




Guide to Explainable AI in Financial Services

Balancing innovation with accountability
in the fight against financial crime

Contents

1. Introduction to AI explainability in financial crime prevention Page 3

Overview of AI in financial services: The role of AI in transforming financial crime prevention.

Importance of explainability: Defining AI explainability and why it is crucial, especially in highly regulated industries like finance.

SymphonyAI's commitment: SymphonyAI's focus on explainability within its SaaS products and AI overlays.

2. The role of explainability in enhancing trust and compliance Page 5

Building trust with explainability: How transparent AI models help build trust with stakeholders.

Regulatory compliance: The importance of explainability for meeting regulatory requirements.

3. Explainability in Anti-Money Laundering (AML) Page 8

Challenges in AML: Common difficulties faced in AML, such as false positives.

Explainability in the AI Overlay for Transaction Monitoring: How SymphonyAI's transaction monitoring overlay provides clear, understandable decisions to users.

4. Explainability in Sanctions Screening Page 10

Challenges in Screening: The many difficulties that mean screening effectively can be challenging.

Explainability in the AI Overlay for Screening: How SymphonyAI's screening overlay helps institutions make informed decisions with clear AI outputs and transparency.

5. Case Study – Absa Bank Page 12

How a global tier 1 bank increased their hit rate and identified new risks by incorporating AI into their case management.

1 Introduction to AI explainability in financial crime prevention

Overview of AI in financial services

Traditional methods of monitoring and compliance often rely on rule-based systems that struggle to keep pace with the complexity and volume of modern financial transactions. AI helps solve this issue, and its use is transforming financial crime prevention.

How? By enhancing the efficiency and effectiveness of detecting and mitigating financial crimes. Put simply, AI is providing institutions with the technology to secure the financial system more effectively.

With its advanced analytical capabilities, AI can process vast amounts of data in real-time, identifying patterns and anomalies that might indicate illicit activities. Using machine learning algorithms, the technology adapts and improves its detection engine as new threats emerge, providing organizations with a dynamic, proactive defense against financial crimes like money laundering, fraud, and sanctions evasion.

Not only content to spot crimes that might previously have gone undetected, AI also enhances the accuracy of financial crime prevention efforts, significantly reducing false positives that plague traditional, legacy systems. Able to learn from historical data and user feedback, AI refines its criteria for flagging suspicious activities, improving over time. For example, in [Anti-Money Laundering \(AML\)](#), AI can identify complex layering techniques used by criminals to obscure funds, which might otherwise have been missed. This process has a knock-on positive effect: It streamlines the workload for compliance teams, allowing them to focus their attention on genuine threats, which also enhances the customer experience by minimizing unnecessary interventions.

And it isn't only detection and accuracy where AI shines. It also fosters greater collaboration and intelligence sharing. Advanced AI platforms should integrate seamlessly with existing systems, facilitating the exchange of insights and data, [breaking down silos](#), and enabling a more unified approach to combating financial crime. Through sophisticated case management tools including generative AI copilots, investigations can be far quicker, leading to improved resolution times.

In short, AI not only bolsters the security and compliance of financial institutions but also contributes to the stability and integrity of the broader financial ecosystem.

This all sounds familiar, doesn't it? But what about the most important bit that doesn't often get explained or answered? Away from all the marketing copy, there's a crucial element to AI systems that doesn't often get talked about. [And that's the...](#)

Importance of explainability

AI explainability is crucial for ensuring transparency and accountability in decision-making. This understanding is essential for gaining the trust of stakeholders (regulators, compliance officers, etc.) and customers, who need to be confident that AI-driven decisions are fair, unbiased, and logical.

Although these are all good reasons, explainability is not just a nice to have. Within financial services (and financial crime prevention in particular), it is vital for meeting legal and regulatory requirements. Rational explanations are needed for all automated decisions (as required by the likes of GDPR and AML directives), helping institutions avoid significant fines, loss of trust, and reputational damage for non-compliance.

