




Defy. Deploy. Run.

Machine Learning in Trade Surveillance

LPA CONSULTING WHITEPAPER

October 2023



Executive Summary

Legacy trade surveillance systems frequently suffer from a high number of generated alerts, which are the result of false positive warnings. Those alerts require extensive attention from compliance officers. As a result, legacy systems do not only impose high cost related to manual alert handling and maintenance but do also open windows for market abuse practices remaining undetected.

Following regulators' call for improved market monitoring practices, Machine Learning and AI-based systems offer a path towards less costly and more accurate trade surveillance. By pointing out the various use cases and benefits of integrating Machine Learning techniques into trade surveillance, this whitepaper provides the outline for the transition to AI-based technologies, explains the arising opportunities, and sheds light on associated challenges and approaches to overcome them.

The application of Data Analytics and Machine Learning can significantly improve the parametrisation of surveillance models, thereby reducing false positives and increasing chances of more effectively detecting cases of potential market abuse. In addition, the handling of alerts generated in legacy surveillance systems can be improved by generating AI-based alert scorings, allowing compliance officers to focus their attention to the most critical cases. Using a higher degree of AI, surveillance systems can be trained on historical trading, market trends, and compliance practices and thus can contribute to the generation of more meaningful and less false positive alerts.

As the information fed into surveillance systems typically varies in data source and format, AI-based systems can incorporate and standardise initially unstructured data through elaborate deep learning models such as Natural Language Processing (NLP). Data availability is also a great point of concern, as the lack thereof can be detrimental for trade surveillance. Synthetic data, generated through AI algorithms, can overcome such information deficiencies.

A future advancement lies in the utilization of Generative AI (GenAI). This technology holds the promise of further enhancing trade surveillance functions across various dimensions, including streamlined development, more efficient implementation and calibration, and improved investigations. This way, Generative AI is poised to enhance the overall effectiveness and efficiency of the surveillance process.

In summary, Machine Learning and AI-based systems offer new opportunities for improved trade surveillance practices. These systems are not only beneficial to the institutions deploying surveillance but also to regulators. In such a transition to state-of-the art technologies, it is of paramount importance to understand the underlying mechanisms, the vendor solutions, and ultimately, through well-defined steps, the outcome goals. Proper training allows proper reporting and a concise and transparent information transmission between compliance officers and regulators.

Table of Contents

1	Current Challenges in Trade Surveillance	4
2	Machine Learning Use Cases in Trade Surveillance	6
2.1	Model Calibration	7
2.2	Alert Scoring	9
2.3	Alert Generation	12
2.4	Integration of Unstructured News Data	15
2.5	Integration of Communications Surveillance	17
2.6	Data Augmentation	19
3	Generative AI: Redefining Trade Surveillance Dynamics	22
4	Opportunities and Risks arising from the application of Machine Learning in Trade Surveillance	24
5	Conclusion	25

1 Current Challenges in Trade Surveillance

With the Market Abuse Regulation (MAR) being in effect since 2016, regulators call for better and more efficient monitoring practices to increase the integrity and resilience of financial markets, as well as to improve investor protection.

- **A lack of common market standards contributes to inefficiencies in legacy systems**

Since the introduction of the MAR, the majority of players in the financial industry have implemented some form of automated market surveillance system to comply with regulatory requirements. Affected institutions face the challenge of selecting a solution that best meets their functional and technical requirements and is also appropriate for their business activities and inherent compliance risks. This typically leads to the implementation of a solution provided by an external vendor with off-the-shelf surveillance models. What is lacking is a common market standard for surveillance approaches, models, and practices.

Our practical experience and recent surveys¹ show that currently implemented solutions are often inefficient, requiring a significant amount of manual effort in order to process what turns out to be mostly false positives of systems that are unable to disentangle the relevant from the irrelevant. The inability to efficiently identify cases of actual market abuse or the intention thereof and the insufficiency of narrowing down false positive alerts is the blind spot of many legacy surveillance systems. Certainly, this is one of the major contributing factors why many institutions are unsatisfied with their current surveillance setups.

