FROM HINDSIGHT TO INSIGHT TO FORESIGHT

FAQ on the Use of AI for Financial Crime Compliance

protiviti®

Global Business Consulting

INTRODUCTION

Ask financial crime professionals what the most challenging part of their job is, and most will likely say it is the timely identification of suspicious activity. As much as companies have worked to improve their detection capabilities given their compliance obligations, their desire to protect their reputations and their understanding, as corporate citizens, of the impact of financial crime on society — success has remained elusive.

Historically, financial crime detection has involved mostly after-the-fact (hindsight) identification of potentially illicit activity gleaned from reviewing massive amounts of alerts — most of which are non-productive. Enter artificial intelligence (AI), which offers the opportunity not only for better detection results (insight) but for predicting (foresight) when suspicious activity may occur. Add the potential process efficiencies that AI offers across many facets of financial crime compliance programs, and the large number of companies at risk of being used to facilitate financial crime — and it becomes clear why enterprises are increasingly eager to understand and explore the opportunities.

For purposes of this publication, we have chosen to focus on financial services and e-commerce, two of the higher at-risk industry sectors; however, much of what is covered will apply to other industries as well given the indiscriminate nature of financial crime and wide-spanning regulatory/compliance mandates. The questions are illustrative and far from exhaustive, but they are the ones that have arisen most often during our discussions and work with clients and others in the marketplace. Some questions apply generally to the adoption of AI for any purpose, while others are very specific to the use of AI to support financial crime compliance. We intend our responses to be plain language, largely non-technical answers to the questions that may be on the mind of financial crime professionals, management, and board members when it comes to whether or how to adopt artificial intelligence.

This booklet is provided for general information only and is not intended as legal analysis or advice. Companies should seek legal counsel on specific questions as they relate to their unique circumstances. Regulatory guidance and industry standards on the use of artificial intelligence continues to evolve and varies across jurisdictions and industries. Accordingly, some of the issues addressed in this booklet may be impacted by future guidance.

November 2024

Carol Beaumier
Senior Managing Director, Risk & Compliance

Constantine Boyadjiev
Managing Director, Risk & Compliance Analytics

01 What is artificial intelligence?

Artificial intelligence (AI) refers to any computer system or machine capable of mimicking human intelligence. In other words, it is the ability of a computer system to emulate human-like cognitive abilities such as learning and problem-solving.

Although AI has been receiving significant attention recently because of relatively new technologies like large language models (LLMs), AI as a scientific domain is not inherently new. The foundational concepts of AI began appearing in scientific literature in the 1940s and the actual term "Artificial Intelligence" was coined in 1955 by Stanford Professor Emeritus John McCarthy. Elements of AI such as machine learning (referring to the ability of algorithms and statistical models to learn and adapt without being explicitly programmed) have been deployed for decades across a variety of use cases and industries.

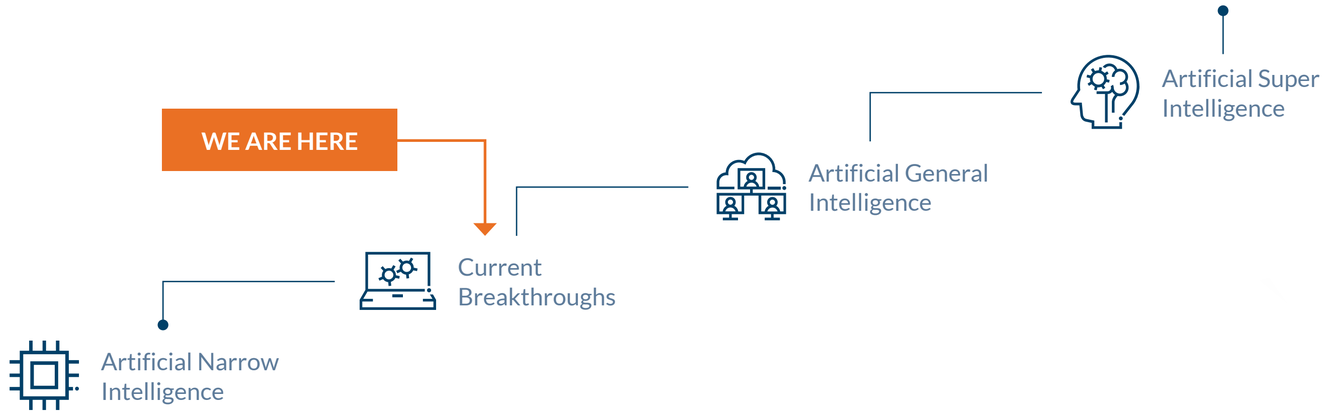

Source: IDB Invest.