Leading with explainability also allows banks and other financial institutions to dissect and audit their own AI operations, enabling compliance teams to continuously improve AI model effectiveness. These tweaks can see better interpretation of outputs, which directly improves model accuracy, reduces false positives, and leads to faster decision making. Such behavior also demonstrates continued adherence to regulators, showcasing an ongoing commitment to transparency, accountability, and reducing unnecessary customer friction.

In essence, everybody gains from explainable AI. It builds trust, ensures compliance, and optimizes the performance of AI-powered financial crime prevention tools, leading to a much more secure financial system.

SymphonyAI's commitment

SymphonyAI is deeply committed to integrating AI explainability into its financial crime prevention software, recognizing it as a cornerstone of responsible AI practices.

In fact, every AI model and application is run through a rigorous checklist to ensure compliance with the company's [responsible AI principles](#):

Accountability

Transparency

Reliability and safety

Security

Privacy

By prioritizing interpretability in its products, SymphonyAI ensures that users can understand, trust, and use the insights generated by its entire software suite.

This commitment to provide clear, understandable rationales for AI-driven decisions aligns with industry best practices and extends to its collaborations with regulatory bodies and clients.

By actively engaging with stakeholders, SymphonyAI is better able to understand the specific needs and concerns within the financial sector. This collaborative approach allows for continuous innovation via refinements and enhancements, always aligning with evolving regulatory standards and expectations.

Through these efforts, SymphonyAI is not only improving the usability and accountability of its solutions but also creating a culture of trust within financial crime prevention, positioning itself as a leading advocate for innovative and responsible AI within financial services.

2 The role of explainability in enhancing trust and compliance

Building trust with explainability

Building trust in AI systems within financial crime detection and prevention hinges significantly on not only transparency but also the clarity provided by explainable AI models, demystifying complex decision-making processes into an understandable and interpretable system. This, in turn, builds significant levels of trust across all levels.

There are also other benefits for those in specific roles, explained below.

For regulators: Explainability not only aids in fulfilling necessary requirements but also demonstrates an institution's commitment to responsible AI use and accountability, a commitment to be operating within legal guidelines, and ensuring fair and unbiased treatment of customers.

For customers: It is important to be able to see that AI models and decision-making is consistent, justified and free from discriminatory biases. The transparency of communication, proactive approach to mitigating concerns, and reduction in false positives also improve customer satisfaction.

For investigation teams: Employees want to be able to believe in the software that they are relying upon. Being able to delve into the decision-making, read exactly why something was flagged and act upon it, helps with productivity and an assurance in the system. Alongside this, such explainability allows for investigators to truly work as one with the AI, ensuring a seamless approach to working as confidence soars.

For all stakeholders: Knowing how and why an AI system arrived at a particular decision increases their confidence in the technology. This feeling of reliability is crucial in an industry where implications of errors can be profound.

All these outcomes are made possible via explainability, which in turn enhances trust and strengthens long-term relationships.

Regulatory compliance

Regulatory compliance is a crucial aspect of financial services, necessitating that institutions adhere to various standards set by governments and regulators around the world. Explainability plays a pivotal role in ensuring compliance with these regulations, especially considering how much they differ based on jurisdiction, providing considerable complexity for global businesses.

Here are some key regulatory requirements where explainability is essential:

General Data Protection Regulation (GDPR):

- [Article 22 of GDPR](#), an EU initiative, emphasizes the 'right to explanation'; Individuals have the right to understand and obtain meaningful information about the logic involved in automated decisions affecting them.
- Institutions must ensure that AI systems can provide clear explanations of how personal data is processed and the rationale behind such decisions. E.g., those made in credit scoring or [fraud prevention](#).

