NICE Actimize

eBook

AI-Assisted Alert Reduction

Demystifying Alert Prediction in Financial Trade Surveillance

1. Introduction

According to the Chartis research survey, The Future of Trader Surveillance, one of the biggest challenges financial compliance and risk teams have to contend with is the high number of false positive alerts. Compliance analysts are on the front line of helping their firms detect employee misconduct and market abuse. Yet, the sad reality is – they waste so much time chasing false positive alerts that it can distract them from the important work at hand.

So, what exactly is a false positive alert? A false positive alert scenario occurs when a transaction (or communication or behavior) initially identified as suspicious by the surveillance system, is later found to be invalid.

How big of a problem are false alerts? Based on our discussions with financial firms, compliance analysts at most tier one banks review about 1000 alerts a day. Even more frustrating, over 99% of these alerts turn out to be false.

Managing the deluge of alerts is a cumbersome task which can tap already stressed compliance resources and add to overall compliance costs. Even worse, the only way to confirm which alerts are false or true (with any degree of certainty) is to sift through each one manually.

One way to solve this problem is through the use of ***alert prediction***.

False alerts out of control?
Time for a new approach

> Compliance analysts review about 1000 alerts a day
>
> – over 99% of these alerts turn out to be false

2. Alert Prediction at a Glance

Alert prediction is a form of predictive analytics which uses machine learning (a type of AI) to transform data into insights. To put it another way, historical alerts are analyzed using supervised machine learning. From this analysis of past alerts, the surveillance system learns how to accurately predict the outcome of new alerts.

The goal of alert prediction is to map specific input variables to specific outcome(s). The alert prediction algorithm uses input data related to alerts (such as alert create reason, confidence scores, communication types, transaction type and alert dispositions) to categorize new alerts into specific classes. These could include categories like "Close-Issue" (aka true alert), and "Close-Non Issue" (aka false alert). The algorithm is able make these predictions accurately because it has already been trained on historical alert data.

When presented with a new alert, the algorithm predicts the outcome of the alert (true or false) based on the cumulative, previous data it was trained on. Because the algorithm continues to learn from newer, cumulative data, the alert prediction only gets smarter over time.

Alerts prediction is a three-step process:

1. **Data Processing and Handling** – First, historical alert data must be cleaned and pre-processed.

2. **Model Training** – The next step involves model training. Included in this step is the balancing of the training dataset which is important to reduce imbalanced data. Imbalanced data happens when there is a high disparity in the number of true and false alerts.

   The number of true alerts is usually much lower (under 1 percent) and this can lead to imbalanced data. The problem can be addressed through various methodologies, such as over sampling of true alerts (a technique used to increase the size of very small samples), and similarly, under sampling of false alerts (a technique used to decrease the size of comparatively larger samples).

   During this step, the model is also trained utilizing the sample data.

3. **Alert Prediction and Explanation** – Once the model is trained, alerts can be sent to the prediction algorithm for analysis. The resulting prediction outcome is then sent back to the surveillance system.

   The prediction and explainer output are then displayed along with the alert so the compliance analyst can review the alert effectively.

> "Surveillance systems can learn how to accurately predict the outcome of new alerts"

3. Alert Prediction: Empowering Better Decision-making

Because alert prediction analyzes a myriad of data, the predictions can be enriched with information which ultimately helps compliance analysts save time and make better decisions. The following example illustrates this point.

Let's assume *Company A* generates **1,000 communication surveillance alerts each day**. In the absence of alert prediction, *Company A*'s compliance team would need to review every single alert and that would take a lot of time. With alert prediction, each alert comes with additional information (see the example on the right). This enriched information enables compliance analysts to prioritize alerts based on their relevance and importance. Analysts can easily segregate true and false alerts and focus on the ones that need immediate attention.

As you can see from the fictitious alert prediction outcome data for the communications and market surveillance alert examples (to the right), the analyst can immediately see that both the alerts are predicted to be false, with **80 and 90 percent confidence** respectively, for the precise reasons stated in the rightmost column. The explanation also includes an analysis of each contributing factor enabling compliance analysts to zero in on specific areas during the alert review process.

