the most fundamental level, AI is any system or tool capable of performing tasks that typically require human intelligence, such as learning, reasoning, perceiving, and problem-solving.

To make informed decisions about AI in banking, risk and compliance professionals need a clear grasp of its foundational pillars — Machine Learning, Deep Learning, Natural Language Processing, Generative AI, and Large Language Models. By breaking down these basics, this section provides the practical understanding needed to spot risks, ask the right questions, and guide responsible adoption.

Machine Learning

Machine Learning is a core subfield of AI and is the foundation of many of the most well-known AI systems and tools, from forecasting models to large language models and other generative tools.

In machine learning, a model is “trained” on a dataset that contains representations of the types of problems and tasks that the model will be expected to evaluate and perform. The model then employs predictive algorithms to find patterns in the data that it then uses to infer the correct response to the problem or task that it is asked to resolve. As a user interacts with the tool, it can also learn from that input data and further improve its performance.

Deep Learning

Deep Learning is a specialized branch of machine learning that uses multilayered artificial neural networks — algorithms designed to mimic the way the human brain processes information to perform complex calculations and map inputs to outputs — to process large volumes of data.

Where traditional machine learning systems utilize simple neural networks with no more than one or two computational layers, deep learning systems typically use hundreds or even thousands of layers.

Deep learning is commonly used in systems and tools that perform heavy analytical tasks or provide complex automation, such as digital and voice-enabled assistants, generative AI and, specifically in the context of the financial services industry, fraud detection and prevention systems.



AI

frontier

Natural Language Processing

Natural Language Processing is a type of machine learning that uses computational linguistics — utilizing algorithms to understand and interpret written or spoken language — and statistical modeling — mathematical models that determine the probability of relationships between elements in a given data set — to enable machines to read, understand, analyze, and generate human language.

Chatbots, virtual assistants, and voice-operated systems and tools are common examples of systems that utilize natural language processing.

Generative AI (GenAI)

Generative AI (GenAI) is a type of machine learning that relies on the computational linguistics of natural language processing and the multilayered neural networks of deep learning to create original content from a user prompt. Depending on the capabilities of the model, this generated content may be text, an image, or even audio-visual content.

For example, a user may ask a GenAI model to “summarize the top regulatory risks for U.S. banks related to third-party vendors in 2025” to begin researching the topic. Later, they might ask the model to “generate a picture of a diverse group of financial professionals reviewing compliance reports and data dashboards in a modern bank boardroom” to use in a presentation.

Large Language Models (LLMs)

Large language models (LLMs) are GenAI models trained on large amounts of textual data that use deep learning and natural language processing to understand and generate natural language text. LLMs utilize a specialized neural network architecture that renders them especially proficient at understanding patterns, contexts, and relationships between words and applying statistical prediction algorithms to predict the probability of the next sequence of words.

Well-known examples of LLMs include OpenAI’s ChatGPT, Google Bard, Anthropic’s Claude, and Microsoft CoPilot.

LLMs are used for tasks such as drafting, wordsmithing, summarization, translation, and conversational simulation. The models can be fine-tuned for domain or industry-specific applications such as legal analysis or financial reporting. In banking risk and compliance, LLMs can be tailored to support compliance reviews, regulatory analysis, and reporting by quickly processing complex regulations and providing clear summaries or recommendations.

AI Myths and Realities

What AI is not

Understanding what AI is also requires an understanding of what it is *not*.