AML regulations:

- Various AML regulations such as the [AML Directives](#) in the EU or [Bank Secrecy Act](#) in the US, require financial institutions to implement effective measures to detect, prevent, and report money laundering activities.
- Explainable AI systems are crucial to demonstrate how suspicious transactions are identified and flagged, ensuring that institutions can provide justifications to regulators for their AML processes.

Financial Action Task Force (FATF) recommendations:

- [FATF recommendations](#) provide a comprehensive framework for countries to implement risk-based approaches for combating money laundering and terrorist financing.
- AI systems must be transparent to allow financial institutions to clearly understand and explain the logic behind risk assessments and compliance measures within their respective countries.

By prioritizing explainability into AI systems, organizations operating within financial services not only comply with these regulatory requirements but also enhance their own operational integrity, enhancing trust and supporting a strong, responsible approach to financial crime prevention.

AI laws and regulations

It is not only industry-specific guidance and laws that companies must follow. With the impressive speed improvements within AI in recent years, many countries and jurisdictions have moved to enact AI-specific regulations.

These include the EU AI Act as well as similar legislation in countries including China, and Colombia. Most other countries, including the likes of the UK, US, Australia, Japan, Canada, and Brazil have drafts or frameworks currently in place. In fact, there are at least [69 countries in the world](#) that have proposed over 1000 AI-related policy initiatives and legal frameworks.

This underscores the increasing importance of explainability in AI systems.

These regulations are designed to ensure that AI is developed and deployed in a safe, responsible, and transparent manner, avoiding any potential risks associated with AI.

Here are a few to help you understand the regulations currently in place (or progressing) across the globe:

EU AI Act:

- The [EU AI Act](#) became law in 2024. It classifies AI systems based on their risk levels and imposes different requirements for each category. High-risk AI systems, which include those used in financial services for fraud prevention and compliance, are subject to stringent transparency and documentation needs.
- Explainability is a core component of the EU AI Act's compliance framework. Companies deploying high-risk AI systems must ensure that these systems are transparent and that their decision-making processes can be understood by users and regulators alike, fostering accountability and trust.

China's Generative AI Measures:

- China was among the first countries to regulate generative AI with [generative AI measures](#) to address privacy concerns, data security, and content screening passing in 2023. Since then, there have been various other laws and regulations passed on the topic of AI.
- Providers of AI services must follow national labeling standards on the data that they are working with and outline exactly where AI has been used (if customer facing). Alongside this, AI algorithms that are used must be filed with the Chinese government.

Colombia's National Artificial Intelligence Policy:

- In February 2025, [Colombia passed its AI laws and regulations](#), which is aimed at developing ethical and sustainable development of AI systems within the country.
- Explainability is key to the policy with all AI systems having to be transparent in design and operation. AI must be free from bias and discrimination while clear documentation and communication with stakeholders is expected.

Although they may differ slightly, the AI laws and regulations already in place, as well as those working their way through governments around the world all emphasize the need for clear, understandable AI.

By proactively incorporating explainability into AI software today, financial institutions can better prepare for compliance with these regulations tomorrow, safeguarding against any future potential legal and/or reputational risks while contributing to a more responsible and transparent AI ecosystem.

3 Explainability in AML

Challenges in AML

AML efforts are critical to stemming the tide of criminal profiteering. However, implementing effective AML measures presents many challenges due to the complex and developing nature of financial crimes:

High volume of transactions:

Financial institutions process an immense number of transactions daily, making it difficult to effectively monitor and flag suspicious activities without advanced solutions

False positives:

Traditional AML systems often generate a high number of false positives, overwhelming compliance teams with alerts that require manual investigation, reducing the focus on genuine threats

Evolving money laundering techniques:

Criminals are constantly developing new methods to disguise illegal funds, such as using cryptocurrencies or complex layering strategies, making it challenging for institutions to keep up

Regulatory complexity:

Compliance with diverse and evolving global, regional, and local regulations requires institutions to maintain rigorous and adaptable AML programs, which can strain resources and expertise

Data privacy concerns:

Balancing the need for thorough surveillance and data collection with obligations to protect customer privacy and comply with data protection regulations adds complexity

Cross-border transactions:

The globalization of finance means that money laundering often involves cross-border transactions, complicating efforts to trace and comprehend the flow of funds across different jurisdictions and regulatory environments

Resource constraints:

Institutions often face limited resources, both in terms of personnel and technology, to manage the growing demands of AML compliance effectively

Integration with legacy systems:

Many financial institutions operate on legacy IT systems, which can make integrating new, advanced AML technologies difficult and may require significant investments in upgrading infrastructure

Addressing these challenges requires a combination of advanced technologies, like AI and machine learning, strong regulatory frameworks, and continuous training and development of compliance personnel to adapt to new threats and technologies.

The [AI overlay](#) for Transaction Monitoring aims to confront these challenges head on to improve financial crime prevention across the industry but it's important that decisions are explainable.


Explainability in the AI Overlay for Transaction Monitoring

The AI Overlay for Transaction Monitoring is an AI overlay solution designed to enhance the performance of existing rules-based transaction monitoring systems. Powered by advanced predictive AI, SensaAI for AML is trained on a customer's historical alert adjudication data to assess transactions and identify those resembling known suspicious activity. By doing so, a far more accurate assessment of risk can be made, enabling three significant benefits:

- 1 False positive reduction:** The AI Overlay for Transaction Monitoring has demonstrated proven success in real-world deployments, helping global financial institutions reduce false positives by 70–80%. It does this by evaluating every alert generated by the rules-based transaction monitoring system and assigning an AI Score that indicates the likelihood of genuine risk. Alerts with low scores can then be automatically suppressed or deprioritized for investigation.
- 2 High risk prioritisation:** Using the same AI Score, alerts with high-risk scores can be prioritized and automatically placed at the top of the Level 1 investigation queue or escalated directly to Level 2 for immediate review. This ensures that the highest-risk alerts are always addressed first and resolved in a timely manner.
- 3 New risk detection:** In addition to reviewing all alerts flagged by the rules-based transaction monitoring system, the overlay goes one step further. Because rules cannot see what they do not know, the AI overlay reviews every transaction that took place throughout the day to find any hidden risk missed in the initial screening.

To give users complete confidence in the AI Overlay for Transaction Monitoring, every AI Score is accompanied by a natural language explanation. Generated by a Large Language Model (LLM), this explanation provides investigators with clear reasoning behind each decision, detailing why an alert was classified as a potential true or false positive and highlighting the specific features and behaviors that influenced the AI Score.

The AI Score and its explanation are seamlessly integrated into the case management interface, displayed alongside the original alert, to make insights intuitive and easily accessible for investigators. By emphasizing transparency and interpretability, SymphonyAI empowers investigators to make faster, more confident risk assessments, ensuring genuine threats are quickly identified rather than lost among a sea of false positives.



4 Explainability in Sanctions Screening

Challenges in Screening

Screening processes play an essential role in AML and [combating the financing of terrorism \(CFT\)](#) as they help identify and mitigate risks associated with names, sanctions, and transactions. However, these processes present a unique set of challenges that financial institutions must navigate:

Inaccurate or incomplete data:

- Inaccurate, incomplete, or outdated information can lead to ineffective screening, resulting in missed matches or false alerts.

Name variations and transliterations:

- Individuals and entities may have multiple name variations or may be represented differently in various languages and scripts. This makes it challenging to detect matches accurately across diverse datasets.

High volume of alerts:

- Screening against extensive sanction lists produces a lot of alerts, many of which are false positives. This overwhelms compliance teams, diverting resources away from investigating true matches or suspicious activities.

Updated sanctions lists:

- Sanction lists are continually updated. Financial institutions must comply with these developments, which is resource intensive.

Jurisdictional differences in sanctions:

- Global institutions must comply with multiple, sometimes conflicting, regulatory requirements.