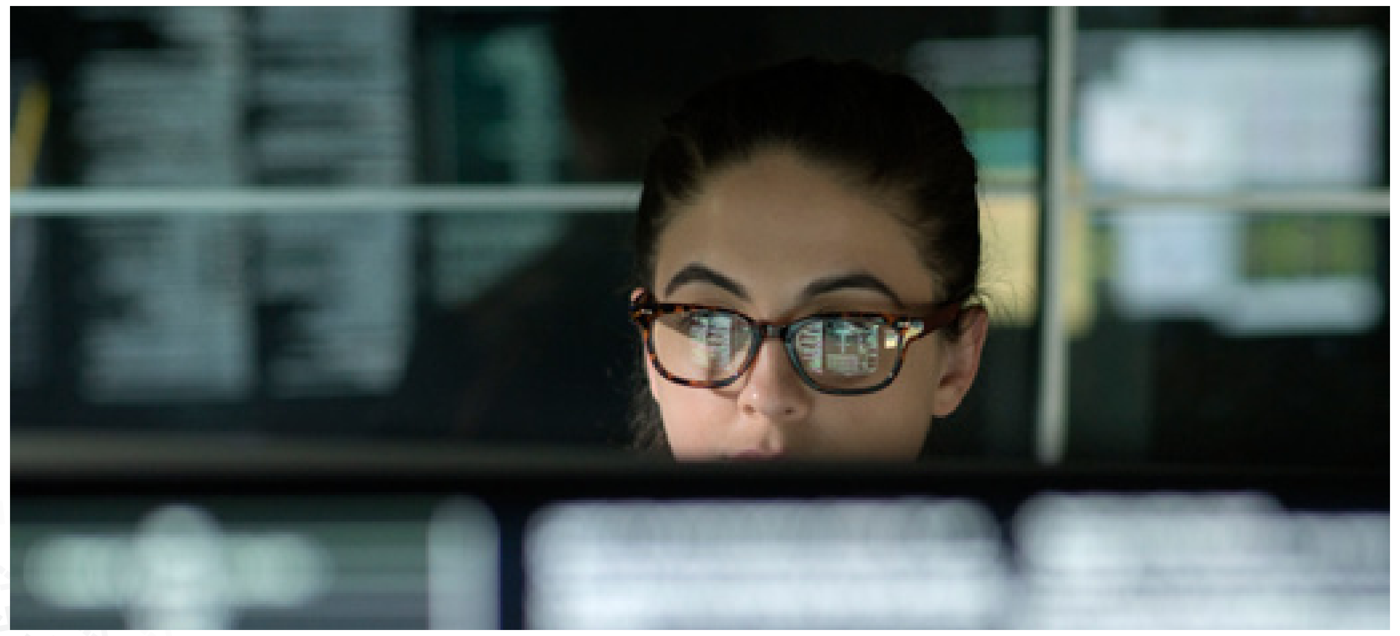

Communication Surveillance Alert Example

Prediction Outcome	Prediction Probability	Prediction Explanation
Close-Non Issue	80%	• Conversation Start Hour / 30% • Interaction Recipients / 15%, • Participant Count / 22% • Communication Types / 10% • Key phrases / 23%

Market Surveillance Alert Example

(Alert Type = Marking The Close)

Prediction Outcome	Prediction Probability	Prediction Explanation
Close-Non Issue	90%	• Look Back Period / 15% • Previous Side / 20% • Quantity / 30% • Side / 10% • Time Window / 25%

Armed with this knowledge, the compliance analyst may choose not to review both of the "Close-Non Issue" alerts at all, or only review them after reviewing all of the true alerts in her queue.

> "Alert prediction can help compliance analysts save time and make better decisions"

4. Alert Prediction by the Numbers: Powerful Results

In our experience working with customers in the financial services sector, I've seen firms achieve significant benefits when alert prediction is integrated into their communications and market surveillance programs. Here is a snapshot of some of these results:

- On average, firms were able to reduce false positive alerts by up to ***70 percent***. Prior to deploying alert prediction, one customer reported reviewing approximately 4,000 alerts weekly. After deploying alert prediction alerts declined by ***78 percent*** (to 900 weekly).

- Another customer reported being able to identify false alerts with ***90 percent*** accuracy. This helped the firm to de-prioritize false alerts and focus on the true ones. The firm's compliance team achieved around a ***25 percent*** time savings by automatically silencing "Close-Non Issue" (aka false) alerts that included a confidence score of ***90 percent*** or above.

- Many customers have also stated that the fully explainable prediction process has enabled them to easily integrate alert prediction into their existing workflows, and that in turn has increased their business efficiency. With alert prediction and full explanations supplementing their decision-making process, they no longer need to make blind decisions.

- Firms also appreciate how alert prediction has enabled their compliance teams to plan their workdays more effectively. They know at the start of each day which alerts are true and which are false, with predictions supported by confidence scores. Analysts can prioritize their alert review queue and make better use of their time.

