Chartis RiskTech AI 50 2024

Table of contents

1. Foreword

Welcome to the inaugural Chartis RiskTech AI 50 ranking and research, which explores the broad landscape of artificial intelligence (AI) development and adoption in the financial services. As part of this comprehensive and extensive research, we examine specific training techniques and machine learning (ML) architectures, and look beyond the hype to identify not only those areas that are maturing, but also those offering promising growth and innovation.

Indeed, while the relative levels of maturity of AI-based applications and tools vary widely, we also believe that AI and AI tools will increasingly become core components of risk and analytics, co-existing and blending with other statistical and mathematical elements.

For now, though, this is an area undergoing rapid evolution, and one thing is clear: AI and ML are now part of the mainstream, and they will continue to have profound impacts on businesses of all types. As new tech providers continue to emerge in the market, with ever more innovative and potentially game-changing offerings, Chartis' goal remains to analyze this dynamic and changing landscape, and recognize the companies operating within it.

With that in mind, please join me in congratulating the featured vendors and category winners.

Maryam Akram, Research Principal

2. Introduction

In its inaugural edition, the Chartis RiskTech AI ranking and research report explores the broad landscape of AI adoption in financial risk management, and provides some context and highlights from our research, as well as some important insights into the current AI landscape. As such, RiskTechAI is a comprehensive study that offers both a panoramic view of the industry and a detailed examination of specific ML architectures and methods. Chartis is adopting a pragmatic approach, assessing the current state of AI in risk management by identifying areas of maturity, promising growth and innovation. From emerging large language model (LLM)-based applications to established credit-scoring and fraud-detection ML models, the report highlights the depth and diversity of AI applications in the field.

Chartis acknowledges that while some hype-driven generative AI (GenAI) use cases in financial risk management may not reach maturity, the deep learning methods and emerging infrastructure underpinning them represent crucial and lasting innovations. Over the years, advances in AI have significantly reshaped the technological and operational frameworks across various financial service business lines. Reflecting this transformative trend, data science teams and ML expertise have become increasingly integral to vendors and institutions.

Key takeaways from the research include a focus on AI projects that yield measurable real-world impacts and the critical role of scalable deployment. Building on this foundation, our 2025 research will broaden to include a wider array of firms and topics, such as:

- Technical trends in AI and their applications within financial services.
- Ongoing developments in AI governance and policy.
- Tracking of investments, in (for example) start-ups, hardware and infrastructure.
- A detailed look at the emerging AI technology stack.

3. Overview

Key ranking highlights and insights

Notable achievements and success stories

Our RiskTechAI research is based on several key principles:

- **Impact**. High impact scores were awarded to vendors whose infrastructure or models/model-based solutions deliver measurable, real-world benefits. Notable use cases include document processing, load forecasting, fraud detection, credit risk analytics and IT risk analytics. In some cases, vendors provided a suite of customizable, pre-trained models embedded within solutions and platforms and designed to support the full model lifecycle.

- **Deployment**. Vendors that excelled in deployment demonstrated robust scalability, diverse compute support and comprehensive AI infrastructure, spanning high-performance computing (HPC), cloud graphics processing unit (GPU) clusters and interoperability across modeling frameworks. Regarding data, notable achievements included low-latency data processing, advanced database architecture and efficient low-latency infrastructures.

- **Strategy**. Vendors with outstanding strategy scores demonstrated significant investment in both classical ML and GenAI, leading to well-defined AI products. LLM-based applications included domain-specific question-and-answer systems and text-to-SQL data search tools. Key differentiators were access to proprietary data, fine-tuned models and integration. A clear understanding of the limitations and contextual suitability of various analytical techniques was also crucial.

- **Innovation**. High innovation scores reflected a range of initiatives, including promising GenAI pilots, pioneering infrastructure design and novelty of approach. Innovation also reflected research and development investment, innovation labs and academic partnerships.

AI investments and practical use cases

- As pilot GenAI projects transition to production, vendors are focusing on revenue growth and measurable return on investment (ROI). Projects that address critical pain points, such as high-volume document processing, are more likely to yield successful outcomes.

- Banks are increasingly developing proprietary GenAI initiatives, often in collaboration with frontier research labs.

- Open-source LLMs have shown benefits in terms of security and customizability, yet adoption remains generally slow for risk management pilot projects. However, regional concentrations of open-source development and implementation are emerging. China and India are developing considerable expertise and pilot projects in this space.

- ML techniques in financial crime detection are moving toward standardization. Here, scalability and data accessibility are essential for implementing effective and efficient solutions.

- Code generation and automation pilots continue to demonstrate the most immediate and substantial productivity gains, making them a leading use case across industries.

State-of-the-art AI

- **Innovation beyond scaling**. The broader ML industry is moving beyond scale alone, focusing on methods to optimize training and test efficiency. Key areas include:
  - **Test-time compute and chain of thought**. These are increasingly significant dynamics in LLM logic and reasoning performance improvements.
  - **Fine-tuning**. Techniques such as direct preference optimization (DPO) and reinforcement learning from human feedback (RLHF) are part of efforts to refine model outputs.
  - **Architectural innovations**. New model architectures aim to reduce computational costs and enhance speed during deployment. **Mamba** and the Multi-Head Latent Attention process used for the **DeepSeek-V2 model** are emerging to enhance the efficiency and scalability of LLMs without compromising performance.
  - **Agentic LLMs**. Future workflows may benefit from agentic LLMs, which could enable more integrated, context-aware processing that addresses memory, query context, multi-step actions and data access. Despite these benefits, persistent security and accuracy issues remain barriers to their production-level use in industry.
  - **Specialized models for specific tasks**. Experiments with smaller models fine-tuned on targeted datasets are enhancing efficiency and accuracy for specific downstream applications.
  - **Retrieval-Augmented Generation (RAG)**. RAG is becoming a standard practice for data integration, offering efficient retrieval mechanisms during model deployment.
  - **Mitigating hallucinations**. Researchers are exploring methods to address hallucination issues. These include using knowledge graphs and **Retrieval Interleaved Generation (RIG)**, which integrates real data into responses to improve reliability.