There are many forms of AI, often grouped by capabilities and functionalities.[1]

AI Capabilities	AI Functionalities
Narrow AI	Reactive Machine AI
General AI	Limited Memory AI
Super AI	Theory of Mind AI
	Self-Aware AI

Today, all AI is Narrow; General and Super AI remain theoretical. Refer to the Appendix for definitions of these capabilities and examples of some of the respective functionalities.

[1] "Understanding the different types of artificial intelligence," (IBM, 2023): www.ibm.com/think/topics/artificial-intelligence-types.

02 Is the use of AI new to financial services or e-commerce?

No. Both sectors have used AI for some time. For example, AI in the financial services sector traces its origins to the 1980s, when it was primarily employed to identify fraud.[2] Other examples of early adoption by the financial services industry include back-office automation, credit scoring/risk underwriting models, portfolio management, structured derivatives pricing and customer service chatbots. Advances in AI have continued to introduce additional functionality and complexity. E-commerce businesses have also used AI for decades to, among other things, analyse customer data and make personalised product recommendations, respond to routine customer inquiries, and predict customer demand and drive dynamic pricing decisions.

03 What is the difference/relationship among AI, generative AI and large language models?

Generative AI, or Gen AI, is a subset of AI that focuses on "generating" new content such as images, audio, video, text, code or even 3D models that are original and not just a variation of existing data.

Despite its increased functionality, Gen AI is considered Narrow AI because it operates under far more limitations than even the most basic human intelligence.

Large language models (LLMs) are a type of generative AI trained on vast amounts of data with a large number of parameters that generate novel text-based responses. Today there are a number of proprietary LLMs built/developed by third parties with a conversational interface (e.g., ChatGPT developed by OpenAI), accelerating user interactivity and ease of adoption.

[2] K. W. Kindle, R. S. Cann, M. R. Craig, and T. J. Martin, "PFPS — Personal Financial Planning System — AAAI," (Proceedings of the Eleventh National Conference on Artificial Intelligence,1989), 344-349.

04 What is financial crime?

Financial crime broadly refers to all crimes that involve taking money or other property that belongs to someone else to obtain a financial or professional gain. The specific activities included as financial crime are called predicate crimes or offences and are generally determined by jurisdictional law.

The extent to which a company is exposed to any of these financial crimes is a function of many variables including the nature of its products and services, its customer base, its geographic footprint, and its control environment.

Predicate crimes or offences include but are not limited to:

- Bribery and Corruption
- Cyber Crime
- Drug Trafficking
- Environmental Crime
- Human Smuggling
- Human Trafficking
- Illegal Arms Trafficking
- Market Abuse
- Organised Crime and Racketeering
- Proliferation Financing
- Tax Evasion
- Terrorist Financing
- Trafficking in Arts and Antiquities
- Violations of Sanction and Export Control Requirements

05 Are there types of financial crime that are unique to the e-commerce industry?

E-commerce money laundering, or transaction laundering, is the process of leveraging e-commerce and merchant processing to create fictitious transactions that appear legitimate. These transactions may involve knowing or unknowing participants in the e-commerce ecosystem, a network of interconnected parties involved in the buying and selling of goods and services.

Transactions may be facilitated by front companies that appear to sell legitimate goods and services but are set up by money launderers to provide cover for their illegitimate activities, pass-thru companies set up by third parties and used by one or more criminals, or funnel accounts in which payment processors may commingle legitimate and illegitimate transactions. They may involve the sale of fake or contraband goods, the value of e-commerce transactions may be over-inflated, or the transactions may simply be nonexistent, a scheme sometimes referred to as ghost laundering companies that offer payment platforms, further attracting criminals as a means to "circumvent" more traditional financial services payment channels.

Think of transaction laundering as an updated version of trade-based money laundering. Transaction laundering is difficult to detect for a number of reasons including the complexity of the payments network, growth in alternative payment methods, the inability of merchants to safeguard their websites from being used illegally, and the use of hidden websites.[3]

Data from Juniper Research indicates that losses resulting from e-commerce fraud will exceed $107 billion by 2029.

Source: "eCommerce Fraud to Exceed $107 Billion in 2029," Juniper Research, Oct. 7, 2024 press release: www.juniperresearch.com/press/pressreleasesecommerce-fraud-to-exceed-107bn-in-2029/.

[3] "Clean Money is a Click Away: The Money Laundering Risks of E-Commerce" Protiviti, 2021: www.protiviti.com/us-en/whitepaper/clean-money-click-away-money-laundering-risks-e-commerce.

06 What makes AI interesting to companies with financial crime compliance obligations?

Organisations with financial crime compliance obligations have made and continue to make massive investments in technology and talent to support their compliance efforts. Empirical evidence however suggests that the level of investment is not, in many cases, supported by results/success metrics. For example, the global spend by banks in 2022 to combat money laundering is estimated at $274.1 billion, yet a study by the United Nations Office of Drug and Crimes suggests "much less than one percent (probably around 0.2 percent)" of the proceeds of crime laundered via the global financial system are seized and frozen.[4]

While publicly-available information on the financial crime compliance spend by e-commerce companies and its correlation to the ability to identify financial crime is more difficult to obtain given the relative newness of the industry, there is no doubt that the rapid growth of e-commerce has given rise to a surge in transaction laundering. While the exact scale of transaction laundering is difficult to quantify, in 2017 one industry observer suggested that the global volume of transaction laundering exceeded $350 billion. That number has likely grown significantly by now.

