INDUSTRY INSIGHTS

State of AI in Surveillance Report 2025

How are surveillance teams integrating artificial intelligence into their workflows in 2025?

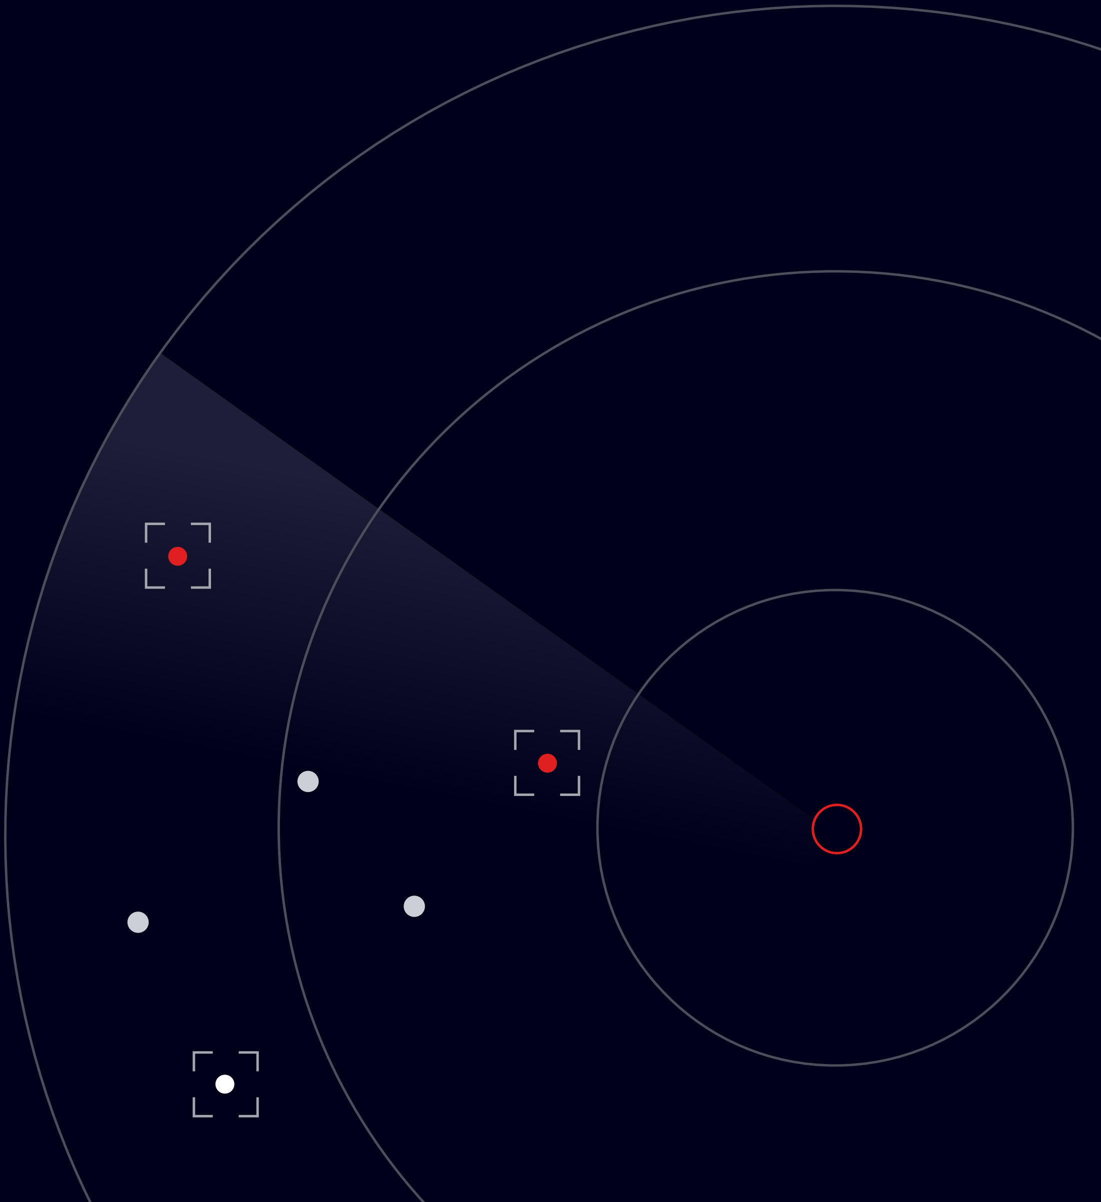

Integrity | Reputation | Control

globalrelay.com

Key findings

Global Relay issued a survey to surveillance professionals across the globe to find out their attitudes and approaches to artificial intelligence (AI) within surveillance workflows.

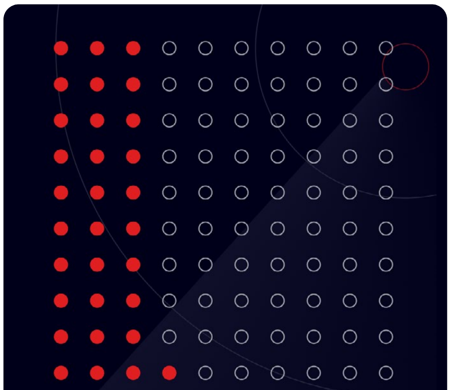

31%

of respondents are using AI for surveillance or intend to implement it in the next 12 months