AI infrastructure

- **Hardware competition and geopolitical stakes**. As companies vie to disrupt Nvidia's AI hardware market dominance, government interventions like the Chips and Science Act underscore the strategic significance of AI hardware development – including at the geopolitical level. Indeed, multiple tech companies have met with the White House to discuss and advocate for the future power requirements **needed to support** evolving AI workload demands.
- **Transformation of the AI technology stack**. The development of GenAI has transformed the AI technology stack, introducing new components and layers and popularizing technologies like vector databases.
- **Specialized hardware and application frameworks**. Beyond hardware innovations like memory optimization, an ecosystem of application frameworks has emerged, offering essential tools for orchestration, integration and support across AI workflows.
- **Intense hyperscaler competition and new GPU cloud companies**. The hyperscaler landscape for model inferencing remains highly competitive. Fast-growing GPU cloud specialists are also competing for market share, as demand continues to grow.

AI governance and policy

- As financial regulatory bodies evaluate and debate the evolving risks of AI, markedly different regulatory approaches – such as the EU AI Act and California's recently vetoed AI Safety Bill (SB 1047) – focus on governing the technology itself.
- Although bodies like the Financial Conduct Authority (FCA) state that they 'do not regulate technologies,' the rapid evolution and increasingly opaque capabilities of AI systems are prompting a reevaluation of existing frameworks to maintain public trust in financial systems.
- Financial regulators are confronting a spectrum of AI-related risks, including market integrity, cybersecurity, consumer and institutional fraud, and discriminatory model processes.
- The debate over responsibility and accountability for AI-related harms extends beyond the financial sector, as policy stakeholders consider the roles of developers, users and companies. Compute thresholds are emerging as a critical factor in shaping AI regulations, although their effectiveness as governance benchmarks is contested.

4. Implementing innovation: the practical use cases of AI

A use case-driven approach

Chartis' perspective on ML trends in risk management – whether they are cutting-edge or well-established – is rooted in a use case-driven approach (see Figure 1). Our research and industry evaluations for both emerging and established technologies are focused on real-world applications. Within the financial services industry, risk management and customer engagement historically have been areas of especially high ML adoption. The COVID-19 pandemic **further fueled investment** in customer engagement applications as the rate of digitalization in banking accelerated. In risk management, a range of supervised, semi-supervised and unsupervised ML methods are being applied alongside other analytics, tools and technologies.

Figure 1: Mapping ML techniques to use cases

Source: Chartis Research

ML tools and systems generally fall into three broad categories: prediction (including classification), generation and clustering or dimensionality reduction. However, the array of ML techniques deployed in risk management is diverse and the maturity and industrialization varies by application. For instance, graph neural networks (GNNs) are utilized for entity resolution, while convolutional neural networks (CNNs) can be adopted for time-series forecasting and transformer-based embedding models can be leveraged for document processing tasks.

Despite the diversity of these techniques and methods, certain domains, such as financial crime prevention, are seeing a trend toward standardization. Best-practice clustering and classification algorithms are now routinely applied to such areas as anti-money laundering (AML), transaction monitoring and Know Your Customer (KYC) processes. Although such issues as false positive rates continue to drive the development of ML techniques, the primary success of applications is often based on data availability, scalability and cost.

The structure and functions of bank units play a critical role in how ML solutions are deployed. Factors such as explainability requirements, regulatory constraints and prevailing modeling paradigms heavily influence the development and adoption of ML in specific areas of financial institutions. For instance, organizational silos in the back and middle office – driven by functional specialization, regulatory compliance and legacy systems – can hinder process automation and slow the integration of ML technologies.

Credit and fraud (overlapping techniques and analytical goals)

Efforts to collaborate and adopt enterprise-wide data frameworks are increasing as banks seek to improve automation and digitalization. In retail banking, the large volume of transaction data (both real-time and historical) has led to significant investment in ML applications, particularly in credit scoring and fraud detection (see Figure 2). The alignment of data, regulatory pressures and shared business goals can foster interdepartmental collaboration, enabling the co-development of ML solutions in such areas as fraud prevention and loan underwriting. Despite the success of ML techniques in these areas, challenges remain in aligning up-to-date models with representative historical data.