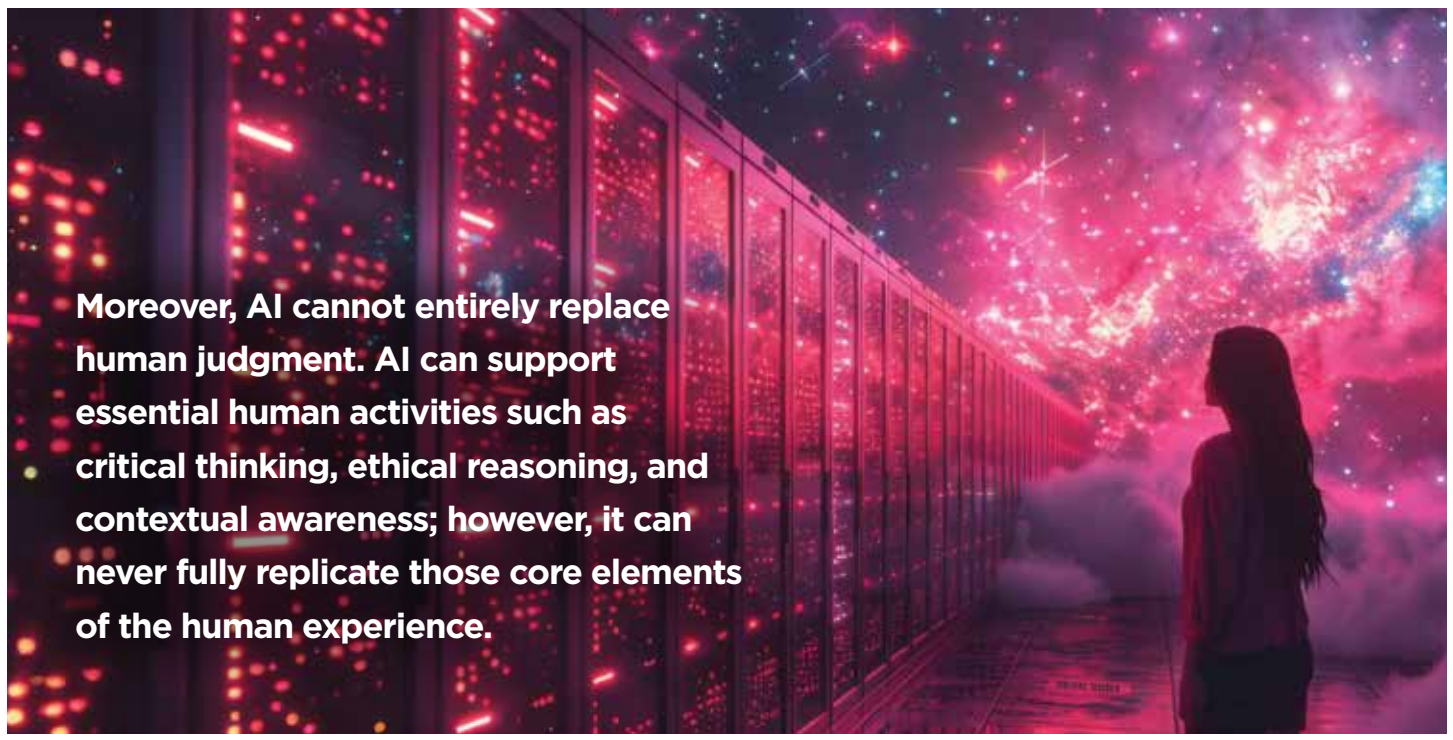
AI is not infallible or even universally applicable. An AI model, like any other system or tool, is constrained by its data and its design.

AI is not sentient — that is, it cannot perceive or feel. It does not have consciousness or self-awareness, and it cannot have motive or intent.

Moreover, AI cannot entirely replace human judgment. AI can support essential human activities such as critical thinking, ethical reasoning, and contextual awareness; however, it can never fully replicate those core elements of the human experience.

Importance of clarity and realistic expectations

It is important neither to overestimate nor to underestimate AI’s capabilities. Overestimation can lead to design failure, validation and oversight errors, misuse, and mistrust. However, underestimation may be equally troublesome as it might leave an organization vulnerable to reputational, compliance, or prudential risk that can be mitigated through thoughtful adoption and use of AI systems and tools.

A woman with long dark hair is seen from the back, looking out at a long, glowing server rack. The server rack is filled with red and orange lights, and the background is a vibrant, colorful cosmic scene with stars and nebulae. The text is overlaid on the left side of the image.

Moreover, AI cannot entirely replace human judgment. AI can support essential human activities such as critical thinking, ethical reasoning, and contextual awareness; however, it can never fully replicate those core elements of the human experience.



It is important neither to overestimate nor to underestimate AI's capabilities. Overestimation can lead to design failure, validation and oversight errors, misuse, and mistrust. However, underestimation may be equally troublesome as it might leave an organization vulnerable to reputational, compliance, or prudential risk that can be mitigated through thoughtful adoption and use of AI systems and tools.