- **Improving existing solutions is challenging, but often necessary**

A further complicating factor is that regulatory requirements leave a lot of room for interpretation, making it difficult to translate them directly into technical surveillance models. Due to high complexity and versatility of products and markets, banks commonly rely on multiple surveillance solutions, calling for the need of a more centralised, holistic solution that integrates the alphabet of market surveillance.

In addition to internal pressures towards more efficient and effective systems, market surveillance has also been in the regulatory spotlight over the recent years and the industry is expecting regulators to further increase their focus on MAR for future regulatory audits.

¹ [Governance of Market Abuse Surveillance Controls](#): An industry perspective, Association for Financial Markets in Europe (afme), January 2021

Improving surveillance capabilities remains a challenging endeavour, mainly because of the following reasons:

- Quality, availability, robustness, consistency, accuracy, and timeliness of data are a common issue.
- Out-of-the box vendor models are overly simplistic and fall short to account for the specifics of the individual business models of the institution that deploys surveillance.
- Many institutions rely on legacy IT systems and make use of a silo approach which does not support an integrated view on the monitored activities. Hence, such institutions do not support an integrated analysis of communications and trade data which could greatly amplify the effectiveness of trade surveillance.
- Technological innovations such as algorithmic and high frequency trading lead to more sophisticated trading strategies and new possibilities of market abuse.

— **The application of Machine Learning is the logical next step to overcome surveillance challenges**

Another reason might be the lack of guidance or transparency which slows down the transition to modern surveillance technologies that leverage elaborate Machine Learning models and make use of Data Analytics techniques.

To counteract the arising concerns of market abuse, an adoption of Machine Learning based surveillance systems seems to be the logical next step that market participants need to take to ensure compliance in an efficient manner. As vendor solutions become more sophisticated, a clear understanding of how these systems work will be necessary. In addition to the use of Machine Learning techniques as an integral part of the surveillance models, technical know-how and data competency will become increasingly more crucial within surveillance divisions.

2 Machine Learning Use Cases in Trade Surveillance

Two key dimensions must be considered when applying AI/Machine Learning techniques in trade surveillance: data processing and alert generation. For each dimension there are several use cases in which Machine Learning can create a business benefit for financial institutions, depending on the business model and surveillance strategy.

When it comes to applying Machine Learning algorithms to data processing, the most promising use cases revolve around data structuring and augmentation. Alert generation on the other hand can be improved by utilising Machine Learning algorithms for model calibration, alert scoring, and alert generation.

Although many financial institutions are planning to adopt artificial intelligence in one way or another in their trade surveillance system, most of them are lacking a specific transition roadmap. This is concerning because the transition from a static rule-based system to an intelligent system provides numerous benefits but also poses some challenges. These challenges need to be identified and actively managed to avoid pitfalls and create actual business benefits from the implementation of artificial intelligence.

A smooth transition from a rule-based system to an intelligent system can be accomplished by implementing use cases with a lower degree of AI to gradually improve the surveillance system. An advantage of a gradual transmission is that it ensures that surveillance experts are still able to monitor the process and intervene where necessary. Comprehension, traceability, and auditability are major concerns when implementing artificial intelligence, particularly in the field of compliance. The importance of human comprehension and intervention in AI systems has been recently highlighted in a report of the European Commission.² Therefore, it is not sufficient to make use of AI if its utilisation and algorithmic logic cannot be explained to a non-technical audience. The same is also highlighted in a recent report published by the German Federal Financial Supervisory Authority, BaFin.³ The latest developments in the area of Explainable AI (XAI) and the application thereof in trade surveillance will contribute to achieving this goal.

LPA has identified five use cases for the application of AI in trade surveillance which vary in their degree of AI-usage as well as in the complexity of their implementation (Figure 1). Each of these use cases is described in more detail in the following chapters.