Figure 2: Areas of overlap: machine learning in fraud detection and credit risk

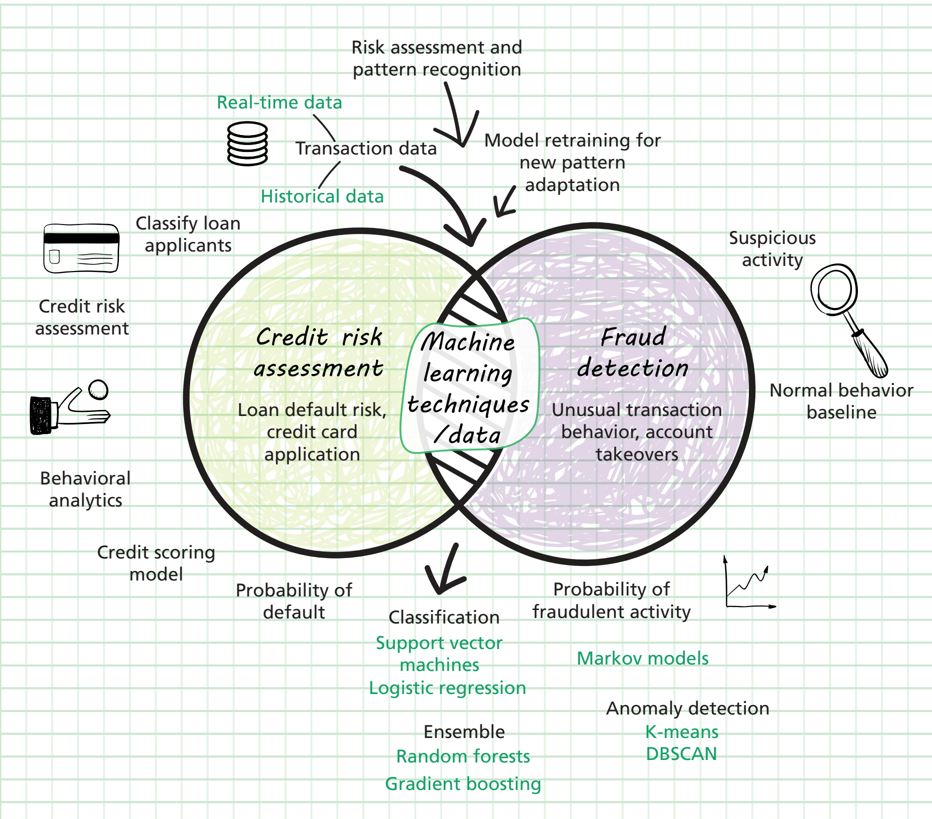

Source: Chartis Research

Supervised fraud detection algorithms rely on historical examples of fraudulent transactions to identify patterns and predict future fraud accurately. However, these datasets are often imbalanced, with fraudulent transactions representing only a small fraction of the total available data. This imbalance can skew model performance, as systems can become biased toward predicting non-fraudulent behavior. Frequent retraining is necessary, as historical data may not fully capture emerging fraud tactics and evolving patterns. High false positive rates also pose a significant challenge, impacting customer experience and operational efficiency. To address these issues, firms often employ a combination of supervised and unsupervised techniques, leveraging labeled transaction data alongside anomaly detection to uncover new or shifting fraud patterns. The rapid uptake of consumer GenAI solutions is also driving new fraud challenges, as tools enable the production of convincing fake documentation and synthetic identities. Technology firms are working on watermarking techniques for generated content, but their current and long-term robustness and effectiveness are unclear.

Making trade-offs: selecting the right ML approach

Choosing the appropriate ML approach depends on several factors, including the nature and volume of available data, the specific goals of the analysis and the computational resources at hand. These considerations guide the selection of algorithms, each of which involves its own trade-offs. For instance, K-means and hierarchical clustering are two widely used clustering algorithms, each with its own strengths and limitations.

Clustering algorithms, as part of unsupervised learning techniques, are applied across a broad spectrum of use cases. These include network analysis for cybersecurity, business process mining to identify inefficiencies or bottlenecks, and transaction monitoring to detect anomalies in financial data. The choice of clustering algorithm depends on the problem's specific requirements. K-means, while favored for its scalability and computational efficiency, assumes clusters of uniform size and shape and is sensitive to outliers. Hierarchical clustering, though more robust to outliers and capable of accommodating clusters of varying sizes, is computationally more intensive and often better suited to exploratory analysis rather than large-scale clustering applications.

Use case scalability and industrialization

Chartis' AI research also emphasizes the importance of the surrounding technology infrastructure that enables AI applications to scale effectively. As firms rapidly adopt LLMs, there is an increasing need for infrastructure that can support AI workloads at scale. This demand is driving innovations in computational capacity, cloud deployment and data management systems.

While the rapid adoption of GenAI has led to a surge in proof-of-concept initiatives, the development of fully scalable, customer-ready applications remains a more gradual process. The gap between proof-of-concept models and production-ready systems is not limited to GenAI applications and well-established areas of ML deployment can also face production challenges. For instance, applications like fraud detection require not only high accuracy but also low latency, robust throughput capabilities and the ability to handle 'bursty' workloads. Achieving industrial-scale ML deployments demands more than simply refining algorithms. Supporting infrastructure (such as distributed computing frameworks, low-latency data pipelines and robust data storage architectures) plays a crucial role in ensuring that ML systems can operate at scale.

AI case studies

Enterprise documents: enhancing document management and transparency with AI

Effective document management and transparency are essential components of modern enterprise operations. Loan agreements, such as leveraged loans, project finance loans and syndicated loans, often involve extensive and complex documentation with low levels of standardization. Loan agreements are growing progressively longer; for instance, leveraged loans are increasing due to the growing complexity of financial structures, greater flexibility in contractual provisions and stricter regulatory and compliance requirements. Commercial due diligence in the US also highlights the document management challenges faced by institutions. The rapid expansion in private credit markets is further driving demand for the collection and analysis of capital statements, contributing to the broader complexities in enterprise document management.