One obvious reason for companies to adopt AI in their financial crime compliance programs is to exploit its analytical (insight generation) and predictive capabilities (foresight of emerging risks) to target potentially illicit activity more accurately and conversely, eliminate the "noise" in companies' transaction monitoring systems while generating "signals."

Another reason companies are interested in adopting AI to combat financial crime is the desire and need to keep up with the criminals, whose methods are becoming increasingly sophisticated and crafty. Criminals have become early adopters of new technologies, including AI, and exploit them fully to serve their nefarious purposes. Organisations therefore are viewing the adoption of AI tools and capabilities as necessary to level (at least somewhat) the playing field that has become an "arms race" of sorts.

A third motivating reason for regulated companies to use AI in their financial crime compliance programs is that regulators themselves are increasingly deploying AI. Institutions that rely solely on traditional methods for managing their financial crime programs may therefore not detect issues that their regulators identify. The level of adoption and the maturity of an institution's AI program may prove to be a key differentiating factor, making some institutions more or less susceptible to being used as an "easy target" for facilitating financial crime.

*Global financial crime compliance costs for financial institutions total more than $206 billion or $3.33 per month for each working-age person in the world, according to LexusNexis Risk Solutions.**

[4] Elisabeth Krecke, "Why anti-money laundering policies are failing" (Geopolitical Intelligence Services AG, GIS Reports, 2024): www.gisreportsonline.com/r/why-anti-money-laundering-policies-are-failing/.

* www.prnewswire.com/apac/news-releases/lexisnexis-risk-solutions-study-reveals-global-financial-crime-compliance-costs-for-financial-institutions-totals-more-than-us206-billion-301937916.html

07 How can AI improve the efficiency and effectiveness of financial crime compliance?

The potential benefits to the use of AI can be grouped into two broad categories: 1) advanced analytical capabilities and 2) elimination of routine tasks. Examples include:

Advanced Analytical Capabilities:

- AI systems can analyse in high velocity massive amounts of transactional data as well as various types of data (e.g., structured/semi-structured/unstructured) to identify unusual patterns or behaviours that may indicate potential financial crimes and over time can minimise false positives.
- AI can support the development and maintenance of dynamic financial crime risk assessments, which afford the opportunity to monitor continuously for changes in a risk profile.
- AI can forecast areas of potential/emerging risks before such materialise.
- AI algorithms can detect anomalies and outliers in datasets, i.e. not only uncover the known-knowns and known-unknowns, but the unknown-unknowns.
- AI can prioritise areas for review based on an analysis of the attendant risks.

Elimination of Routine Tasks:

- AI can identify and summarise regulatory requirements to aid in the training and upskilling of financial crime personnel.
- AI tools can be used to automate customer due diligence and negative news screening processes.
- AI can generate requests for information (RFIs) where an analyst or investigator needs more information to understand customer activity.
- AI can generate narratives evidencing the review of suspicious activity.
- AI can generate regulatory filings/reports.

08 AI may present significant opportunities to financial services and e-commerce companies, but how will it impact customers?

The use of AI by financial services and e-commerce companies offers the possibility of reducing customer friction. Financial services companies, including e-commerce companies registered as money services businesses, spend considerable time updating customer profiles and investigating potentially suspicious activity, both of which often involve direct customer outreach. With the improved capabilities offered by AI to develop and maintain dynamic customer profiles and better target potentially suspicious activity, there should be a reduction in company outreach, which is often viewed as intrusive. This should empower positive customer experiences while protecting the enterprise against financial crime threats vectors and vulnerabilities.

09 What are additional applications of AI when it comes to customer identity assurance and detecting account take over?

AI has a number of applications when it comes to detection of fraud stemming from customer identity compromise. An ever-increasing threat vector faced by both FSI and e-commerce enterprises is account take over (ATO). ATO results from a customer's data/account being compromised, which can originate from a variety of sources — social engineering, phishing, vishing, smishing, data leaks (both within and outside the enterprise), etc. It negatively affects customers and enterprises alike, and such effects have both tangible/financial ramifications, as well as intangible/reputational harm dimensions. Additionally, deploying traditional/static controls (e.g., multi-factor authentication) as defence mechanisms for managing the customer's identity can yield only partial benefits, while also bringing customer friction and negatively impacting client experience. Enter AI, which can aid customer authentication while simultaneously helping identify suspicious activity related to ATO.