² [Report on the safety and liability implications of Artificial Intelligence, the Internet of Things and robotics](#), European Commission, February 2020

³ [Big data and artificial intelligence: Principles for the use of algorithms in decision-making processes](#), BaFin, June 2021

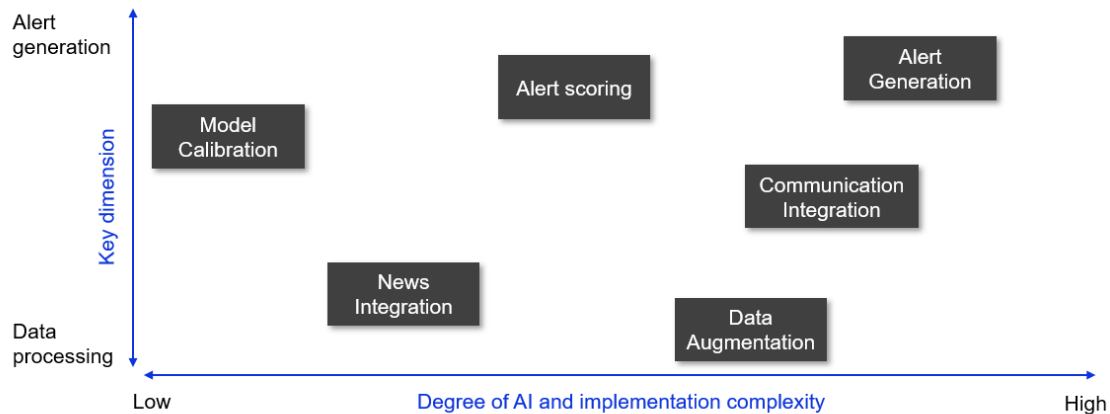


Figure 1. AI use cases in trade surveillance.

2.1 Model Calibration

Rule-based surveillance models typically require certain parameters (e.g., minimum trading volumes, price deviations, etc.) that, when breached, generate an alert. The calibration of these models is the process of finding optimal parameter thresholds and filters that reduce the risk of having false negatives to a minimum and on the other hand limit the number of false positives. The initial conditions specified greatly influence the generated outcomes.

— Challenges with current Model Calibration

Traditionally, a rather static approach based on expert judgement has been used by many institutions, requiring manual adjustments of parameters across products, monitored entities, and time. Due to the high amount of human effort and involvement required, this is a time-consuming process that is not necessarily justified by its outcomes. The challenges can be decomposed as follows:

1. When it comes to the parametrisation of alert rules, there is often a strong reliance on the respective vendor solution. In many cases, the parametrisation methods and functions of the software solution are not appropriate for the specific business model, market, or product.
2. An expert estimate is often used to determine alert rule parameters. Besides the risk of being subjective, such a qualitative approach is often less accurate and difficult to justify both internally and externally.
3. Re-calibration is the process necessary to adjust the initial parameters or filters of an alert rule due to changes in the business model or market environment. In the case where a qualitative approach is used, re-calibration is a time-consuming process that needs to be frequently implemented to derive sound parameters for the prevailing market

conditions, business model, and trading strategies of the financial institution or its clients. In fact, alert rule parameters are already outdated at the day they go in production, because important market and business factors are changing dynamically all the time. Therefore, when it comes to the re-calibration process, financial institutions are forced to position themselves in a trade-off between excessive resource allocation and risk exposure related to false negatives (Figure 2).



Figure 2. Trade-off with re-calibration.

— Analytics-based Model Calibration

An analytics or AI-based approach can help to solve the above trade-off so that financial institutions do not find themselves caught in a high-risk or high-effort dilemma.

There are two major steps that need to be taken:

1. Transition from qualitative to quantitative calibration approach: The first step to a systematic parameter calculation is a transition to a quantitative approach instead of basing parameters on expert estimates. One common way to do this is to determine similar instruments and/or monitored entities with the means of statistics and Machine Learning algorithms (e.g., k-means clustering), and to derive parameter thresholds for each group based on the trading characteristics of the group's constituents.
2. Automation of parameter calculation: Parameter derivation can be (partially) automated to simplify (re-)calibration. Besides implementing parametrisation routines that can handle new data and automatically generate suggestions for optimal parameter thresholds, an automation of required data deliveries and interfaces facilitates a more time efficient calibration process.

— Challenges & Risks of an Analytics-based Model Calibration

- In order to transition to quantitative models for parameter calculation, order, trade, and market data need to be available in a sufficient quantity and quality. Quantitative approaches for parametrisation are only as good as the utilised databases. Especially for smaller banks or buy-side companies, the available order and trade data for specific products will

often not be sufficient to reliably derive parameters with a quantitative approach and market data might be costly to obtain.