Automation challenges

Lengthy contracts and documents filled with detailed terms, conditions and irregular layouts present significant challenges for automated analysis. Over time, the level and sophistication of document transparency have evolved from manual document processing to advanced natural language processing (NLP). While optical character recognition (OCR) is designed to extract text from scanned images or documents, modern NLP tools go beyond text extraction, interpreting the meaning of the text in context and analyzing its nuances. Earlier models, such as bag-of-words (BoW), struggled to capture the meaning or context of words, but developments like Word2vec (2013) marked a significant step forward by encoding the semantic relationships between them.

Before bidirectional encoder representations from transformers (BERT), released by Google in 2018, catalyzed a paradigm shift in NLP, embeddings from language model (ELMo), released in February 2018, played a key role by capturing both the syntactic (structural) and semantic (meaning) nuances of text, allowing for improved contextual understanding. BERT advanced this further by introducing bidirectional context, enabling models to consider both the preceding and following words in a sentence.

Periodic advances in NLP have driven innovation in document transparency, with vendors seeking new methods to capture text from complex document layouts and low-quality images of legal contracts. OCR and NLP solutions often need to be customized for specific use cases, incorporating domain-specific taxonomies and recognizing entities and contractual details that are unique to loan agreements and financial documents.

A multimodal problem

Enterprise document processing presents a fundamentally multimodal challenge, requiring the integration of both text and visual data. The goal is to create tools that can effectively synthesize vast amounts of information while providing efficient ways to manage and understand complex, non-uniform documents, such as credit agreements, financial reports and legal contracts.

The evolution of OCR and image processing techniques

Over time, the methods used for OCR have diversified and advanced. While early OCR relied on such techniques as hidden Markov models (HMMs) for simpler text recognition, CNNs enhanced OCR for more complex image recognition tasks. And transformer-based architectures now enable sophisticated multimodal processing that integrates both text and visual data.

However, fine-tuning models for specialized document types presents ongoing challenges due to the limited availability of domain-specific datasets. The large datasets used for pre-training CNNs and transformers do not always translate directly to niche applications with smaller, highly specialized fine-tuning datasets. While this data mismatch can reduce model effectiveness, advances in transfer learning, zero-shot and few-shot learning are increasingly bridging this gap. Data augmentation and synthetic data generation have also been used to enhance model performance for specific tasks.

Emerging approaches: task-specific and multimodal models

Research into task-specific modeling and multimodal GenAI architectures continues to advance, with initiatives from major financial and technology firms leading the way. Projects such as JP Morgan's **DocLLM** and Bloomberg's **BloombergGPT** are notable for leveraging vast, domain-specific datasets – ranging from legal contracts and financial statements to enterprise documentation – to tailor AI models for specialized applications within finance. DocLLM, in particular, underscores the move toward multimodal approaches, combining text and visual data to better analyze and interpret complex, non-uniform documents in enterprise settings. Meta has also contributed to this space with open-weight models in its **Llama** series (notably the 11B and 90B).

Challenges in domain-specific applications

While these multimodal models have strong potential, implementing them effectively in financial document processing requires specific domain expertise, especially for complex or highly specialized tasks. Analyst support and knowledge of financial regulations are essential to interpret nuanced document content accurately and to ensure compliance in regulatory-heavy applications. Although efforts to mitigate hallucination using methods like RAG have been somewhat successful, hallucination continues to be a persistent issue in text generation tasks. Vendors that focus on accuracy measures and the development of validation protocols will have a competitive advantage in this space.

Natural language search dominates

Natural language search and document summarization have become standard features in GenAI-powered copilots. However, solutions offering measurable ROI often need to go beyond marginal productivity gains. Achieving success can depend on whether the copilot solution addresses a pressing client pain point, typically driven by specific document characteristics or high-volume demands. For instance, automating due diligence for investors or efficiently collecting structured and unstructured data on credit information and fixed-income assets are promising use cases in which productivity gains align closely with tangible client needs. AI startups that offer enterprise search are also developing in this space, utilizing content ranking techniques, RAG and vector search.

Traditional distributed search engines rely heavily on indexing and keyword matching to deliver real-time results. While index-based systems can handle varied data schemas and high volumes of data, these engines are not necessarily designed to understand context and semantics. In contrast, LLMs excel at interpreting natural language queries and generating structured responses. Through such methods as RAG and vector search, LLMs provide context-aware answers that go beyond keyword matching (see Figure 3 on page 12). However, this capability comes with a higher computational cost and a risk of 'hallucination,' i.e., producing context-appropriate responses that may be subtly incorrect or completely fictional. Hybrid approaches, which combine the efficiency and scalability of structured indexing with flexible context-aware LLM solutions, are emerging in the developing enterprise search space.