Through the use of AI and machine learning, analytical principles such as Social Network Analysis (SNA) and Link Analysis can be readily deployed in identifying nefarious behaviour/activity. SNA is the process of investigating social structures through use of networks and graph theory. Networked structures are characterised in terms of nodes (e.g., individual actors, people, or things within the network), and the associative ties/linkages that connect them (e.g., relationships or interactions). Similarly, link analysis is a data-analysis technique used to evaluate connections between nodes. Relationships may be identified among various types of objects, including organisations, people and transactions. AI empowers the use of such approaches given its ability to ingest a massive amount of structured/unstructured data and metadata (e.g., transactional, behavioural, "digital fingerprint"/device data, geolocation, etc.) and in real-time/near-real time develop anomaly detection and behavioural algorithms — which in turn helps organisations in their identity assurance efforts and to uncover "unknown-unknowns," all while protecting their customers and reputation.

10 What are some of the risks associated with the use of AI?

The use of AI carries some of the same inherent risks that financial institutions already face, such as protecting the privacy and security of the data used and relying on incomplete or inaccurate data to form judgements, although it can be argued that the use/misuse of AI exacerbates these risks.

AI also poses the following risks, among others:

- **Ethical considerations/fairness/bias:** If the data used to train an AI system contains embedded biases, while unintended, the AI may replicate or even amplify these biases in its outputs, leading to inaccurate, unfair or discriminatory decisions.
- **Lack of transparency, interpretability and explainability:** Some AI models, particularly those based on deep learning, function as "black boxes" that provide little insight into how they make decisions; this raises questions about whether the developers and users of these models actually know what the models are doing, and also complicates efforts to validate and modify the models given such opacity.
- **Evolving regulatory frameworks:** Because AI policy is being developed at the local, national and regional levels and there is no global AI "policy," inclusive of legal frameworks, ethical standards and principles, the risk exists that a company's decision to deploy AI, while made in good faith based on available guidance, may subsequently be determined to fall outside of acceptable parameters/norms/guidance.
- **Regulator acceptance:** In part related to lack of transparency, regulators may have concerns about the use of AI by financial institutions, putting a firm on the defensive to prove that AI produces better results than the methodologies and more "traditional" tools previously deployed.
- **Trustworthiness:** A famous mathematician (George Box) once stated: "*All models are wrong, some are useful.*"[5] While Gen AI models are generally trained on vast data sets, they may present incorrect or misleading results (such as "hallucinations") as fact for a number of reasons, including bad training data and bad assumptions. In addition, given their inherent design, these models can produce different results with each iteration and are not subject to traditional methods of model validation (such as replication, as per above stated rationale).
- **Over-reliability on technology and lack of "human-in-the-loop" factor:** There may be a tendency to defer to the technology and overlook the continued importance of human oversight and supervision.
- **Impact on the workforce:** AI raises concerns about job loss from an employee perspective, and concerns about obtaining new competencies/upskilling the existing workforce from an employer perspective.

On a more philosophical note, and not specific to the financial services industry, AI raises an existential risk — that artificial general intelligence (AGI) surpasses human intelligence.[6]

Apart from existential risk which cannot be managed by an individual company, the lynchpin of an institution's AI risk management program should be its own AI Use Policy. (See Question 12 below.)

[5] James Clear, "All Models Are Wrong, Some Are Useful" (JamesClear.com, 2016): www.jamesclear.com/all-models-are-wrong.
[6] Bernard Marr, "The 15 Biggest Risks Of Artificial Intelligence" (Forbes.com, 2023): www.forbes.com/sites/bernardmarr/2023/06/02/the-15-biggest-risks-of-artificial-intelligence/.

11 What are some of the questions and considerations companies should ask before deploying AI?

Among the key questions institutions should ask when considering the deployment of AI are the following:

- **Do we have a specific problem/use case we are trying to solve, and can AI solve it?** Without a clearly articulated goal and objectives, the risk is high that AI deployment will be unsuccessful.
- **Should we buy, build or "borrow" an AI solution?** Determining whether an institution should "buy or build" should include due diligence on commercially available options, and consideration of in-house expertise/resourcing, extent of customisation needed, cost, data security and privacy, scalability, opacity/transparency of the solution, and time to market. These are fundamentally the same factors that apply to any technology "buy or build" decisioning, although in the case of AI "borrow" becomes a relevant approach by starting with foundational models and AI capabilities available through hyper-scalers and tuning them to a specific use case.
- **Do we have the skillsets, competencies, talent/resource know-how and cadre experience to manage the AI implementation internally, or do we need to engage outside assistance?** In addition to basic problem-solving and change management skills, AI implementation, depending on its nature, may require a number of specialised skills including programming languages, data modelling, data warehousing and data processing, understanding of machine learning, advanced analytics, data science, and knowledge of intelligent user interfaces (IUIs).
- **Do we have the data we need to train AI models?** Although some AI tools may use publicly-sourced information, many need a sufficient amount of reliable and relevant internal data to learn. Without adequate data, the proverbial "garbage in, garbage out" still rings true. While many enterprises are "data rich," the data still needs to be easily "consumed" (e.g., centralised, complete, in good hygiene, etc.) to be usable.
- **Have we considered the challenges and costs of accessing the data we need?** Data ingress/egress is expensive across multiple clouds; ideally AI models are co-located architecturally with the needed data.
- **Does the potential benefit of the AI justify the cost?** Alongside of the costs of AI implementation itself, it is important to identify the key performance indicators (KPIs) that will be used to measure results and assess the effectiveness of your AI initiatives. (More on this below).
- **Does the planned implementation align with the organisation's principles and guidance on AI governance and usage?** Many organisations have implemented board-approved, enterprise-wide AI governance standards that delineate approved uses of AI; establish the information required to make an informed decision about an AI tool, including identification of all attendant risk; and prescribe monitoring requirements. Implementing AI on an ad hoc basis, absent company governance standards,

may expose the implementation decision to second-guessing internally from senior management and the organisation's board of directors and by external parties, including regulators.

- **Do we understand and are we prepared to manage the regulatory expectations for the use of AI?** Understanding regulatory expectations, especially for financial institutions and e-commerce companies that operate in multiple jurisdictions where expectations may differ, is critically important for regulated institutions.
- **Do we understand all of the downstream effects of the use of AI?** The implementation of AI can be transformative and may require changes to policies, procedures, data management programs, other technologies and internal training programs, among other potential impacts.
- **How will our organisation manage change during this transformation?** Adopting AI may necessitate significant changes across processes, roles and cultures, which should be managed proactively through effective communication and training programs.
- **Have we identified all key risks that may arise from our use of AI (see Question 10) and developed appropriate risk mitigation plans?** Being able to manage AI risks effectively requires a solid understanding of specifically how the risks manifest in AI. That means that risk management, compliance and internal audit personnel responsible for designing and testing the AI control framework must have a solid understanding of AI, and should have a "seat at the table" as these AI initiatives are launched/rolled out.
- **Will the planned use of AI have a direct impact on customer engagement?** If yes, does there need to be some advance communication with customers to help them understand and accept these changes?
- **Do we have a plan for ensuring the continued reliability of the AI model?** As with any other models, institutions need to ensure on an ongoing basis that an AI model is operating as intended and remains conceptually sound. This requires frequent testing/validation, performance monitoring, and outcome analysis to assess the accuracy of the AI model's output and whether it is operating per its prescribed/intended use; data drift monitoring to identify whether the nature of the data that the AI model interacts with is faulty, thereby potentially requiring adjustments to the model; and bias and fairness checks.

12 What should be included in a company's AI Use Policy?

An AI Use Policy should document the company's responsible use of AI and should include the following content:

- **Purpose and scope —** goals of the use of AI, aligning with strategy, any limitations on where in the institution AI may be used
- **Roles and responsibilities —** governance and oversight of the development, acquisition and use of AI
- **AI development —** standards for the in-house of AI
- **Due diligence of third-party providers —** standards for performing initial and ongoing due diligence on third-party providers
- **Authorised AI tools —** permissible and prohibited AI tools
- **Regulatory requirements —** relevant laws, regulations and guidance applicable to the use of AI, including ethical considerations
- **Monitoring —** requirements for evaluating the ongoing integrity, reliability and suitability of AI tools
- **Training and awareness —** necessary training and upskilling for employees about responsible use of AI technology
- **Exceptions to policy —** the institution's exception management policy and procedures

13 There are many different types of financial crimes. Does each type use different AI tools?

AI capabilities such as chatbots, natural language processing (NLP), audio signal processing (ASP), computer vision, Gen AI, and machine learning algorithms and more have potential benefits across different types of financial crime detection.

14 What are examples of AI use cases for AML/CFT?

By analysing vast amounts of data and identifying complex patterns, AI can significantly improve the accuracy of detecting illicit activity, resulting in fewer false positives (legitimate transactions flagged as potentially suspicious) and false negatives (suspicious transactions that are not identified).

By automating the process of capturing, documenting, and organising the alert/case narrative in a standardised and traceable format using AI, the alert review team of one bank was able to increase its productivity 5X.

Value: Efficiency, cost effectiveness

15 What is an example of an AI use case for sanctions and export controls?

A financial institution uses AI to risk score its sanctions alerts, dispositioning those that pose little risk, and directing higher risk alerts to humans to resolve.

Value: Efficiency, regulatory compliance

16 What are examples of AI use cases for fraud?

Fingerprint scanners, facial recognition and voice recognition technologies can be used to offer an extra layer of security, making it more difficult for fraudsters to impersonate legitimate customers.