- Firms also report better alert accuracy because they're able to deploy dedicated models for different alert types (rather than having one global model for all alert types). For example, a firm might have one model for "Large Order Entry" and another for "Painting the Tape."

- Firms have also benefitted from increased model accuracy over time. Models are constantly retrained to incorporate compliance analyst reviews that either validate or invalidate the prediction outcomes.

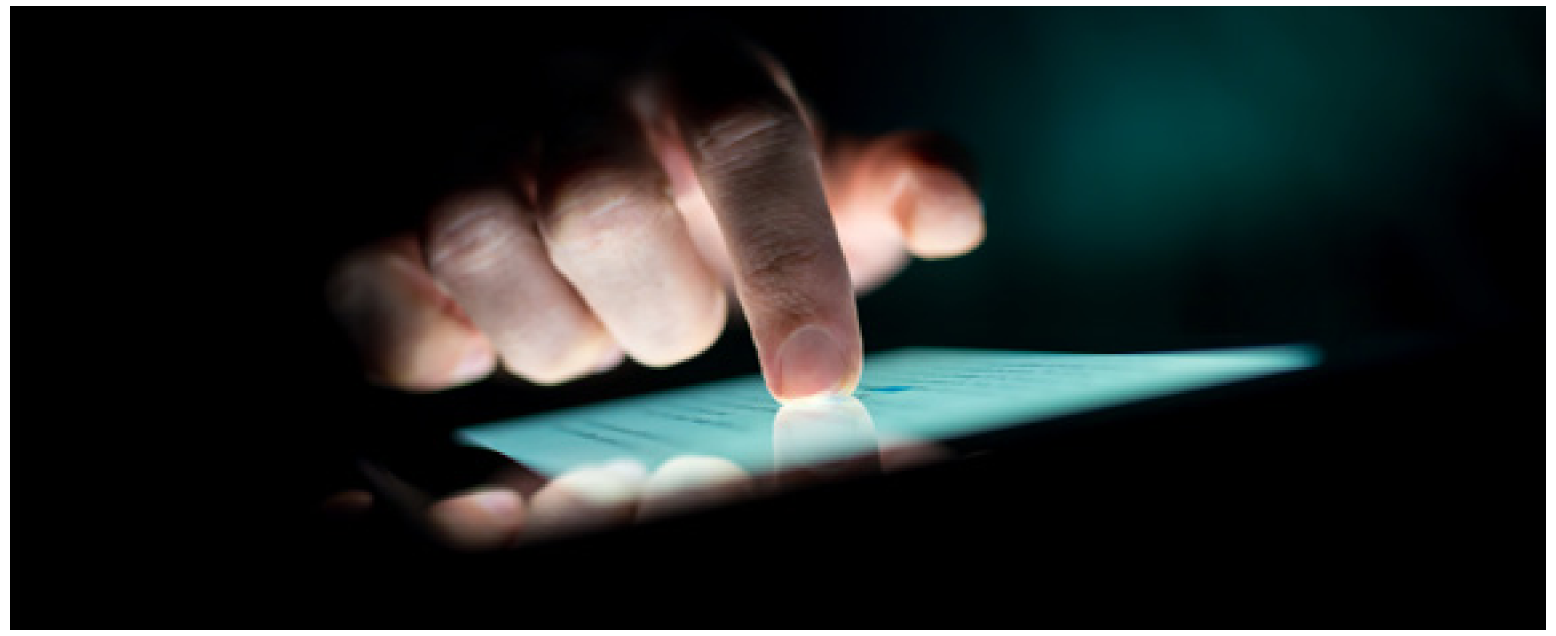

> "Firms can achieve significant benefits when alert prediction is integrated into their communications and market surveillance programs"

5. Steps to Implementing Alert Prediction

As you implement alert prediction there are a number of factors you need to consider. Careful consideration of these four factors can help your firm avoid common missteps and increase your chances of success.

Step 1 - Focus on Data Quality

The accuracy of alert prediction is only as good as the underlying data that feeds the machine learning model. Fortunately there are steps your firm can take ahead of time to improve data quality. Here are three key things that should be reviewed at least three to six months prior to your firm's deployment of predictive alerting:

- **Data Labelling** – Alert prediction relies on supervised machine learning. The surveillance system uses supervised machine learning to 'learn' from the labels that compliance analysts attach to historical alerts (e.g. "Closed-Non Issue," "Closed-Issue," etc.). Without accurately labeled data for all alerts, it's impossible to train the machine learning models. Additionally, alert dispositions marked as "In-progress," "Ready," and "Pending" can't be used to train models. The implication for your firm: you need to be vigilant in ensuring that historical alerts have been thoroughly and accurately dispositioned.