Figure 3: Text-to-SQL RAG process and challenges

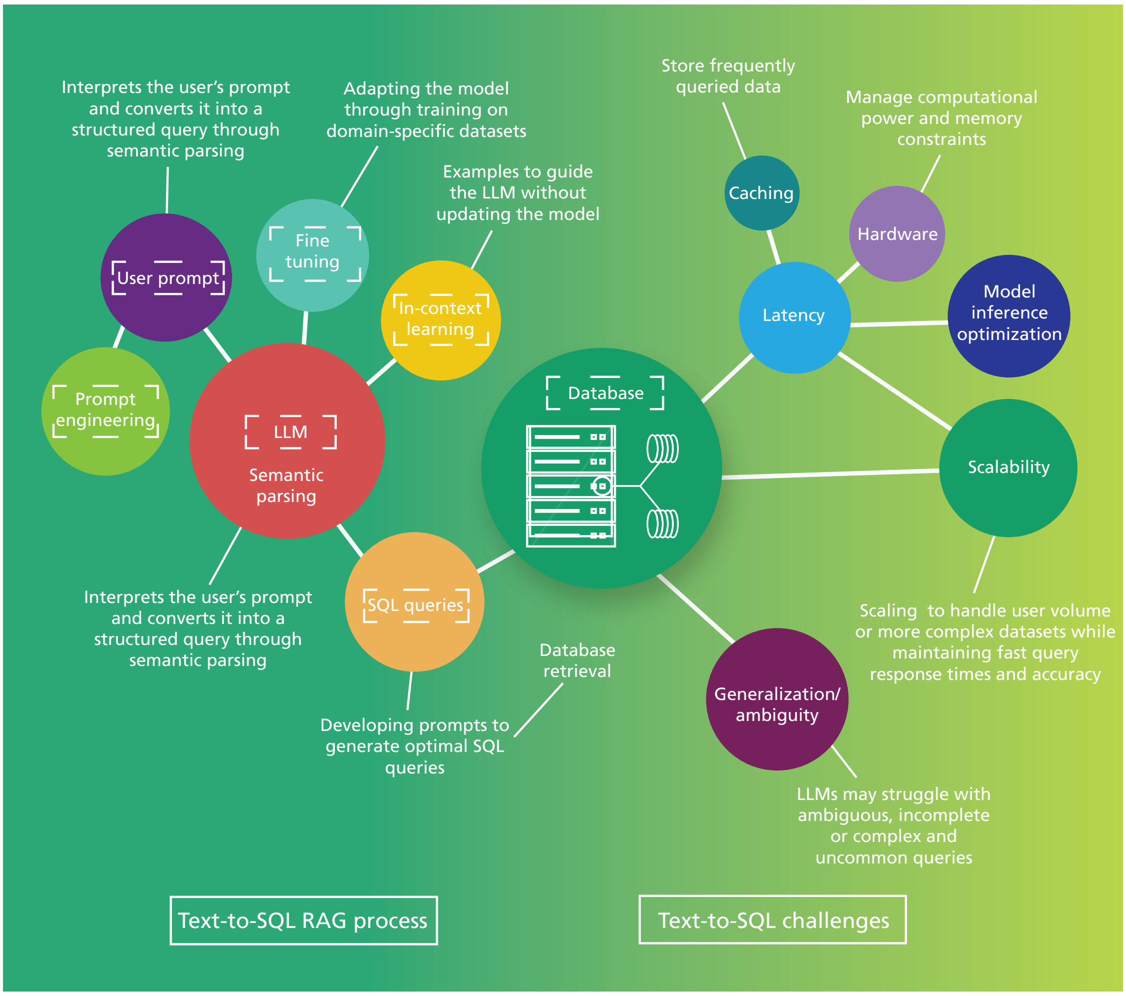

Source: Chartis Research

5. Navigating the evolving AI landscape: product insights from coding assistants and copilots

Redefining technological development

The latest wave of AI innovation is characterized by the rapid succession of watershed moments, redefining technology development across industries. Some key highlights include:

- **AI in scientific research**. DeepMind's AlphaFold program made significant advances in protein folding, marking a pivotal achievement in scientific research. AlphaFold's co-creators, Demis Hassabis and John Jumper, were awarded the 2024 Nobel Prize in Chemistry for their work in protein structure prediction.

- **Robotic labs revival**. Despite declining investment in the robotics industry since 2021, there are instances of renewed interest in research and automation.

- **AI-generated media**. The rapid proliferation of AI-generated media is driven by a combination of incumbents and start-ups. As competition among major players accelerates, creative professionals now have a host of image- and video-generation tools available.

- **AI in healthcare**. GenAI systems are being piloted across various applications, including patient care, diagnostics and outcome prediction. But despite promises of industry-wide transformation, barriers to widespread adoption remain.

Promises of transformation: profitability and sustainability?

While these technologies may promise transformation, questions relating to their profitability and sustainability loom large. The rapid pace at which frontier labs and tech companies are releasing new models and products showcases fierce competition; however, successive funding rounds are revealing the capital-intensive nature of GenAI ventures. Industries must navigate issues such as fair use, accuracy, privacy and trustworthiness, all of which impact the development and deployment of AI products. While these issues are broadly applicable, the way they manifest depends on the specific sector in which technology is deployed. LLMs are sometimes referred to as foundation models (FMs), a definition that highlights how they can be used for a variety of downstream tasks. However, relying on FMs for downstream tasks may generate a host of supply chain and compliance concerns.

GenAI strategies: new dependencies

Many companies featured in Chartis' 2024 RiskTech AI ranking and awards had well-established AI strategies that shaped their product development and use cases before the emergence of LLMs. While these products often rely on deep learning libraries, and to a lesser extent pre-trained models, the current dependence on LLM providers marks a new shift. Software as a service (SaaS) platforms benefit from being able to integrate new AI technologies on an ongoing basis. However, the reliance on Frontier Labs for LLM access introduces unique industry challenges. Frontier Labs has reported significant revenue growth from enterprise sales, which include both application programming interface (API) access and customized solutions. Pricing for these enterprise plans varies based on provider, model type, customization needs, volume and throughput consistency.

Inferencing costs, often measured by input/output tokens and provisioned throughput units (PTUs), have proven lower than initial projections—especially considering the staggering cost of training LLMs. However, the way these costs evolve will depend on the developing business models of Frontier Labs, the way it competes and AI hardware and power dynamics. Technology vendors also face such strategic decisions as whether to opt for smaller, customizable models, or to mitigate risks associated with vendor lock-in as models undergo rapid iteration.