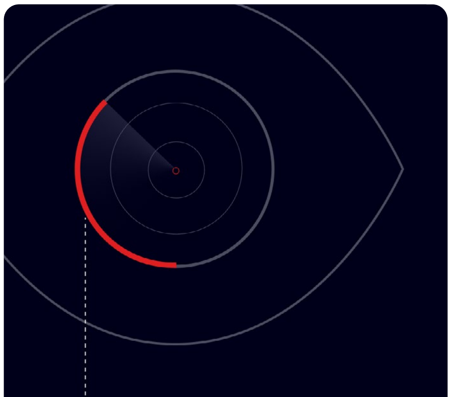

38%

of respondents are watching how the industry develops to determine whether to use AI in surveillance processes

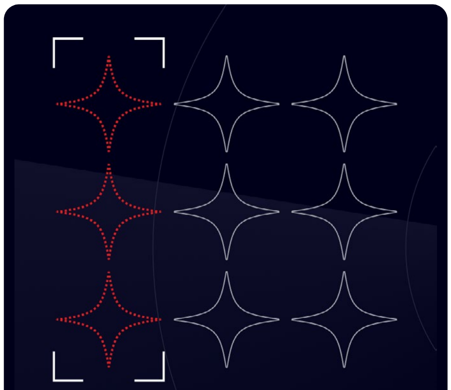

31%

of respondents do not intend to implement AI

Public understanding of AI in surveillance

1 No understanding
3.6
10 Full understanding

On a scale of 1 to 10, the average number of respondents' understanding of how AI works within surveillance is 3.6.

Top 3 use cases for AI implementation

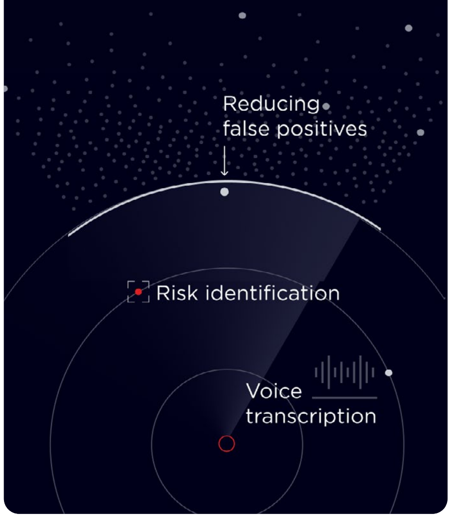

Reducing false positives

Risk identification

Voice transcription

Data security is the main barrier to AI adoption...

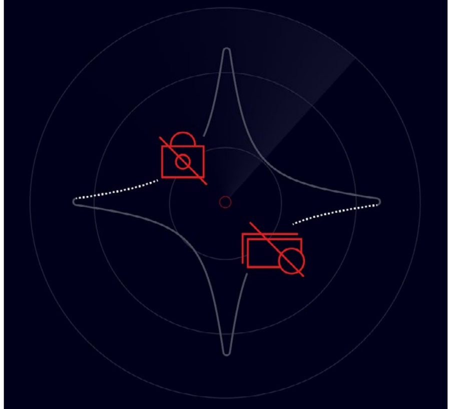

...with budget a close second

Surveillance as we know it: Outdated systems and regulatory expectation

Regulatory expectation increases as legacy systems become outdated

Surveillance teams are increasingly in the spotlight as regulatory expectation demands that financial services implement effective tools to manage an ever-growing risk landscape.

Regulators across the globe have implored organizations to monitor business communications, citing repeatedly that they are failing to "reasonably supervise…personnel with a view to preventing and detecting…violations" within operations.

The Department of Justice's (DOJ) Evaluation of Corporate Compliance Programs (ECCP) asks U.S.-based organizations to assess "the extent to which [they] have access to data and information to identify personal misconduct." It wants to see that companies are "proactively identifying either misconduct or issues with its compliance program at the earliest stage possible."

The U.S. Securities and Exchange Commission's (SEC) 2024 Enforcement Results revealed that failures to supervise personnel have led to fines of over $2 billion against financial organizations. Meanwhile, the Financial Industry Regulatory Authority's (FINRA) Annual Oversight Report for 2025 showed that it has uncovered "surveillance deficiencies" across its monitored population.

In the U.K., the Financial Conduct Authority (FCA) has said that "conduct and culture" are priorities for 2025, and expects to see that firms have "suitable controls in place to detect misconduct and to take appropriate action against those found to be committing misconduct."

Capturing the right risks: The move away from lexicon-based surveillance

Communications surveillance is currently facing a period of transformation. Traditional, lexicon-based approaches have shown varying levels of effectiveness, and the rapid innovation of AI is changing the way that firms approach communication monitoring.

Lexicon-based approaches are becoming outdated for myriad reasons. Lexicons require surveillance teams to manually define all the keywords their system needs to watch for. This means it is difficult to capture every term or phrase variation that personnel may use relating to high-risk areas – let alone identify misconduct attempts that don't fit into prescribed lexicon buckets.

A comprehensive lexicon-based model can capture a vast range of risk-related phrases, but also produce a high volume of false positives – flags that incorrectly identify risk – which means more cases for reviewers to examine. Similarly, if wrongdoers attempt to utilize evasion techniques to dupe lexicons, it could lead to the oversight of true risk areas, known as false negatives.

Outside of creating lexicon lists to target risk areas specific to a firm's appetite, compliance teams are tasked with revising and updating keywords regularly as risks emerge. This fine-tuning helps in sharpening a surveillance system's aim, though it's expected that false positives will inevitably materialize.