Putting AI to Work

Armed with a practical grasp of AI, risk and compliance professionals can better appreciate the ways these technologies are being deployed across the financial sector. In several areas, AI-based solutions are helping banks achieve revenue growth, efficiency, customer satisfaction, and more accurate risk management.

Fraud detection and prevention is an area where machine learning models help analyze transaction patterns to detect and flag anomalies in real time. Where traditional rule-based systems (explicitly programmed “if-then” statements) would flag a transaction over \$10,000 for manual review, an AI-based solution might flag a \$9,500 transaction as suspicious because it deviates from a customer’s typical behavior. As a result, most global financial institutions utilize AI-driven systems, and as of late 2025, these systems were intercepting a reported 92% of fraudulent activities before they are approved.²

Credit underwriting is another area where ML technology helps to improve loan approval accuracy. The use of AI is helping financial institutions go beyond traditional credit scoring models to include more data points, including non-traditional sources such as online behavior, employment histories, education, and accessing data from different types of online payment systems. This allows lenders to reach unbanked or underbanked markets while still maintaining risk management practices.

The rapid emergence and evolution of generative AI, particularly large language models (LLMs), are being used by regulatory and compliance professionals to manage complex legal and regulatory frameworks more efficiently. Properly trained LLMs can ingest and analyze large volumes of regulatory data, creating clear summaries, extracting legal obligations, and analyzing existing risks, controls, and procedures to aid compliance and risk officers in swiftly addressing existing and evolving regulatory requirements.

Regulatory Guidance: From Guardrails to Greenlights

Federal

Following a deregulatory pivot by the current administration — for example, Executive Order 14148 and Executive Order 14179, both signed in January 2025, revoked or otherwise directed the identification, suspension, revision, or rescission of prior AI mandates — the Federal Reserve’s SR 11-7 (2011) remains the most significant and clearest statement of model risk management (MRM) regulation.

Nonetheless, even in the absence of any significant federal statutory or regulatory frameworks specifically applicable to AI, banking regulators have previously demonstrated consistent oversight of AI technology through application of existing legal standards.

The Office of the Comptroller of the Currency (OCC) incorporated AI findings in its enforcement actions in 17 matters since fiscal year 2020, and in its Fall 2023 Risk Perspective, the bureau explicitly stated:

“Advances in technology do not render existing safety and soundness standards and compliance requirements inapplicable. Although existing guidance may not expressly address AI use, the supervision risk management principles contained in OCC issuances provide a framework for banks that implement AI to operate in a safe, sound, and fair manner.”

In an August 2024 comment on Request for Information (RFI) on Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector, the Consumer Financial Protection Bureau (CFPB) asserted that:

“Although institutions sometimes behave as if there are exceptions to the federal consumer financial protection laws for new technologies, that is not the case. Regulators have a legal mandate to ensure that existing rules are enforced with respect to all technologies, including those marketed as new or novel.”

Following the RFI, the Treasury Department, in issuing its Report, recommended that:

“Financial firms prioritize their review of AI use cases for compliance with existing laws and regulations before deployment and that they periodically reevaluate compliance as needed.”

State-related actions

In the absence of comprehensive federal legislation and in light of the deregulatory actions and reduction in enforcement activities under the current administration, the states have begun stepping up to offer regulatory solutions — with California, Colorado, Utah, and Texas leading the way.

Of note for financial institutions are the enactments in Colorado and Texas.

The Colorado AI Act (Consumer Protections for Artificial Intelligence), now set to take effect on June 30, 2026, targets “high-risk systems” that influence consequential decisions like loan approvals and credit scoring. It requires institutions deploying such systems to implement risk management programs, conduct impact assessments, provide transparency notices to consumers, and self-report algorithmic discrimination to the Attorney General. Penalties can reach \$20,000 per violation, or \$50,000 if senior citizens are affected. Notably, the Act does include language that deems a financial institution compliant with the Act if the institution is subject to and compliant with other federal laws or regulations that provide substantially similar protections. Legislative amendments are under consideration, which may refine definitions and compliance obligations, so financial institutions should monitor developments closely.