- As of 2023, most vendors do not have an integrated solution for automated parameter calculation within their trade surveillance system. Therefore, required calculations need to take place outside of the surveillance tool. This requires additional efforts to avoid inconsistencies or mismatches between data used for actual surveillance and parametrisation.

— **Business Benefits of an Analytics-based Model Calibration**

- By increasing the quality of alert rule parameters, the risk of false negatives can be significantly mitigated, and the number of false positives can be reduced.
- (Re-)calibration efforts can be reduced when a quantitative approach is used as implemented routines for the determination of optimal parameter thresholds can be applied repeatedly and with little additional effort to new data.
- A quantitative approach to calibration allows for an improved justification of the surveillance approach used, both internally and externally, e.g., in communication with the regulator.

2.2 Alert Scoring

Alert scoring refers to the system-based application of statistical methods to the full spectrum of alerts, in order to make a prediction on the relevance of each single alert, either as a numerical or a categorical value. Alert scoring provides a standardised way to classify generated alerts and allows compliance officers to take this additional piece of information into consideration when assessing them.

— **Challenges with current Alert Scorings**

Rule-based surveillance systems usually generate vast amounts of alerts, with a majority being false positives. This leads to many drawbacks. To handle and investigate all the alerts, the investment of a large amount of human capital is required. A high number of alerts might be a contributing factor for alert misclassification, implying a heightened risk of true positives being overlooked which in turn leads to a reduced effectiveness of the surveillance mechanism and a lack of oversight.

Most of today's systems usually offer limited support to compliance staff in allocating their time according to the severeness and importance of alerts. A common approach for alert scoring is to use the magnitude by which the parameters have been breached. However, such approaches only rely on order, trade, and market data directly related to the alert. The information on how a

similar alert has been handled in the past is neglected. This leads at least to the following two disadvantages:

1. As alerts can significantly differ from each other and thus require a deeper look into asset specific characteristics, transaction history, and alert determinants, a compliance officer must investigate each alert individually until he or she has collected all the information required to finally judge the alert. This is a manual and highly time-consuming process.
2. Highly manual alert processing entails operational risks as similar alerts might be classified and documented inconsistently over time.

— **AI-based Alert Scorings**

Machine Learning algorithms score alerts not only based on the narrow spectrum of data that is directly related to the alert, e.g., parameter breaches regarding volumes or prices, but also on how similar alerts have been classified by compliance officers in the past. An AI-based alert scoring is particularly useful when alerts are generated through a traditional rule-based approach. In this case, the scoring can be considered as a second layer that is implemented on top of the regular alert generating process and as such can also optimize the outcome of legacy trade surveillance systems.

At a high level, the process for an AI-based alert scoring in conjunction with a rule-based system could look as follows (Figure 3):

1. Alerts are generated with a traditional rule-based approach.
2. Generated alerts are analysed by an AI algorithm and scored with a probability of being a true or false positive, based on current and historical alert data. Explainable AI can be used to provide the compliance officer with the rationale why the algorithm has come to a specific classification decision.
3. Alerts are manually investigated by compliance staff and classified.
4. Processed alerts are stored in the alert database and will be used for future alert scoring, continuously improving the quality of the scoring through a feedback mechanism.