Goals in implementing AI surveillance

> **False positives in surveillance review queues can lead to increased costs, as compliance teams must review every flagged message, often requiring additional personnel. AI-enabled surveillance reduces false positives and accurately pinpoints undetected real risks, which enhances the efficiency of the review process.**

Nazy Alborz
Group Product Manager – AI
Global Relay

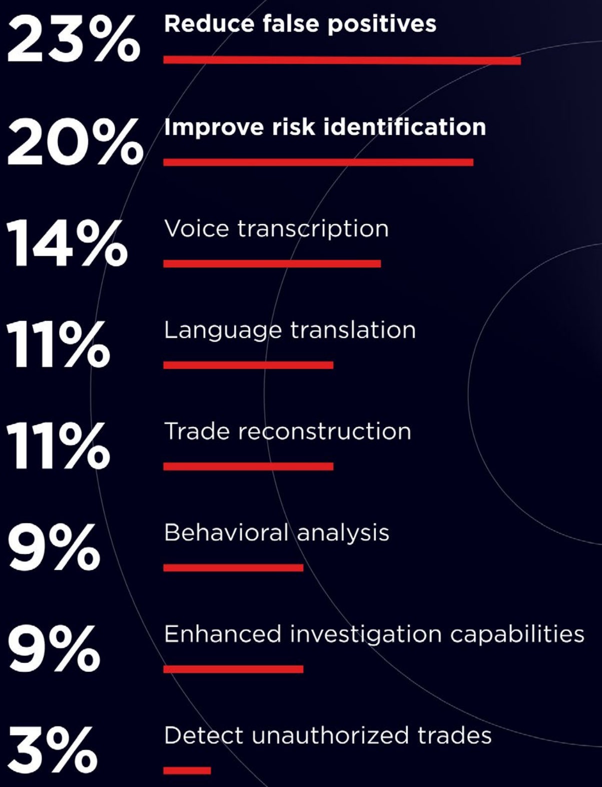

The top two goals respondents hope to accomplish by implementing AI-enabled surveillance are **reducing false positives** and **improving risk identification**, demonstrating the desire to target common issues that can stem from the nature of utilizing lexicon-based mechanisms.

False positives and risk identification challenges impact surveillance and compliance teams' accuracy and efficiency, which directly correlates to time and budget – leading to a vicious cycle of business justification. When attempting to rationalize the benefit of buy-in needed to enhance surveillance programs, it's important that executive boards see an evident payoff.

How does AI enhance surveillance?

Outside of solely utilizing lexicon-based systems, there's a second tier of surveillance that has become available for firms to use. As firms begin to grasp the infusion of innovative technologies into business workflows, this method involves a combination of lexicon and deep learning models or Large Language Models (LLMs) to optimize risk detection through keyword implementation and substantiate identified risk through AI analysis.

> **There's a bit of an evolutionary step between lexicon models and LLMs. Lexicons are more using brute force – you cast the net wide. You know you're going to get a lot of irrelevant noise. You make your lexicon inclusive because you want to get all the risk. In order to get all the risk, you have to deal with a bunch of false positives, which means investing a lot of time in carving off all the noise.**
>
> **Though, time for refinement takes more effort to get even marginal improvements. Using smaller LLMs that aren't as capable of performing high-scale analysis is like an enhanced lexicon. It knows problematic phrases and looks for phrases that are similar by analyzing and breaking apart conversations. When it determines that something is a risky message, it then asks a generative AI model to explain what the problem is. It is more limited use to do a phrase-based comparison with a small model and then use a big model for the final step.**
>
> **In contrast, a [generative] LLM will analyze messages and give compliance teams a decent explanation of what it found. Instead of breaking [conversations] up into phrases, LLMs read like a human to understand what's going on in a whole conversation. This give us the power of the LLM as a risk detector, and not just an explainer of risk.**

Donald McElligott
VP of Compliance Supervision
Global Relay

Evolving from lexicon-based and deep learning systems, LLMs are the most advanced models for surveillance. LLMs offer an unprecedented enhancement to surveillance methodology, and divert from traditional risk identification models.

Where lexicon lists would target every instance of a keyword or phrase being used in conversations, which yields an overabundance of alerts for reviewers to evaluate, AI thinks as a reviewer would by analyzing complete messages to determine if risk is inherent. To do this, AI contextualizes conversations and uses its intelligence capabilities to understand the meaning behind what is being said.

In lieu of a keyword list, LLMs use prompts to identify risk. These models draw from an endless amount of data points when analyzing communications, which make it possible to ascertain sentiment.

Due to its contextual aptitude, AI also has the ability to recognize patterns within communications – an ability that previously would be assigned to human reviewers when sorting through risk alerts. In the case of keyword evasion or code language, this usage means that surveillance systems can flag suspicious behavior that would otherwise only have been spotted through random sampling.

What are compliance teams’ goal in implementing AI?

> **The biggest benefit we’re hoping to gain from using AI in surveillance is improving risk identification. We have a lot of lexicon hits a day, and many of them are false positives. We spend a lot of time clearing these false positives and need a better way to manage this to save time. I’m most excited to reduce false positives and better identify concerning communications that might not have come up through a lexicon-hit but were determined as risky by AI.**

Elise Kindig
Deputy Chief Compliance Officer
FCM Compliance
StoneX

Anonymous survey respondents:

> **To reduce the number of man hours spent on monitoring and to create a more precise review process.**

> **Efficiency, allowing experts to use their expertise, job enhancement, and risk-based control framework.**

> **Advanced surveillance capabilities such as sentiment analysis, risk scoring of alerts, and AI lexicon policies.**

> **Improve the accuracy and productivity of the compliance team.**

> **Improve detection, reduce false positives, and gain the ability to spot concerns at an earlier stage.**

> **Free up time.**

Up in the AI(r)

Keeping an open mind

Global Relay's Industry Insights: Compliant Communications 2024 report, published in June 2024, found that 42% of financial services said that they would integrate AI into compliance over the course of 2024 while 57% did not plan to.