The Texas Responsible Artificial Intelligence Governance Act, which became effective on January 1, 2026, applies across industries to AI systems used in business, commerce, healthcare, government, and consumer interactions. The Act broadly prohibits any discriminatory or manipulative practices and improper handling of personal information in the use of AI — exempting financial institutions only to the extent that they are following existing industry-specific state or federal regulations — and also requires a disclosure to consumers who may interact with AI systems and tools.

For risk and compliance professionals, the message is clear: while formal rulemaking may be slow, particularly at the federal level, legacy expectations still endure. Institutions must remain accountable to existing expectations, while also keeping an eye on the ever-evolving regulatory landscape, particularly at the state level.

Given the complexity of the legal landscape surrounding the development, implementation, validation, and use of AI in the financial services industry, it is incumbent upon risk and compliance professionals to develop a working understanding of AI technology sufficient to effectively evaluate



By pairing human insight with strong governance, compliance teams can turn AI from a regulatory challenge into a catalyst for resilience and trust.

risks, appropriately assess compliance obligations, and provide meaningful oversight of their institutions’ programs.

Shifting compliance from reactive to proactive Facing the Risks

A proactive risk and compliance program strengthens your institution’s resilience and integrity in an increasingly complex regulatory environment. In this space, AI systems and tools can provide transformative gains in operational efficiency and cost reduction.

As discussed above, AI is currently used to automate repetitive tasks, perform rapid, complex data analysis, and streamline workflow. This reduces manual effort, improves accuracy, accelerates decision-making, and enables compliance teams to focus on higher or more material risks.

Automation and scalability also allow risk and compliance teams to shift from a reactive to a proactive orientation to investigation, audit, monitoring, and reporting activities. Rather than relying on manual effort, samples, and spot checks, AI systems and tools can be employed to provide continuous, real-time, and on-demand results.

While AI’s foundational pillars offer significant benefits for risk and compliance, they also introduce new risks and challenges. The following section highlights what teams need to watch for as adoption accelerates.

Data quality, privacy, and governance

AI is only as good as the data it receives. If the data is incomplete, inconsistent, incorrect, or outdated, then AI outputs will be similarly compromised, flawed, and unreliable. Lack of or insufficiency of data standards, insufficient validation processes, and fragmented systems of unintegrated legacy software can all contribute to a data environment that undermines the utility of even the best AI systems and tools.

Data privacy concerns must also remain at the forefront. Sensitive customer data and information is protected by strict privacy laws that continue to apply even in an age and environment where new, innovative tools and processes are changing how we think about collecting, processing, accessing, and using that data.

Bias, fairness, and ethical considerations

AI models may unintentionally replicate or amplify historical biases embedded in training data, risking discriminatory outcomes in credit decisions, fraud detection, and customer interactions. Continuous testing, transparent evaluation, and mitigation strategies are necessary to promote fairness and equity. Ethical AI governance involves ensuring systems operate transparently and accountably, respecting consumers’ rights and regulatory requirements.

Transparency, explainability, and accountability

Transparency, explainability, and accountability are essential components of trustworthy AI development, deployment, and use.

Explainable AI (XAI) aims to open the “black box” of complex AI models through the application of specific processes, techniques, and methods that ensure that each decision made in the design of the model, how data is processed by the model, and how the output is rendered can be followed and understood.

These concepts are at the heart of trustworthy and responsible AI use. Ensuring that employees, customers, and regulators understand how technology is used and how it works is essential in fostering trust, accountability, and the expected

degree of rigorous oversight necessary for these tools to be used effectively in the financial services industry.

AI's next chapter

As AI rapidly evolves, risk and compliance teams must look beyond today's capabilities. While it is always risky to break out the crystal ball, especially given the speed of change we are experiencing, here are some of the emerging trends and innovations that can be expected to reshape the industry in the coming years.

Agentic AI: Autonomous intelligence in action

Agentic AI is a type of artificial intelligence that can take action and make decisions on its own. Unlike traditional AI models that require explicit instructions or operate within narrowly defined parameters, agentic AI leverages the foundational pillars of Machine Learning, Deep Learning, and Natural Language Processing to figure out what needs to be done and how to do it.