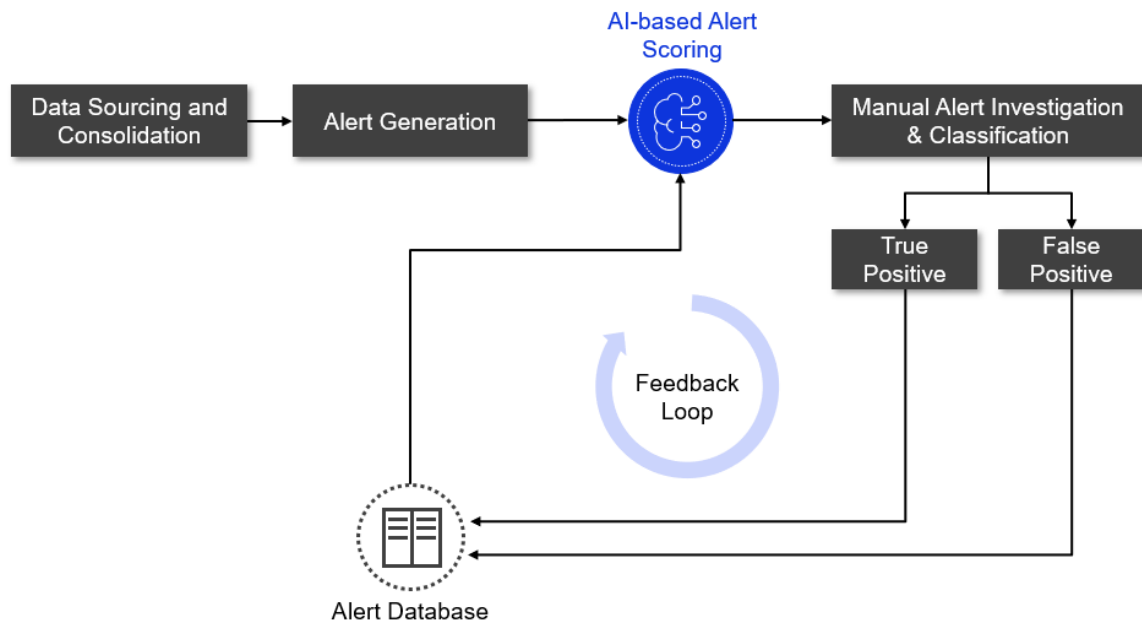


Figure 3. Mechanism for AI-based alert scoring in conjunction with a rule-based system.

In case alerts are directly generated through an AI algorithm instead of a rule-based system, the score would most likely be derived from the alert generation model itself.

— Challenges & Risks of AI-based Alert Scorings

- A sufficient history of processed alerts must be available for reliable scorings. This might be challenging when new rules are implemented, or new products are integrated into the surveillance system.
- A lack of true positives might lead to a pronounced imbalance of alert classes, posing a challenge for predictive modelling as most of the Machine Learning algorithms used for classification were designed around the assumption of an equal number of observations in each class. A possible solution to counteract this issue might be the selection of models that give more weight to the minority class or to use under- and over-sampling techniques, e.g., SMOTE. However, depending on the size and business of the institution deploying surveillance, training the system solely on the narrow spectrum of true positives that have been detected in the own system, if any, might still be problematic. One possible solution could be a consolidated true positive database that is maintained by a central institution or organisation, containing all market abuse patterns that have been detected in the past. Such a database could then be used by other institutions for training purposes and thus could be a first step to more common market standards.

- Additional effort for back testing is required because it must be ensured that a higher alert score indeed indicates a higher probability of a true positive alert.
- Enriching an alert only with a score value might be unsatisfactory from the perspective of compliance officers as the determinants of the score might be opaque. Explainable AI can be a way out but comes with additional implementation efforts.
- Too much reliance on the alert score bears the risk that compliance officers investigate alerts with a low score less thoroughly.

— **Business Benefits of AI-based Alert Scorings**

- Using the additional piece of information provided by the alert score, compliance officers can focus more thoroughly on alerts that might be true positives.
- Explainable AI can support compliance officers during alert investigation and classification by identifying the most relevant information. This can also provide insightful information that can feed-back to the parametrisation process and thus can be used for false positive reduction.
- An AI-based alert scoring approach is a collaborative process between humans and AI. The alert scoring serves as one additional assistance during alert handling while the final decision and oversight remains with the compliance officer. Such an approach is designed to ensure a high level of acceptance, both internally and from regulators. As it involves a medium degree of AI-usage, it can be considered a soft transition towards an entirely AI-based solution in the future.

2.3 Alert Generation

Alert generation entails the system-based process of identifying and reporting potential cases of market abuse based on order or trade characteristics such as volumes, prices, number of events, etc. and market data.

— **Challenges with current Alert Generation**

Traditional alert generation models use a rather static approach with pre-defined thresholds. These thresholds are often based on expert judgment and when certain thresholds are breached an alert is generated. Due to the static approach, there is usually a superfluous number of alerts that requires a lot of attention from compliance officers.

Furthermore, alert generation processes are usually implemented in an underlying daily framework and therefore provide limited control for historical

information. Seldomly historical alerts are matched with contemporaneous ones based on the similarity of their key characteristics.