- **Uniform Alert Review Process** – The lack of uniform alert review processes can also impact data quality and hamper accurate alert prediction. Alert review processes and the meaning of data labels often vary from financial institution to financial institution, and can even vary across analysts (within the same compliance department/institution).

  Without uniformity, it's hard to infer what data labels mean. For example, the review reason "Investigation" could mean different things to different institutions and/or analysts. For example, one client we consulted with told me that they consider "Investigation" to be the same as "Close-Issue," while other clients have disagreed, citing that investigated alerts don't always result in a closed issue. Additionally, some analysts may insert a comment when reviewing an alert without changing the status from "In progress" to "Closed." Still other analysts may change the alert status, without noting any reasons. These inconsistencies can make it difficult to train the machine learning model, and as a result, can hinder alert prediction. The implication: you want to make sure your firm carefully considers how it instructs analysts to label alerts before training the prediction model.

- **Balanced Data Set** – In working with some financial institutions we've found that historical alert data can be imbalanced – in other words significantly more likely to be false than true. For example, with some clients we've worked with, out of 10,000 alerts, only a very small fraction (around 50) were found to be true or close to true (aka "Close-Issue"). This dramatic imbalance between true and false positive alerts is a hindrance to machine learning models which rely on correctly labeled large data sets to learn.

  One way to overcome this problem is to work with the client to figure out if other types of alerts, other than "Close-Issue" can be categorized as true alerts to reduce the data imbalance. For example, in discussions with one client, it was determined that alerts categorized as "Close-Analysis" and "Investigation" tended to be high quality alerts, so the client requested these alerts be reclassified as true alerts to address the problem of data imbalance prior to model training.

> "The accuracy of alert prediction is only as good as the underlying data"

Step 2 - Federated Versus Individual Models

The next important consideration for firms is selecting a federated or a client-specific, individual model. Both have their advantages, but the ultimate deciding factor is your firm's business needs. The federated approach to modeling enables prediction models to train on different datasets (multiple clients/ extensive industry data) whereas individual models are trained solely on individual client data sets.

Each model type has distinct advantages. For example, federated models offer data diversity because they are trained on industry data (originating from a number of firms). This enables creation of shared global models that leverage broader, diverse data sets. On the other hand, because they are trained on specific client data, individual models can be customized and tailored to specific client needs. For example, multiple individual models can be customized to different lines of business (within a firm) or even tailored to compliance teams scattered across different geographies.

When working with firms, we generally advise them to start with a federated model first due to scarcity of labeled data. Over time, the alert prediction model can be tuned to the firm's individual data, when data is more abundant. This deployment methodology combines the benefits of both approaches – federated and individual.

Step 3 - Fully Explainable Machine Learning Results

Firms also need to carefully consider if, and how they want their compliance analysts to use machine learning results in business decisions. Careful consideration must be given to this question as it can alter business processes that the analysts need to follow. For example, will compliance personnel make decisions simply using the alert predictions? Or will they also need an explanation about the predictions – why a prediction was made, and what factors contributed to it?

Consider the example of machine-based prediction to approve or reject loans. To make better decisions, loan officers might want to review all the factors that contributed to the bank's decision to extend or not extend a loan – including the applicant's credit history, consistent income, and so on. Firms need to think the same way about alert prediction. To ensure analysts can make business decisions with a high degree of confidence, firms may want to provide fully explainable machine learning results. If this is important to you, make sure your technology vendor provides this capability.

Step 4 - Know Your Firm's Tolerance for Error

Alert prediction solutions are not 100 percent accurate. Their purpose is to assist compliance teams, not replace them. Knowing your firm's tolerance for error is crucial to which alert prediction solution you select.

Carefully consider your tolerance for the following scenarios:

1. **An alert is predicted to be true but turns out to be false.**
2. **An alert is predicted as false/negative and turns out to be true.**

If your firm's tolerance for error is low, I recommend partnering with a company with technology that backs up alert predictions with confidence scores that are fully explainable. Decisions to prioritize true alerts over false alerts, or suppress false alerts automatically are far more clear-cut when you know the confidence level of alert predictions, and can see that the predictions are fully explainable.