Our latest survey, conducted in January 2025, shows a 19% reduction in the number of firms reluctant to implement AI, from 57% to 38%, suggesting that attitudes around AI are changing. If teams are not adopting AI models entirely, there is a wider group that is considering them.

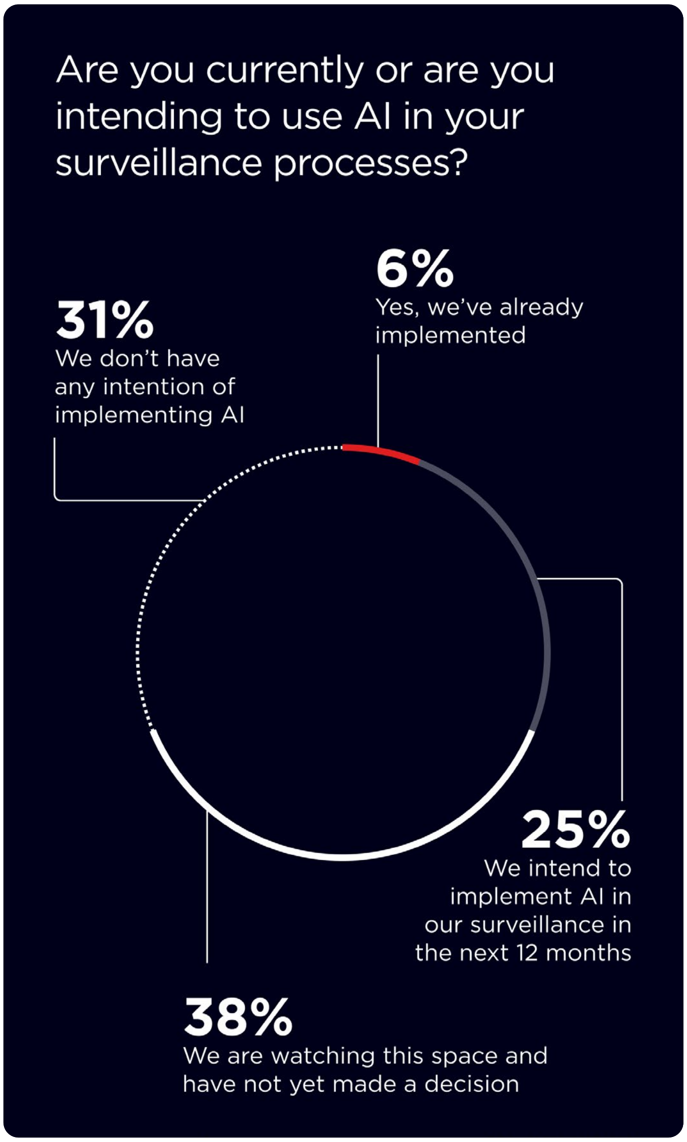
Are you currently or are you intending to use AI in your surveillance processes?

> The whole world of surveillance has been about dealing with false positives. Everything you did to this point was getting 100 false positives for every one interesting thing to review. All the resources were poured into just attacking false positives. This flips that over. LLMs don't produce false positives. If you feed an LLM clean data and reduce the noise by asking it questions, it's going to produce a lot of relevant risks that need to be analyzed.
>
> The shift now is attacking that real risk. You might get 100 messages a day you need to spend time investigating and working on, and this will change the culture of how surveillance teams operate. The days of throwing people at a big pile of data and sifting through to find a couple of needles are gone. Now, it's about looking into real risks to address. The toolbox of investigation and follow up is going to be important now. It's not just looking at one or two risks a day after sifting through massive piles of data – now, it's looking at 20 to 30 risks and determining how to explain it. Regulators are going have to adapt to follow this workflow. That's going to be the most radical change.

Donald McElligott
VP of Compliance Supervision
Global Relay

Blocks and barriers to adopting AI

The main barrier for firms adopting AI into surveillance workflows is data security, which has long been cited as a compliance concern for AI. Firms must be assured that AI and generative models, which function through data analysis, are not covertly storing or processing information. Organizations need to know how their data is processed, where it is stored, and who has ultimate ownership. For firms looking to implement third-party models, this is particularly critical.

A SecurityScorecard analysis found that 97% of the top 100 U.S. banks experienced a third-party data breach over 2024. In the case of AI-enabled surveillance, where firms enlist the support of third-parties to run generative models, sound third-party governance is nonnegotiable.

A close second to data security is budget as a barrier to implementation. Compliance and surveillance teams face barriers when requesting that executive boards invest in new tools. Surveillance tools, especially those using AI, can carry sizeable upfront costs, and surveillance teams must show return on investment for such tools.