Code assistants: tangible benefits and ROI

Widespread attention on high-profile GenAI model and product releases has spurred significant interest among technology vendors and financial institutions looking to integrate and extend their current systems with GenAI capabilities. Amid the surge of AI-powered copilots, one application has emerged as particularly impactful: code assistants and automation. While the promise of industry-wide transformation remains complex and challenging, code assistants are already delivering measurable returns.

According to Google's CEO, more than a quarter of the company's new code is written by GenAI models, before being reviewed and accepted by engineers. GitHub Copilot has been adopted by more than 77,000 organizations and **accounted for more than 40%** of GitHub's overall revenue growth in 2024. A **Microsoft research paper** has discussed controlled experiments to evaluate the performance difference achieved by using Copilot. It reported a 55.8% improvement in coding productivity, with the greatest benefits observed among less experienced developers, older programmers and those who coded more each day. Start-ups are also active in this space, with substantial ongoing funding rounds, and security and disclosure concerns are also driving some financial institutions to develop internal projects. Most existing projects in this area rely on partnerships with established tech firms.

Financial institutions are increasingly exploring the use of code assistants and AI-powered copilots (see Table 1). Code assistants can significantly enhance productivity by offering code autocompletion and snippets, reducing time-intensive tasks (such as debugging) and boosting overall coding efficiency. The growing adoption of code assistants in finance provides insight into the kinds of GenAI applications likely to deliver the highest ROI across the industry. Key factors for success include ease of integration within existing systems, automation of repetitive tasks and measurable productivity gains. Scale also appears to be a factor here, as large banks and financial firms have extensive in-house development teams that manage elements such as proprietary trading platforms, risk management systems, third-party system customization and client-facing applications. Pricing models for code assistant enterprise plans vary, but **GitHub Copilot currently charges** $39 per user per month. Challenges to adoption, such as compliance, data privacy and security considerations, provide valuable lessons for further implementation efforts.

Table 1: AI projects by financial institutions – some examples

Bank	GenAI project	Partner/provider
ANZ	GitHub Copilot (3,000 software developers and engineers). Copilot for Microsoft 365.	**Microsoft**
Bank of America	Erica (virtual financial assistant – client-facing).	
BNP Paribas	Multi-use case.	**Mistral**
Citi	GitHub Copilot.	Microsoft
Commonwealth Bank of Australia	GitHub Copilot, Copilot for Microsoft 365.	**Microsoft**
Danske Bank	GitHub Copilot, Copilot for Microsoft 365. DanskeGPT (internal smart assistant).	**Microsoft**
Emirates NBD	GitHub Copilot, Copilot for Microsoft 365, ChatGPT.	**Microsoft**
Goldman Sachs	GitHub Copilot.	**Microsoft**
Morgan Stanley	Debrief (meeting assistant).	OpenAI
Banco Itaú	GitHub Copilot.	**Microsoft**
JPMorganChase	LLM Suite (60,000 employees).	**OpenAI**

N.B. This list is not exhaustive, is focused on copilot projects and is limited to Frontier Lab and Hyperscaler partnerships. Details of RiskTech AI vendor copilot projects are not included.
Source: Chartis Research

6. AI governance and policy

The rapid pace of innovation in today's AI industry is prompting financial supervisory bodies to assess systemic and emerging risks. Notably, the European Union's AI Act and California's recently vetoed AI Safety Bill SB 1047 sparked significant debate about how to regulate AI. The EU AI Act, which came into force in August 2024 and is expected to be generally applicable by 2026, has already led some major tech firms to reconsider product releases in the EU. Key sticking points in the Act include issues around data consent, transparency, emotion recognition and other biometric data processing concerns. In the financial services sector, certain AI systems, including credit risk scoring, are classified as high-risk under the Act's tiered risk framework. The Act aims to harmonize with existing financial regulations, supervisory bodies and the Digital Operational Resilience Act (DORA).

Financial regulators are particularly focused on managing AI-related risks, including those impacting market integrity, cybersecurity, fraud prevention and the potential for bias in AI models. In addition to providing ethical guidelines, many supervisory bodies are continuing to take a risk-based approach. This emphasizes regulation around the applications and effects of AI rather than focusing on the underlying technology.

Currently, supervisory bodies are actively gathering information and shaping preliminary policies (see Figure 4). This includes such efforts as the US Treasury Department's request for proposals (RFP), public consultation documents and the AI Consortium being organized by the Bank of England (BoE). These fact-finding and policy-shaping efforts are likely to inform the next stages of regulatory and supervisory frameworks for AI in financial services.

Figure 4: Regional regulatory snapshot

Source: Chartis Research

US agency regulatory actions on AI:

- Commodity Futures Trading Commission (CFTC)
  - 2024: Issued a request for comment on the use and risks of AI in derivatives markets.
  - Technology Advisory Committee: Published a report on responsible AI in financial markets.
- Financial Stability Oversight Council (FSOC)
  - 2023: Designated AI use in financial services as a potential vulnerability, recommending ongoing monitoring and assessment.
- Securities and Exchange Commission (SEC)
  - 2023: Proposed new regulations to address conflicts of interest related to the use of predictive data analytics (including AI) by broker-dealers and investment advisors.
  - 2024: Clamped down on 'AI washing' (making false or misleading statements about AI capabilities).
- Financial Crimes Enforcement Network (FinCEN)
  - AML Act 2020: Established standards to promote the use of AI in AML and counter-terrorism financing (CFT) compliance.
  - 2023: Organized industry discussions on AI's role in fraud detection and protecting data confidentiality.
- Federal banking agencies (OCC, FRB, FDIC, CFPB, NCUA[1])
  - 2021: Issued joint guidance on AI use in financial institutions, focusing on robust risk management frameworks.