A payment processor uses time, location, device and GPS data to determine whether activity occurring in distant geographies may be fraudulent. The company believes that AI will eventually learn to evaluate certain behaviours, including swiping speed and gestures when assessing the likelihood of fraud.

Value: Effectiveness, reputational harm minimisation, customer protection, revenue leakage aversion

17 What is an example of AI use cases for market manipulation?

A broker-dealer uses AI to analyse large datasets from multiple sources, such as market data, transactional data, social media and news feeds, to identify deviant trading patterns or anomalies in real time. This enables firms to undertake real-time monitoring, detect and deter potential violations, and send out timely alerts for investigation across a series of use cases (e.g., rogue trading, insider information, market abuse, collusion, sales malpractice, elder abuse, etc.)

Value: Effectiveness, regulatory compliance, reputational harm minimisation

18 What is an example of AI use cases for anti-bribery and corruption?

An institution uses AI which learned from historical data to analyse large data sets, flag transactions and establish links between entities that deviate from established patterns and may indicate improper payments.

Value: Effectiveness, prudent risk detection, compliance

19 In addition to some of the use cases cited above, how else can AI be used to detect transaction laundering in e-commerce?

Computer algorithms can be used to examine merchant sites electronically and can spot indications of front companies that the human eye might not be able to detect.

Value: Effectiveness, risk and loss mitigation

20 How would a company measure the impact of AI on its compliance effort?

Measuring ROI for AI investments can be complex as many benefits are long-term and difficult to quantify precisely. Among the metrics institutions may consider are the following:

a. Improved efficiency as evidenced by better productivity and/or reduction/reallocation of staff

b. Reduction of false positives/improved detection rates (i.e., more signal/less noise)

c. Better regulatory outcomes including better examination results, fewer violations of law and penalties

d. Reduced customer friction such as faster client onboarding and quicker resolution of questions about customer transaction activity, less need to contact customers

e. Greater agility to manage new threats

21 What are the expectations and requirements for the use of AI?

Governments, regulators and standard-setting bodies are all developing guidelines and frameworks for the use of AI.

As governments and regulators across the globe consider the transformative impact that AI will have, they are developing governing frameworks and communicating their expectations for the ethical and responsible use of AI. The EU AI Act is one significant example of a government framework. Other jurisdictions and regulators are still in an information-gathering phase and have not published final guidance, although most have at least signaled through speeches and in industry for what they are thinking.

Examples of emerging standards for AI governance issued by standard-setting bodies include:

- **NIST AI Risk Management Framework** is designed to equip organisations and individuals with approaches that increase the trustworthiness of AI systems, and to help foster the responsible design, development, deployment and use of AI systems over time.
- **ISO/IEC AI Framework** provides guidance on managing risks associated with the development and use of AI. The document offers strategic guidance to organisations to assist in integrating risk management into significant activities and functions.

The policy paper published in August 2023 by the UK's Office of Artificial Intelligence and Department for Science, Innovation and Technology does a great job of succinctly offering five guiding principles for the responsible development and use of AI, which are common to most of the published and emerging guidance:

- Safety, security, resilience and robustness
- Appropriate transparency, interpretability and explainability
- Ethics and fairness
- Accountability and governance
- Contestability and redress

As governments and regulators across the globe consider the transformative impact that AI will have, they are developing governing frameworks and communicating their expectations for the ethical and responsible use of AI.

22

What are some of the lessons learned by companies that have adopted AI tools?

Some companies that were early adopters of newer AI tools learned the hard way of the importance of making sure they address all questions provided in the response to Question 11.

Some early adopters overestimated the functionality and benefits of an AI tool; in some cases, this included misjudging time savings and the extent to which staff could be reduced.

A valuable lesson learned by early adopters was the benefit of starting small and scaling. This allowed them to prove their value proposition, gain necessary experience without being overwhelmed by the process, and test the technology before scaling to larger initiatives and ultimately industrialising new-found capabilities.

23

What impact will the use of AI likely have on the staffing of financial crime compliance departments?

Depending on the nature and extent of adoption, the use of AI may allow for staff reduction, principally among the staff who perform routine tasks. This would leave compliance professionals more time to focus on what's really important — the activities that require human judgement and experience. The use of AI will also prompt the need to add (or upskill) staff in more specialised roles, including individuals who understand how to use AI tools and effectively evaluate AI outputs, and who can evaluate the ongoing performance.

24

What is the future potential for the use of AI to fight financial crime?

The use of AI to fight financial crime can be a game changer — achieving cost-savings while driving efficiency and improving efficacy that the industry has been unable to achieve to date. Given the continued evolution of AI capabilities, potential use cases are limited only by our imagination and institutions that don't leverage AI will find themselves at a disadvantage.

25

In the race to achieve foresight, who will win — financial services or e-commerce?