Complex ownership structures:

- Understanding ultimate beneficial ownership can be difficult, especially in cases involving shell companies. This complexity makes it challenging to determine whether an entity is indirectly linked to a sanctioned party.

Transaction structuring:

- Suspicious transactions may be structured in ways that evade conventional screening so screening algorithms should be updated and constant monitoring required.

Data privacy and security:

- Safeguarding customer information while conducting effective screening poses challenges, especially when handling sensitive data across borders. Compliance with data protection laws must be balanced with screening requirements.

Integration of screening tools:

- Incorporating screening tools into existing, legacy systems can be difficult without significant investment into technology upgrades.

Addressing these challenges required a multi-faceted approach that includes investing in advanced technology such as AI for improved accuracy, maintaining up-to-date compliance knowledge, and enhancing data management practices.

Explainability in the AI Overlay for Screening

[The AI Overlay for Screening](#) significantly improves screening match accuracy and allows financial institutions to upgrade their existing sanctions screening solution with AI.

The software uses generative AI to extract and classify relevant entities from unstructured free text found within payment messages. This provides context that otherwise wouldn't be available, meaning that situations like mixing up names with places or names with vessels can be avoided, which are common sources of false positives. These extracted entities are then passed to a predictive AI model, which can do an exhaustive check looking at all the information available to be used (name differences, date of birth, addresses, etc.)

As with any AI solution, explainability is very important. Using SymphonyAI, a prediction is returned alongside a probability of a match, with a human readable explanation on why the AI model came to its conclusion. Proof of concepts using the AI overlay have seen an 80% reduction in false positives while retaining 100% of true positives.

By increasing the quality of alerts going to investigators we can reduce investigator fatigue and help prevent mistakes in the investigation process. Because the AI model is not a rules-based system designed and understood by humans, the level of scrutiny and governance must be high.

AI investigation summary - offering explainability around the conclusions

The screenshot displays the SymphonyAI interface. At the top, the logo 'SymphonyAI' is visible. Below it, the header shows 'John Rolands Inc.' and a status 'Unexplained Difference USD 290,000/year'. The main content area is divided into three panels. The left panel, titled 'Agent Work Summary', lists four steps: 'Analyzing transaction context', 'Gathering relevant data', 'Evaluating Regulatory and Policy', and 'Generating Report'. Each step has a 'Step Detail' link and an 'Ask Agent' button. The middle panel, titled 'AI Investigation Summary', contains an 'Assessment of \$45,000 payment from London imports to Volga Trading Inc.' and a list of findings: 'Overview: A \$45,000 transaction to a Russian import-export firm raises potential compliance concerns.', 'Background: Volga Trading Inc., based in Moscow, is co-owned by Sergei Ivanov, who is the brother-in-law of Dmitry Sidorov—a sanctioned individual.', 'Compliance Concern: The family link to Sidorov suggests possible indirect ties to a sanctioned party, requiring closer review.', and 'Recommendation: Flag this transaction for further investigation.' The right panel, titled 'Transactions Intelligence', asks 'What data sources did you use for this assessment?' and lists 'Internal transaction records', 'Sanctions lists from OFAC, EU, and UN', and 'Relationship and ownership databases'. It also includes an 'Export the full audit trail for compliance review' button and a 'Ask about another issue...' link.

5 Case Study

Absa Bank

Read the full case study

The global tier 1 bank incorporated AI into their financial crime management, identifying new risks and increasing their hit rate

The challenge

[Absa](#) is one of the largest banks in Africa, working with a global client base. They partner with SymphonyAI for anti-financial crime compliance solutions, including transaction monitoring, KYC/CDD, and sanctions screening.

The finance industry and its regulations are becoming increasingly complex, so a dynamic, modern, and forward-looking approach is required. Alongside this, there is an urgent need to combat the threat of rising financial crime, with criminals using AI tools to test and threaten bank processes.