In banking, this means agentic AI could operate in fraud detection to continuously learn and adapt to new types of attacks and take instant counteractive measures. Agentic systems can learn from both historical and real-time data — making adjustments from subtle patterns that static rules would miss — and they can score the identified risks in real time and automatically adjust account limits, block transactions, or escalate for human analysis.³

Agentic AI systems can also be set up as a network of agents designed to interact with each other. This is referred to as an Agentic AI Mesh — a distributed network of agentic AI agents that collaborate, reason, and act independently across multiple systems and workflows within an organization.

Some banks have active programs to adopt agentic AI mesh technology. In one example, a robust mesh helps enable autonomous, interconnected agents to manage fraud, compliance, and client services in real time.⁴ This architecture aims to transform banking operations from reactive to proactive while maintaining transparency, auditability, and human oversight — key requirements for compliance officers.

Emerging risks: AI workslop & AI debt

As agentic AI systems drive new efficiencies in banking, they also bring fresh risks that require attention. Two standouts are AI workslop, the proliferation of low-quality, unchecked outputs from generative AI tools — and AI debt — the buildup of poorly governed or undocumented AI systems.

In the rush to adopt AI, institutions can move fast without implementing sufficient governance and controls. When automation is scaled without sufficient oversight and planning, AI workslop can appear. Automated reports or chatbot responses that look polished but contain errors or lack substance go through the operational process unchecked, forcing teams to spend extra time correcting or clarifying the work.

Meanwhile, AI debt can accumulate if organizations deploy AI solutions without proper oversight, leading to compliance gaps and operational headaches as outdated or unmonitored systems

pile up. To deploy a solution rapidly, documentation and testing may be limited or rushed. If good operational controls and governance are lacking, then the ability to timely identify and resolve issues as they arise can create a mess to clean up and result in compliance and audit risks that may require costly remediation efforts.

Unlocking human/AI synergy: Best practices for compliance teams

As artificial intelligence transforms financial services, compliance teams are uniquely positioned to lead the way in building resilient, ethical, and innovative organizations. By actively shaping how AI supports regulatory integrity and operational excellence, professionals can turn potential risks into opportunities for growth and trust. Thoughtful collaboration between human expertise and AI capabilities — grounded in strong governance — ensures that technology becomes a true partner in advancing the industry.

The following best practices will help compliance teams harness AI's benefits, navigate change confidently, and set the standard for responsible adoption.

- Define clear roles and responsibilities for both human and AI participants.
- Establish cross-functional governance that includes compliance, legal, technology, risk, and audit professionals.
- Maintain transparency and explainability by using AI models that can be explained to non-experts and ensure documentation exists and is maintained.
- Enable Human-in-the-Loop oversight by ensuring that humans can intervene, override, or halt AI-driven decisions and ensure that feedback loops for continuous improvement are in place.
- Monitor and learn — monitor regulatory changes that are sure to come and stay abreast of emerging technology and the associated risks.

Embracing AI is not just about adapting to new technology; it's about leading the transformation of financial services with integrity, foresight, and purpose. By championing responsible innovation and continuous learning, risk and compliance professionals can shape a future where AI strengthens trust, enhances resilience, and unlocks new possibilities for their institutions and the customers they serve. ■

ABOUT THE AUTHORS:

KRIS STEWART is a Senior Director with the Compliance Program Management solutions team with Wolters Kluwer Financial & Corporate Compliance. Reach her at Kris.Stewart@wolterskluwer.com.

AMY BAISDEN, JD, is a Specialized Consulting Manager with the Compliance Center of Excellence with Wolters Kluwer Financial & Corporate Compliance. Reach her at Amy.Baisden@wolterskluwer.com.

Endnotes

1. <https://blogs.nvidia.com/blog/ai-in-financial-services-survey-2024/>
2. <https://coinlaw.io/ai-in-banking-statistics/>
3. <https://www.gofast.ai/blog/machine-learning-fraud-detection-case-study>
4. <https://aiexpert.network/ai-at-jpmorgan/>

ABA MEMBER RESOURCES

Banking Topic: AI

aba.com/banking-topics/technology/artificial-intelligence

Workshop: AI in Banking

aba.com/training-events/online-training/ai-banking-workshop