> **I know the adoption of AI is new in the industry, so the biggest barrier for us as a compliance team is the unknown. How can we trust AI and what it is reviewing? Is it able to review communications and properly flag them? I would love to see a full month review and understand the difference between what AI caught versus what lexicons caught.**

Elise Kindig
Deputy Chief Compliance Officer
FCM Compliance
StoneX

Main barriers to AI adoption

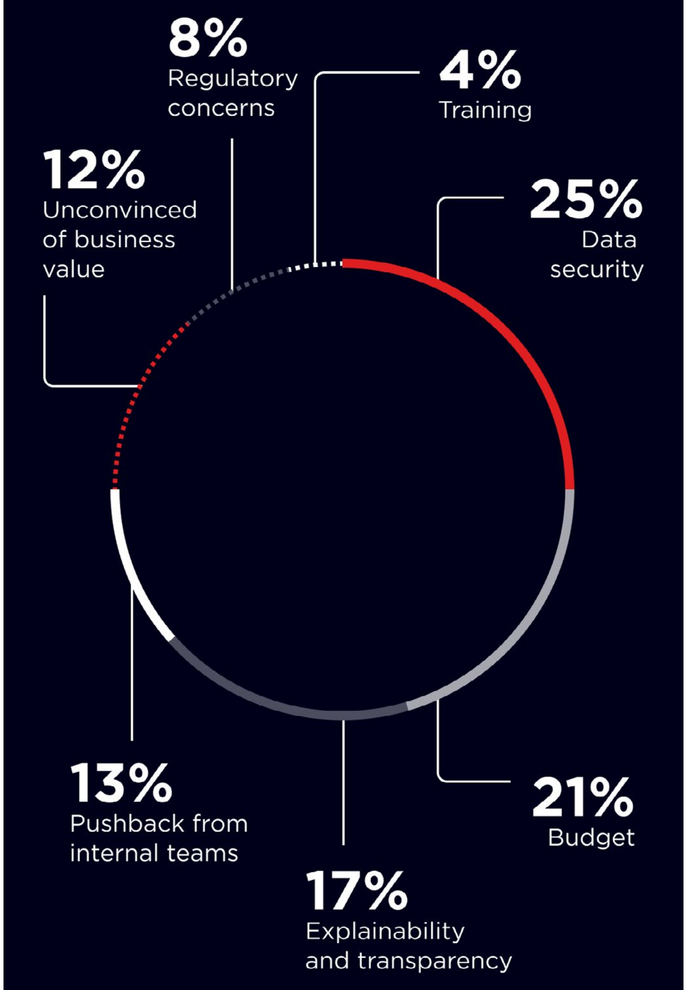

How does AI-enabled surveillance work?

Tackling explainability and transparency

Understanding of explainability

The third-highest barrier to AI adoption is a perennial concern surrounding AI – explainability. On a scale of 1 to 10, with 1 being not important and 10 being very important, respondents' average score was 9.14.

Respondents' average score was **9.14**

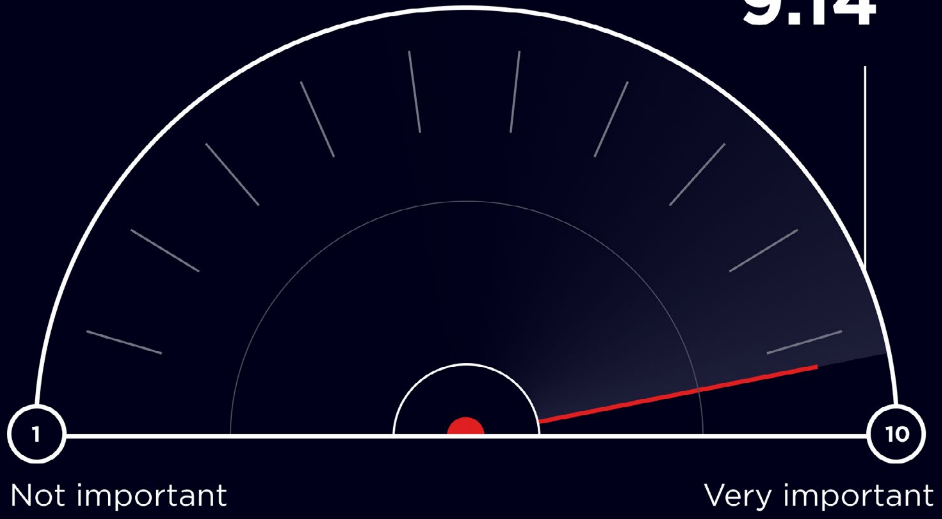

Though transformative, firms are apprehensive to use AI to completely restructure surveillance methodologies. The foundation of this concern is often that, in the event of regulatory investigation, firms using AI-driven processes will be unable to explain to the regulator how a decision was made.

Explainability has long been a barrier to AI adoption within surveillance and compliance. However, with latest AI advances offering significantly better results and far outperforming traditional or lexicon-based approaches, the question arises – is explainability more important than accuracy?

Is the benefit of accuracy outweighing the risk of explainability?

> Accuracy outweighing explainability has to happen. Anyone who has played around with Gen AI – especially in the comms space – will see how much it could add on top of traditional methods. So, it would be very hard not to use something that actually significantly enhances our detection capabilities.

Ugne Willard
Global Head of Communications Surveillance
Barclays

> The key is having the vendor documentation and vendor information up front when you're choosing to implement these tools. That's the key. Regarding explainability, from our perspective, we're seeing clients coming to us to independently test and assess how these models are working. We are being approached by model governance teams – not just compliance and surveillance teams – regarding independent testing.

Hannah Bowery
Senior Manager
Surveillance and Market Abuse SME
PwC

Anonymous survey respondents:

> I think it's extremely important to understand the processes, but at the same time, you can't ignore the fact that AI is significantly more accurate than humans and that should be utilized.

> You need to have both. I believe it is very dangerous to have anything you cannot explain to a regulator or internally to a trader. For example, when asking about an alert or activity. If you are not aware why the alert was generated, how can you be sure the trader response is correct?