In the battle to fight financial crime, we believe the collective efforts and lessons learned from all interested parties — both public and private sector — will drive the most progress. Both the financial services industry with its extensive experience fighting financial crime in a highly regulated environment and the e-commerce industry, which includes many tech-savvy digital natives, have much to contribute to the common goal of stopping the bad guys.

*In one survey of 356 experts, half believe human-level AI will exist by 2061, and 90% said it will exist in the next 100 years. But, for now at least, it is important to remember that AI is a tool, not a replacement for humans. It allows humans to focus on what's really important — the things that require human experience, judgement and creativity.**

* AI timelines: What do experts in artificial intelligence expect for the future? — Our World in Data

HOW PROTIVITI CAN HELP

Whether your organisation is just getting started with AI technologies or is far along on its journey to explore advanced use cases, Protiviti can provide support and guidance to help lead your organisation to successful outcomes along the entire lifecycle of AI adoption. We can assist with:

- **Considering AI:** Protiviti can assist with the identification of potential areas where AI can bring value to operations, products or services, and outline a deployment plan. Considerations include Business Value Definition, Data Accessibility, Operational Readiness, Strategic Alignment and Governance.
- **Implementing AI:** Protiviti can help the successful development and deployment of AI solutions that address specific business challenges or opportunities. Considerations include proving initial hypotheses and technical challenges, defining clear objectives for AI projects, selecting appropriate AI techniques and allocating resources, ensuring close collaboration between all relevant stakeholders, and designing and effectuating strong governance.
- **Monitoring AI:** Protiviti can aid the measurement of the effectiveness and impact of AI solutions on operations and/or goals. Considerations include establishing relevant performance metrics, revising risk management taxonomies and processes to cover AI holistically, conducting thorough testing and validation of the AI models, and creating feedback loop and continuous monitoring.
- **Securing AI:** Protiviti can help organisations protect AI systems and data from potential threats and ensure their ethical and responsible use. Considerations include implementing robust cybersecurity measures to safeguard AI models, addressing AI bias and fairness concerns, and adhering to ethical guidelines and regulatory requirements (e.g., transparency, privacy).

In summary, Protiviti can help empower your organisation's AI journey.

About the authors

CAROL BEAUMIER

Carol Beaumier is a senior managing director in Protiviti's Risk and Compliance practice and leader of the firm's APAC Financial Services practice. Based in Metro D.C., she has more than 30 years of experience in a wide range of regulatory issues across multiple industries. Before joining Protiviti, Beaumier was a partner in Arthur Andersen's Regulatory Risk Services practice and a managing director and founding partner of The Secura Group, where she headed the Risk Management practice. Before consulting, Beaumier spent 11 years with the U.S. Office of the Comptroller of the Currency (OCC), where she was an examiner with a focus on multinational and international banks. She also served as executive assistant to the comptroller, as a member of the OCC's senior management team and as liaison for the comptroller inside and outside of the agency. Beaumier is a frequent author and speaker on regulatory and other risk issues.

CONSTANTINE BOYADJIEV

Constantine Boyadjiev leads Protiviti's Global Regulatory and Compliance Analytics data science practice. He is responsible for architecting and delivering Protiviti's Risk, Fraud and Compliance Analytics offerings across geographies. Boyadjiev brings extensive experience across industries and has held executive roles in Financial Services and Advisory ventures, having built robust risk, fraud management and analytic enterprise capabilities.

Appendices

Glossary of Terms

The following non-exhaustive glossary defines terms[7] that are used throughout this booklet:

AI Algorithms — a subset of machine learning that provides instructions for machines to analyse data, perform tasks and make decisions. The three major categories of AI algorithms are supervised learning, unsupervised learning and reinforcement learning.

AI Bias — also known as machine learning bias or algorithm bias, this refers to AI systems that produce results that reflect and perpetuate human biases within a society, including historical and current social inequality.

AI Ethics — a multidisciplinary field that studies how to optimise AI's beneficial impact while reducing risks and adverse outcomes.

Artificial Intelligence — the theory and development of computer systems able to perform tasks that normally require human intelligence, such as visual perception, speech recognition, decision-making and translation between languages.

Chatbot — a computer program designed to simulate conversation with human users, especially over the internet.

Data Mining — the process of analysing large sets of data to find patterns or trends.

Deep Learning — a type of machine learning that uses multilayered neural networks, called deep neural networks, to simulate the decision-making power of humans.

Fairness — the importance of using accurate data and preventing discriminatory effects.

Generative AI — deep-learning models that can generate high-quality text, images and other content based on the data on which they were trained.

Hallucinations — incorrect, misleading or nonsensical results presented as facts by large language models (LLMs).

Human Level AI — a state of AI in which unaided machines are able to accomplish tasks better and more cost efficiently than humans.

Hyper-scaler AI — refers to the immense capabilities of the world's largest technology companies, such as Microsoft, Google and Amazon.