Keen to continue leading the industry with an innovative approach to anti-financial crime, the bank wished to test the impact of AI on their teams' productivity and risk reduction, specifically using the technology to test against their current transaction monitoring rules-based solution. If successful, it would provide enhanced alert quality and risk detection, valuable use cases across the bank, and cement Absa as an industry pioneer.

The SymphonyAI Solution

With AI exploding in popularity, productivity and efficiency claims of software are everywhere. Absa was interested in using the technology but wanted to see the results with their own dataset. The SymphonyAI response was a series of proof of concepts (POCs) designed to prove the likely positive impact of using AI technology across the company, from efficiency and effectiveness of alerts through to potential cost savings.

SymphonyAI worked with Absa's masked data, developing new AI models and software features. Proposed advancements included improving and speeding up new risk detection beyond the bank's current achievements alongside improving alert quality – reducing false positive alerts while continuing to identify all known suspicious activity.

Alongside this, SymphonyAI combined five different language learning models that crossed the full spectrum of manual to supervised analysis, while also developing new features to further enhance the project.

The result

The outcome far exceeded expectations of the proof-of-concept project, persuading Absa to further collaborate with SymphonyAI to fully implement the solution and roll out to other jurisdictions across the continent.

Highlight results included:

- SymphonyAI reduced false positive alerts by 77% while also capturing all suspicious activity found using their current transaction monitoring system.
- Many potential new risks were discovered, which were filtered down to the 200 highest scoring risks based on transaction amount, average score, and maximum score.
- The bank confirmed 21 new risks were identified quickly versus a slower manual adjustment of rules or manual reporting methods. Alongside this, the new risk identification hit rate of 10.5% was significantly more effective than using rules alone.
- The project was so successful that the partnership won an International Compliance Association award for breakthrough collaboration in AI-driven risk management.

Enjoy truly explainable AI with SymphonyAI

To fully appreciate the transformative capabilities of AI, and the AI Overlays for Transaction Monitoring and Screening in particular, we encourage you to dive deeper by exploring our [product pages](#), watching informative [videos](#), and engaging with our interactive click-through demos. By doing so, you'll gain valuable insights into how these tools can revolutionize your financial crime prevention efforts.

Our resources offer a comprehensive look into how SymphonyAI leverages advanced AI technologies to enhance detection, streamline compliance processes, and minimize false positives, ultimately helping you stay ahead in the dynamic landscape of financial regulations.

Whether you are a compliance professional or an operations manager, these materials will provide you with a clearer understanding of how SymphonyAI can be seamlessly integrated into your existing systems to enhance efficiency, reliability, and transparency in your AML and sanctions screening efforts.

[Get in touch](#) to find out more.

About SymphonyAI

SymphonyAI, 2024 Microsoft Partner of the Year for Business Transformation – AI Innovation, is building the leading enterprise AI SaaS company for digital transformation across the most critical and resilient growth verticals, including retail, consumer packaged goods, finance, manufacturing, media, and IT/enterprise service management. SymphonyAI verticals have many leading enterprises as clients. Since its founding in 2017, SymphonyAI has grown rapidly to 3,000 talented leaders, data scientists, and other professionals. SymphonyAI is an SAI Group company, backed by a \$1 billion commitment from Dr. Romesh Wadhvani, a successful entrepreneur and philanthropist.

SymphonyAI provides AI-focused SaaS solutions developed from the ground up for fighting financial crime. An end-to-end offering compiled from more than 25 years of knowledge in tackling money laundering and fraud, our award-winning product ecosystem comprises an innovative and modern approach to financial crime investigation using world-leading predictive and gen AI. From KYC/CDD and transaction monitoring through to payments fraud, entity resolution, and sanctions, PEP, and adverse media screening, SymphonyAI is a trusted name in finance technology with more than a third of the world's top 100 banks enjoying the power, efficiency, effectiveness and productivity that our solutions provide.



www.symphonyai.com