> No, because confidential data needs to be protected, hence the existing data/AI framework needs to be strong. Explainability is a key step in this.

> Humans will need to check AI for accuracy and test.

Going hand in hand with explainability is transparency. To feel confident that the models being employed within their surveillance processes are carefully assessed and authenticated, firms want to see guidelines depicting how AI-enabled systems are managed, how they're trained, and how information is documented.

Respondents shared that a transparent AI governance framework is of the upmost importance. On a scale of 1 to 10, with 1 being not important and 10 being very important, respondents' average score was 9.4.

Third-party providers have a huge influence on a firm's business operations. With regulators across the globe advising firms to stay vigilant in their oversight of third- party risk, such as FINRA in its 2025 Annual Oversight Report and the Canadian Investment Regulatory Organization in its 2025 Annual Report, a thorough framework fosters necessary trustworthiness between clients and vendors. It also assures firms that the AI technologies they're investing in are being implemented responsibly.

Importance of third-party governance framework

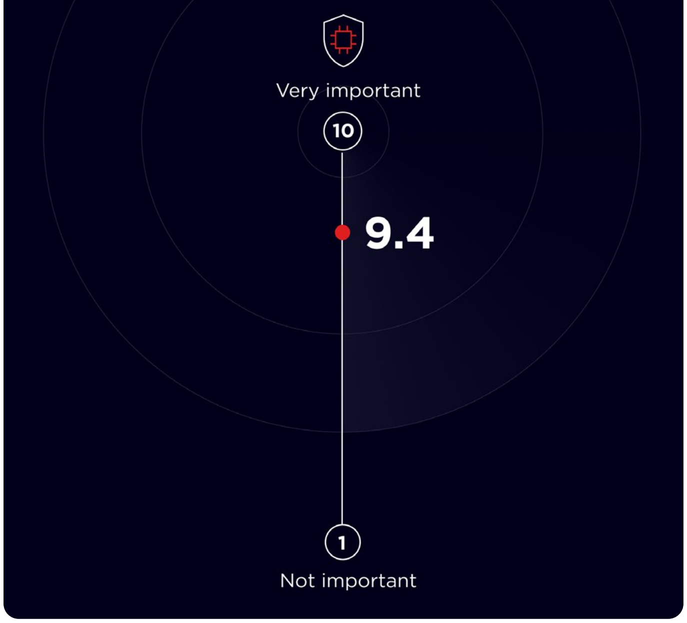

> **Going hand in hand with explainability is transparency. To feel confident that the models being employed within their surveillance processes are carefully assessed and authenticated, firms want to see guidelines depicting how AI-enabled systems are managed, how they're trained, and how information is documented.**

Alec Senne
Compliance Officer
Communication Surveillance
StoneX

Understanding how AI works remains a barrier

What's holding compliance teams back from taking the leap from lexicons to AI-enabled models? Before all else, firms want to first understand how modern technologies operate before implementing them into supervision processes.

In late 2024, a report published by the FCA and Bank of England showed that:

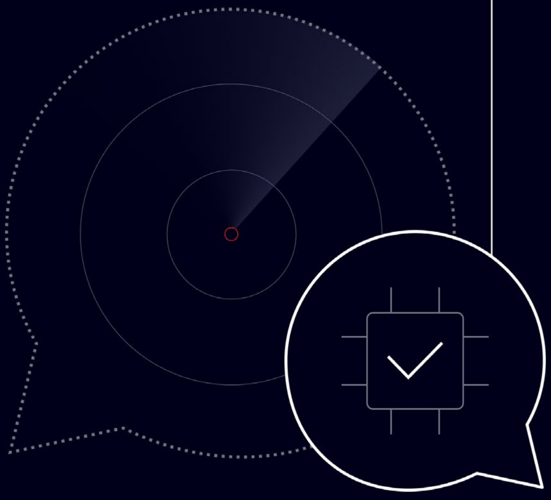

Only 35% of firms that use AI within their systems are able to "confidently explain" how it works.

2024 Report
FCA and Bank of England

When asked about the extent of their understanding of how AI works as a tool for surveillance, respondents to Global Relay's survey indicated there's still a way to go.

Public understanding of AI in surveillance

On a scale of 1 to 10, with 1 being no understanding and 10 being full understanding, respondents' average score was 3.6.

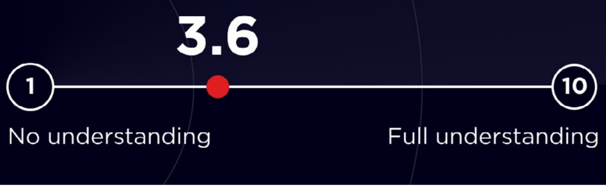

We also asked respondents if their surveillance and technology teams have the requisite knowledge to implement and manage AI-driven surveillance. As depicted in the graph below, a sizeable number of firms are keen to train employees before they are equipped to manage new systems, as there is still a limited understanding of how AI-enabled technology works.

Do you feel your Surveillance and Technology teams have the requisite knowledge to implement and manage AI-driven surveillance?