[1] *Office of the Comptroller of the Currency, Federal Reserve Board, Federal Deposit Insurance Corporation, Consumer Financial Protection Bureau, National Credit Union Administration*

7. Vendor ranking and awards: overview

For this ranking and awards report, Chartis conducted a detailed analysis of vendors across the risk management landscape, assessing their capabilities and use of AI. Through this analysis, Chartis provides valuable insights into the market leaders and emerging players, guiding stakeholders as they navigate the complex and rapidly evolving field of AI-driven risk management.

Scoring criteria

The scoring criteria for the RiskTech AI ranking are based on four key areas: impact, deployment, strategy and innovation. These criteria are designed to focus on the application of AI itself rather than a solution's breadth of coverage.

- **Impact**. This measures the originality of an approach and its effect on a particular use case. While the approach doesn't need to be a 'new' ML technique, it should be applied in a way that has provided tangible benefits to the solution/system and its users or clients.
- **Deployment**. This criterion assesses the computational infrastructure and its alignment with ML workloads, as well as the specific methodological framework used. It evaluates how effectively the solution is deployed and its readiness to handle AI tasks.
- **Strategy**. This evaluates the effectiveness of a firm's strategy in utilizing various AI models and techniques. It also considers how well a firm leverages AI in relevant use cases.
- **Innovation**. This criterion is an overall measure of an organization's creativity and ingenuity in its AI model approach. It also assesses how innovative an organization is in developing, applying or integrating AI into its solutions.

Award categories

The award categories span various risk management areas and business lines, including:

- **Retail banking**. Behavioral modeling, credit scoring, customer segmentation and customer service.
- **Capital markets**. Trade and capital optimization, pricing, optimization and risk analytics.
- **Insurance**. Underwriting, claims management and customer service.
- **Governance, risk management and compliance (GRC)**. Regulatory intelligence, document summarization/search, process mining and controls management and compliance testing.
- **Financial crime**. Cybersecurity/cyber risk management, fraud detection, anti-money laundering (AML), client screening, case analytics, sanctions and transaction monitoring.

By identifying excellence across these categories, the Chartis RiskTech AI report not only ranks vendors but also sheds light on the specific contributions and advances they bring to different facets of financial risk management. This recognition also helps stakeholders understand the capabilities of leading vendors and the potential impact of emerging players in the industry.

8. RiskTech AI 50 2024: ranking

2024 Rank	Company	HQ	Overall score	Impact	Deployment	Strategy	Innovation
1	Oracle	US	78.38	84.00	77.50	70.00	82.00
2	MathWorks	US	77.88	82.00	82.50	70.00	77.00
3	Moody's	US	74.50	71.00	73.00	80.00	74.00
4	Feedzai	Portugal	72.75	72.00	77.00	68.00	74.00
5	Wolters Kluwer	Netherlands	72.00	73.00	68.00	70.00	77.00
6	Hitachi	Japan	70.25	70.00	70.00	66.00	75.00
7	FIS	US	69.25	67.00	70.00	70.00	70.00
8	NICE Actimize	US	69.00	62.00	68.00	78.00	68.00
9	IBM	US	68.75	72.00	72.00	59.00	72.00
10	SS&C	US	68.50	65.00	65.00	73.00	71.00
11	Broadridge	US	68.25	64.00	64.00	67.00	78.00
12	MetricStream	US	68.00	62.00	62.00	77.00	71.00
13	TCS	India	66.00	60.00	66.00	68.00	70.00
14	Prometeia	Italy	65.75	60.00	62.00	76.00	65.00
15	SymphonyAI	US	64.00	53.00	60.00	78.00	65.00
16	Quantexa	UK	63.75	60.00	60.00	70.00	65.00
17	Hawk	Germany	62.75	60.00	56.00	70.00	65.00
18	Appian	US	62.55	58.00	60.20	67.00	65.00
19	Fenergo	Ireland	62.50	55.00	68.00	61.00	66.00
20	Akur8	France	62.38	55.00	60.00	66.00	68.50
21	Global Valuation	UK	62.25	61.00	62.00	57.00	69.00
22	ValidMind	US	60.75	55.00	49.00	69.00	70.00
23	Finastra	UK	60.25	48.00	60.00	68.00	65.00
24	Riskfuel	Canada	59.50	50.00	54.00	67.00	67.00
25	ThetaRay	Israel	59.25	54.00	56.00	62.00	65.00