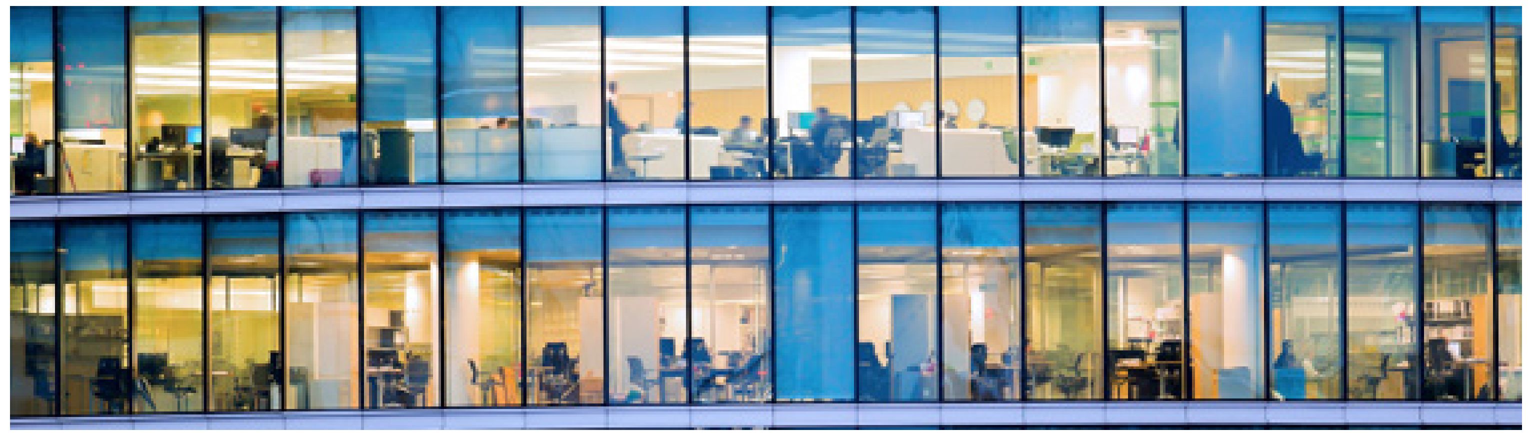

> “Firms need to carefully consider if, and how, they want their compliance analysts to use machine learning results in business decisions”

6. Exploring Alert Prediction for Your Compliance Organization

This eBook touched on the benefits of alert prediction, and some best practices and tips to get you started. But as you embark on your Artificial Intelligence (AI) journey toward more accurate alert prediction, challenges will come up. As the largest and broadest provider of financial crime, risk, and compliance solutions for regional and global financial institutions, **NICE Actimize** can help.

Our supervised and unsupervised machine learning solutions have been successfully deployed at many leading financial institutions. No company is better equipped to help you understand where and how to apply AI and machine learning for optimal surveillance results. Finally, our surveillance drill-down dashboards remove the mystery of AI by providing complete explainability and confidence scores for every alert.

Still have questions? Contact the author:

Nitin Vats
Product Specialist
NICE Actimize
nitin.vats@nice.com

> “As you embark on your AI journey toward more accurate alert prediction – challenges will come up”

NICE Actimize

Financial Markets Compliance

NICE is a leading financial compliance solution provider, serving more than 90 percent of the largest investment banks globally. NICE's compliance solutions assist customers in the capture of trade conversations and trades, analyzing them for potential risk, and correlating trade conversations with trades for trade reconstruction. The company's compliance solutions make automated and intelligent holistic trade compliance programs possible and enable FSOs to more efficiently comply with regulatory requirements, including the future Consumer Duty rules , MiFID II, MAR, FX Code of Conduct, Dodd-Frank and future directives.

NICE Actimize's SURVEIL-X Holistic Conduct Surveillance offers unparalleled risk coverage for online brokers, buy-side and sell-side firms, insurance companies, crypto exchanges, regulators, and more by enabling accurate detection and rapid, thorough investigation of market abuse, inappropriate sales practices, conduct risk, and otherwise undetectable compliance risks to insulate firms from fines and reputational damage.

www.niceactimize.com/compliance

Download the **SURVEIL-X** Brochure

About NICE Actimize

NICE Actimize is the largest and broadest provider of financial crime, risk and compliance solutions for regional and global financial institutions, as well as government regulators. Consistently ranked as number one in the space, NICE Actimize experts apply innovative technology to protect institutions and safeguard consumers' and investors' assets by identifying financial crime, preventing fraud and providing regulatory compliance. The company provides real-time, cross-channel fraud prevention, anti-money laundering detection, and trading surveillance solutions that address such concerns as payment fraud, cybercrime, sanctions monitoring, market abuse,customer due diligence and insider trading.