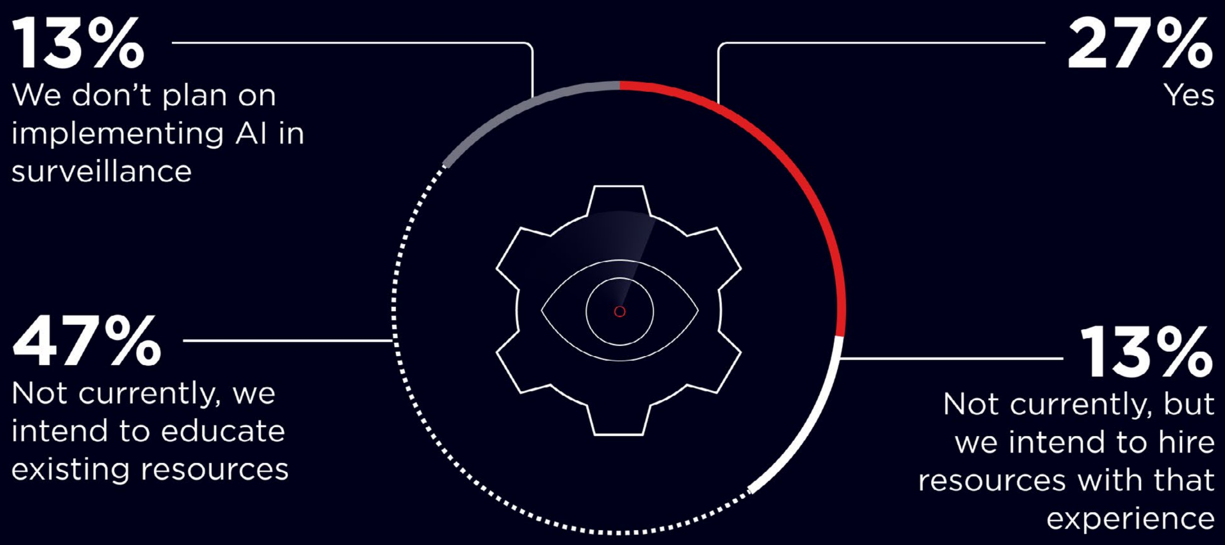

Mythbusting misconceptions: How do LLMs compare to other surveillance systems?

Explainer

Lexicons vs. AI-enabled models

When it comes to lexicon-based models and LLMs, the distinction lies in the scope of capability.

While the benefit of a lexicon is its traceability, the lack of flexibility outside of prescribed keywords means suspicious conversations that evade surveillance attempts could go by unnoticed. In the same sense, an influx of false positives could cloud a reviewer's ability to pinpoint true risk.

If a compliance team is trying to detect insider trading, they would have to create a list of all possible words or phrases that could suggest related misconduct, such as "leak" or "tip." Yet, if bad actors use code words or purposefully misspell related language to evade surveillance, it could lead to missed risk flags. Alternatively, if personnel used these words in innocuous conversation, it would lead to an abundance of false positives.

LLMs modernize this process by replacing lexicon implementation with prompts, which can analyze context and read between the lines to get to the heart of a conversation and determine if risks are present.

For example, by telling an LLM "You are a compliance officer analyzing risky conversations" and directing the system to justify potential detections by pulling key phrases, context, related risk categories, and implications, it can accurately substantiate evidence of possible misconduct.

AI-enabled models also enhance the ability to perform voice surveillance. Pulling from the datasets that equip it with a comprehensive knowledge to recognize speech patterns and differentiations, LLMs can rapidly transcribe and translate calls into written conversation. This saves compliance teams from having to listen to hours of voice recordings to identify possible risk areas.

> **This is a capability that is particularly pertinent being that survey respondents elected voice transcription as the third area they hope to address by implementing AI into surveillance practices.**

Explainer

Industry-specific vs. open source LLM models

Though the difference between industry-specific (also known as deep learning) and open-source models isn't as initially apparent, an open-source LLMs' database is exponentially more comprehensive than an industry-specific model, giving it a substantial edge.

A common misconception is that since industry-specific models are trained with information pertaining to a specific domain, such as with finance, they are better suited for financial firms. Yet, an LLMs' knowledge base is nearly limitless.

The entire scope of an industry-specific model makes up only a small subset of what comprises an entire LLM system. An LLM is 200-300 times the size of a small-scale model, and its universe of knowledge means it already has all the necessary datapoints to advise it of the risk areas on specific industries.

Since an industry-specific model is trained on data that is related to a specific market, it is not as efficient at identifying behaviors that exist outside of that market. This limited scope of data could also lead an industry-specific model to conceive risks that aren't actually there by creating an echo chamber of parameters that are biased toward a certain set of risks.

Like lexicon-based systems, industry-specific models must be fed information to define risk areas, whereas LLMs don't need training and instead are guided by prompts to analyze content.

> There's the sentiment that you need a curated dataset to make models efficient, but this isn't true of LLMs. They give you a better understanding and are in a better position to distinguish between what's really a hit and what isn't.
>
> We have to tackle this sentiment that's dated back to machine learning. People wonder how an open source LLM might be able to distinguish financial slang from someone discussing their weekend plans. That's because it has enough reference points to do that. It's not over trained. It's more likely to have references similar to anything that we input.

Robert Nowacki
Senior SME
Global Relay

Explainer

AI adaptability vs. static lexicon

LLM models are constantly evolving to adapt new techniques that aid in deciphering data. While the techniques of bad acting in the financial world are relatively static, the techniques to locate these risks are progressing. Each new LLM model that is released is built with the objective to reach a higher level of data completeness.

New models are given a larger volume of knowledge to work with, and are also programmed to run faster and more efficiently. All of the knowledge needed to feed AI models is already existing within the scope of the internet, and now the focus is on massaging techniques and capabilities into each newly released model.

With lexicons, compliance teams need to manually update keyword lists to adapt to emerging risks, which requires more effort from compliance teams to sharpen their review process, detracting from productivity.