2024 Rank	Company	HQ	Overall score	Impact	Deployment	Strategy	Innovation
26	Oxane Partners	UK	59.00	53.00	58.00	60.00	65.00
27	Ripjar	UK	58.38	50.50	51.00	64.00	68.00
28	BCT Digital	India	58.25	50.00	55.00	63.00	65.00
29	KPMG	UK	58.13	48.00	53.50	65.00	66.00
30	Eastnets	UAE	57.88	48.00	53.50	65.00	65.00
31	Dasseti	US	57.75	50.00	49.00	63.00	69.00
32	Xapien	UK	57.50	50.00	54.00	63.00	63.00
33	Complytek	Cyprus	57.38	50.00	55.50	64.00	60.00
34	CogNext	US	57.25	48.00	54.00	60.00	67.00
35	Evatech	Armenia	57.00	57.00	50.00	56.00	65.00
36	Cognitive View	Australia	56.75	60.00	55.00	51.00	61.00
37	NetGuardians	Switzerland	56.50	45.00	48.00	68.00	65.00
38	Owlin	Netherlands	56.13	38.00	49.50	72.00	65.00
39	Opensee	France	56.00	45.00	55.00	59.00	65.00
40	GFT	Germany	55.88	54.00	46.50	60.00	63.00
41	Iguazio	Israel	55.63	50.00	52.00	60.50	60.00
42	MoCaX Intelligence	UK	55.50	55.00	50.00	56.00	61.00
43	Solytics Partners	US	55.00	44.00	45.00	66.00	65.00
44	Facctum	India	54.75	39.00	52.00	63.00	65.00
45	WorkFusion	US	54.00	45.00	50.00	58.00	63.00
46	CleverChain	UK	53.80	35.00	47.20	64.00	69.00
47	Decision Focus	Denmark	53.75	47.00	53.00	56.00	59.00
48	Lucinity	Iceland	53.50	41.00	44.00	64.00	65.00
49	smartKYC	UK	51.88	43.00	44.50	58.00	62.00
50	CALPANA	Austria	50.25	48.00	48.00	50.00	55.00

9. RiskTech AI 50 2024: category winners

Category award	2024 winner
Chartis awards	
Deployment	MathWorks
Impact	Oracle
Innovation	Oracle
Strategy	Moody's
Solution category awards	
Advanced tax fraud capabilities	Quantexa
AI-driven anti-fraud platform	Feedzai
AI-driven buy-side documentation management	Oxane Partners
AI-driven credit analytics	Moody's
AI-driven credit data enrichment	Moody's
AI-driven customer onboarding	Fenergo
AI-driven cyber risk management	IBM
AI-driven data management framework	Oracle
AI-driven document management framework	Oracle
AI-driven insurance risk analytics	Moody's
AI-driven insurance underwriting	Akur8
AI-driven IT risk analytics	TCS
AI-driven legal workflow	Zeidler
AI-driven pricing and valuation	Riskfuel
AI-driven process control and process mining	Appian
AI-driven regulatory intelligence	Wolters Kluwer
AI for audit risk management	MetricStream
AI for government applications	Quantexa
AI for workflow and automation	WorkFusion
Code generation and control	MathWorks

Category award	2024 winner
Solution category awards	
Computational infrastructure for AI	Oracle
Cross-vertical GRC analytics	TCS
Financial crime investigation support and copilot	Xapien
Innovative use of AI for anomaly detection	Hawk
Machine learning development environment	MathWorks
ModelOps	Iguazio
Model validation tools for AI	ValidMind
Operational risk dashboard and workflow	MetricStream
Unified research AI copilot	Moody's
Use of AI for climate risk analytics	Moody's
Use of AI for commercial banking-related workflow	FIS
Use of AI for compliance policy control and automation	CleverChain
Use of AI for customer service	NICE Actimize
Use of AI for energy analytics systems	Hitachi
Use of AI for energy portfolio management	Hitachi
Use of AI for finance and accounting	Oracle
Use of AI for institutional brokerage	Broadridge
Use of AI for markets-related workflow	SS&C
Use of AI for private banking and wealth management	NetGuardians
Use of AI for risk data	Opensee
Use of AI in AML and transaction monitoring	NICE Actimize
Use of AI in behavioral modeling	Prometeia
Use of AI in capital markets	Broadridge
Use of AI in communication and control	NICE Actimize
Use of AI in cross-industry applications	MathWorks
Use of AI in retail banking	Oracle

10. RiskTech AI 50 2024: ones to watch

Ones to watch	2024 winner
	• AI Risk • Alertspeed • Cleareye.ai • Fineksus • Flagright • Focal • Kobalt Labs • Napier AI • Refine Intelligence • Stratyfy • Think360 • Zeidler

11. Further reading

STORM 2024

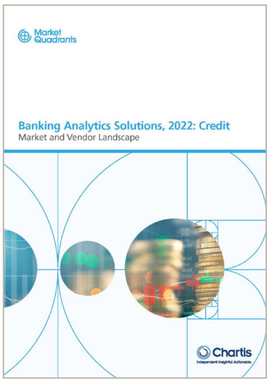

Banking Analytics Solutions, 2022: Credit; Market and Vendor Landscape

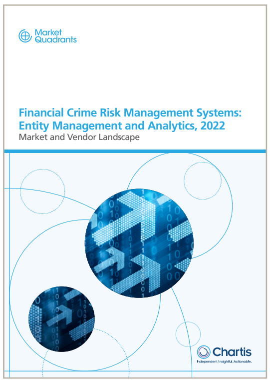

Financial Crime Risk Management Systems: Entity Management and Analytics, 2022; Market and Vendor Landscape

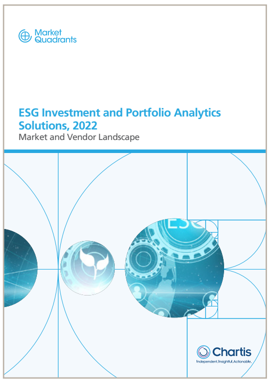

ESG Investment and Portfolio Analytics Solutions, 2022: Market and Vendor Landscape

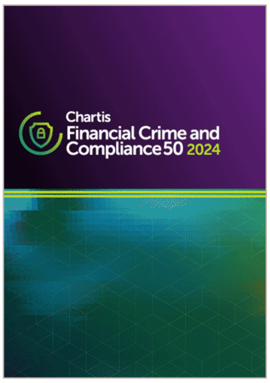

FCC50 2024

RiskTech100 2025

For all these reports, see **www.chartis-research.com**