Large Language Models (LLMs) — type of artificial intelligence model that has been trained through deep learning algorithms to recognise, generate, translate and/or summarise vast quantities of written human language and textual data.

Limited Memory AI — a type of AI that can use past data or experiences to inform future decisions but cannot retain to use over a long period of time. Examples include virtual/digital assistants (e.g., Siri, Cortana and Alexa) and self-driving cars.

Machine Learning (ML) — a type of AI and computer science that focuses on using data and algorithms to enable AI to imitate the way that humans learn, gradually improving its accuracy.

Narrow (or Weak) AI — a type of AI that can be trained to perform a single or narrow task, often far faster and better than a human mind can. This is the only type of AI that exists today.

Natural Language Processing (NLP) — a machine learning technology that gives computers the ability to interpret, manipulate and comprehend human language. Spellcheck is a familiar type of NLP.

Neural Network — a method of artificial intelligence that teaches computers to process data in a way that is inspired by the human brain.

[7] A variety of sources were used for these definitions, including IBM, The European Central Bank Knowledge Centre on Interpretation, Stanford University Human Centered Intelligence, the HubSpot Blog, Science Direct, Alark Kingdom, AWS and various dictionaries.

Phishing — a type of cyberattack that seeks to trick individuals into supplying sensitive information such as usernames/passwords or credit card numbers.

Reactive Machine AI — an AI system with no memory that is designed to perform a very specific task by analysing vast amounts of data and using statistical mathematics. Since it cannot recollect previous outcomes or decisions, it only works with presently available data. One example would be a system that analyses a customer's buying patterns and recommends products that may be of interest.

Reinforcement Learning — a type of machine learning process that focuses on decision making by an autonomous agent, i.e., any system that can make decisions and act in response to its environment independent of direct instruction by a human user.

Responsible AI — a set of principles that help guide the design, development, deployment and use of AI aimed at building trust in AI solutions and align the output with the values, legal standards and ethical principles of society at large.

Robotic Process Automation (RPA) — an intelligent automation technology that can perform repetitive office tasks of human workers, such as extracting data, filling in forms and moving files.

Self-Aware AI — a theoretical type of AI that would possess super AI capabilities. If ever achieved, it would have the ability to understand its own internal conditions and traits along with human emotions and thoughts; it would also have its own set of emotions, needs and beliefs.

Smishing — a social engineering attack that uses fake mobile text messages to trick people into downloading malware, sharing sensitive information or sending money to cybercriminals.

Social Engineering — the use of deception or manipulation to trick people into compromising their personal security or the security of an enterprise network.

Social Network Analysis — a research method that examines the structure of relationships among people, organisations or other entities.

Strong AI (or Artificial General Intelligence) — a theoretical type of AI that can use previous learnings and skills to accomplish new tasks in a different context without the need for human beings to intervene/train the underlying models. This would allow AGI to learn and perform any intellectual task that a human being can.

Super AI — a theoretical type of AI that would think, reason, learn, make judgements and possess cognitive abilities surpassing those of human beings. Applications using Super AI capabilities will have evolved beyond the point of understanding human sentiments and experiences to feel emotions, have needs and possess beliefs and desires of their own.

Supervised Learning — a type of machine learning that uses labelled datasets (i.e., raw data that has been assigned one or more labels to add context or meaning) to train algorithms to predict outcomes and recognise patterns.

Theory of Mind AI — a theoretical type of Strong AI, Theory of Mind functionality would understand the thoughts and emotions of other entities and be able to personalise its interactions based on an individual's unique emotional needs and intentions.

Transparency and Explainability — respectively, the communication of how, when and for which purposes an AI system is being used, and the extent to which it is possible for relevant parties to access, interpret and understand the decision-making processes of an AI system.

Unsupervised Learning — often referred to as "deep learning," a type of machine learning that analyses and clusters unlabelled data sets to discover hidden patterns or data groupings without the need for human intervention.

Vishing — the *fraudulent* practice of making phone calls or leaving voice messages *purporting* to be from *reputable* companies in order to induce individuals to reveal personal information.

Relevant Resources from Protiviti

Enabling Enterprise AI Adoption Through Next-Generation Governance

Success with Generative AI Requires Balancing Risk With Reward

The Director's Playbook for Generative AI

Understanding the Impact of the EU AI Act: A Primer for Financial Institutions

Establishing a scalable AI governance framework (co-authored with OneTrust)

AI Investments Require the CFO's Expertise — and Vice Versa

Podcast: The Rise of Generative AI — with Christine Livingston

For additional insights, visit Protiviti's Artificial Intelligence thought leadership collection and Artificial Intelligence Hub on our website.

Face the Future with Confidence®

protiviti®
Global Business Consulting

© 2024 Protiviti Inc. An Equal Opportunity Employer M/F/Disability/Veterans. PRO-1124