> Having an LLM that understands all the data on the internet, including financial knowledge, provides a deeper understanding within multiple domains. This enables transfer learning, where knowledge in one domain is applied to another one.
>
> The LLM we use for surveillance is hundreds of times bigger than deep learning models typically used, meaning it does not require additional training, whereas smaller scale deep learning foundation models cannot be used for a specific task without additional training.
>
> A deep learning foundation model is like buying a car chassis, which you would have to assemble before you can drive it. You need to put the engine, seats, and wheels on the car. With an open LLM, you can buy a car from the dealer that's ready to drive directly.

Yifan Xia
Product Director – AI
Global Relay

How will AI change the future of surveillance?

Is there enough AI guidance?

Though optimistic, regulators across jurisdictions have taken several different approaches to AI's adoption. Where the U.K. is concerned, that means being pro-innovation, and on the EU side, it means taking a direct, risk-based approach.

U.S. regulators have taken a cautious stance on AI adoption, focusing more on breaking down how AI technologies – specifically generative – work to establish parameters around its utilization. Being that AI is so complex, the main concentration seems to be on its safe use and accuracy.

We asked survey respondents if they feel there's enough guidance from regulators on expectations around AI, how it should be used, and what they perceive as risk areas.

When asked whether regulatory guidance was sufficient, respondents commented:

> **I feel like this is still somewhat of a gray area, but the regulators are quickly getting up to speed.**

> **Not at all.**

Despite cautions, regulators have remarked that it's undeniable AI is and will continue to have a huge impact on industry operations, such as by optimizing manual tasks or streamlining customer service capabilities. In an interview with Politico Tech, previous SEC Chair Gary Gensler said that AI is "the most transformative technology of our time, on par with the internet and mass production of automobiles."

With AI presenting opportunities to modernize how organizations perform complex operations, regulators urge industry players to contemplate how it could allay ongoing issues. During the CFTC's "AI Day" meeting on May 2, 2024, the Chief Innovation Officer of the Federal Reserve Board commenced the event by asking a pertinent question about AI innovation:

> **Are we looking at this new technology and new capability…in the context of solving gnarly problems? Are we starting with a thesis and thinking about how this technology will help us solve it? … Are we designing meaningful optionality solutions that might help us level up the institution for the present…but also for the future?**

Sunayna Tuteja
Chief Innovation Officer
Federal Reserve Board

Where do attitudes stand?

Respondents still have yet to fully trust the conclusions of AI models when making context-driven decisions on whether communications carry a risk. On a scale of 1 to 10, with 1 being not very likely and 10 being very likely, respondents' average score was 4.92.

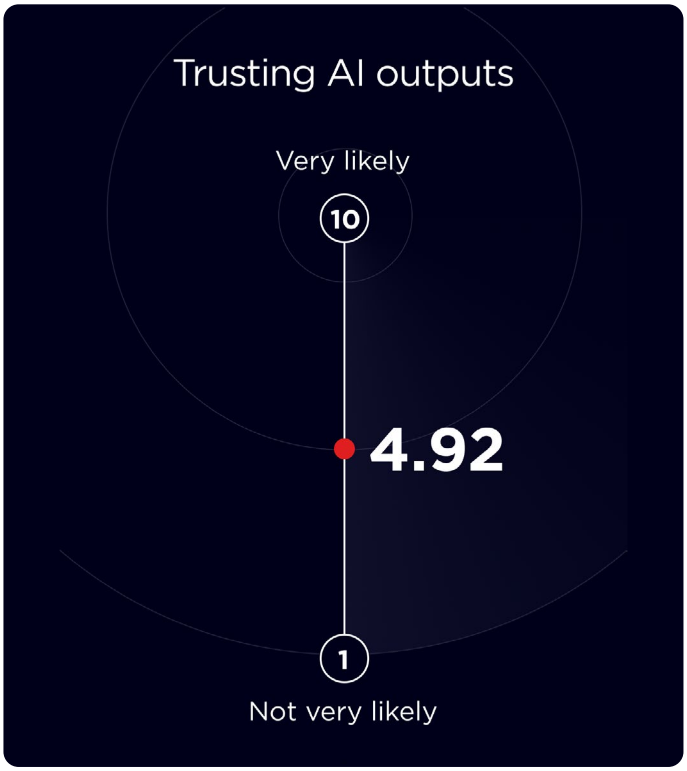

Nearly halfway between minimal trust and complete trust, industry feelings toward generative AI seem to be at a crossroads. As with responses we've seen suggesting a gradual attitude shift, it seems the industry is coming upon a revolutionary change to compliance workflows.

As techniques and documentation become more thoroughly polished, as well as with generative AI capabilities vastly advancing compared to existing surveillance models, it will be a matter of time before generative models gain real momentum.

Find out more about Global Relay Surveillance →

Disclaimer

This report and the graphs therein refer to data collected by Global Relay Communications Inc. The information, materials, and opinions contained in this report are for general information purposes only, are not intended to constitute legal or other professional advice, and should not be relied on or treated as a substitute for specific advice relevant to particular circumstances.

Global Relay Communications Inc makes no warranties, representations, or undertakings about any of the content of this report (including, without limitation, any as to the quality, accuracy, completeness, or fitness for any particular purpose of such content), or any content of any other website referred to or accessed by hyperlinks through the report. Although we make reasonable efforts to update the information on our reports, we make no representations, warranties or guarantees, whether express or implied, that the content is accurate, complete, or up-to-date.

March 2025

global RELAY

Integrity | Reputation | Control

North America: +1 866 484 6630
Europe: +44 (0) 20 3206 1850

globalrelay.com