— AI-based Alert Generation

To overcome the above challenges, a process that incorporates historical information and projects it into daily data, while controlling for the differences in past and present trends is deemed necessary. Machine Learning based surveillance processes are fully integrated models that account for the above issues, while offering the potential to outperform traditional processes in terms of efficiency, speed, and flexibility.

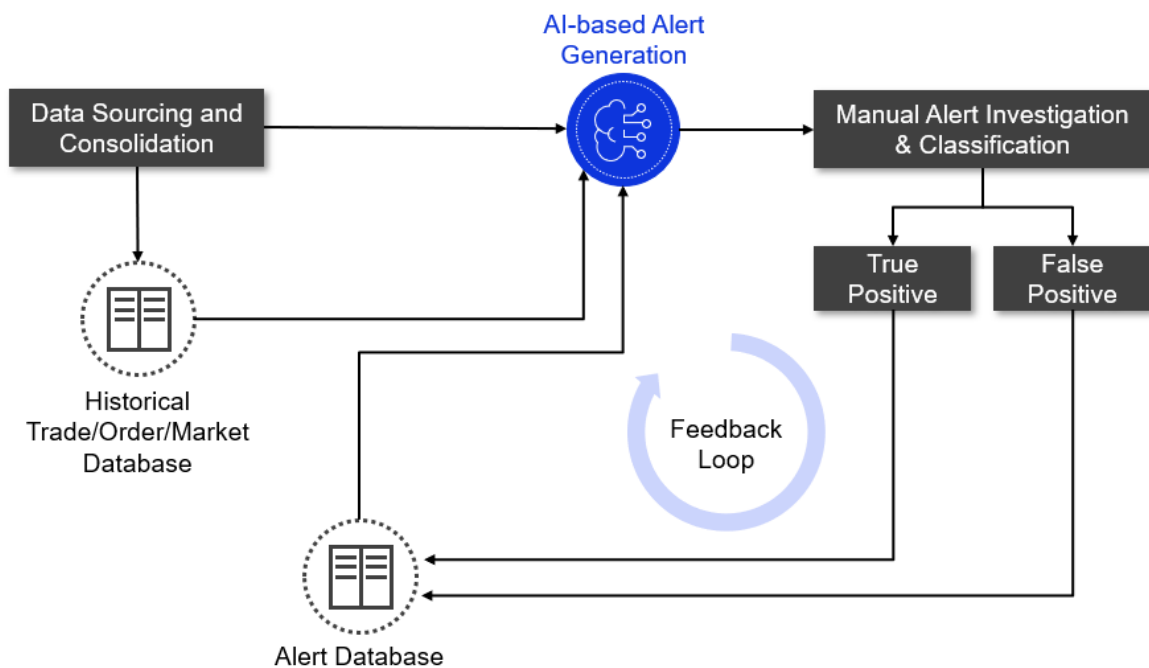


Figure 4. Mechanism for AI-based alert generation.

An exemplary mechanism for an AI-based alert generation is shown in Figure 4 and can be characterised as follows:

- AI-based alert generation is a bottom-up process that analyses a series of factors as well as historical and contemporaneous information to generate alerts.
- An AI-based alert generation process can for example identify a transaction, then compare it with similar historical transactions and consequently generate an alert based on confounding factors between actual transaction(s) and those used for training purposes.
- AI-based models allow to focus on broad underlying market abuse patterns such as false signalling, price manipulation, or insider dealing instead of covering several narrowly defined scenarios only.

- To minimise false positive alerts, AI-based systems can easily move from linear to non-linear, higher dimensional approaches.
- Adequate data inputs allow for the identification of the driving factors between true and false positives.
- AI-based solutions also highlight the causal reasons for alert generation, making the case for concise identification and investigation of market abuse cases. Learning from past experiences, i.e., historical data, the algorithms can disentangle and weigh the alert factors and carry this information in current and future applications. This in turn increases transparency and reduces the cost of compliance.

— **Challenges & Risks of an AI-based Alert Generation**

- In general, a sufficiently large data set must be available as a basis for training. This might pose a challenge for institutions that are currently in the process of integrating new asset classes or products into their surveillance system.
- Compared to the implementation of a rule-based system, an even more gradual implementation process is required which allows for proper back-testing, alert comparison, and interpretation by compliance officers.
- Whereas completely AI-based processes do not require human intervention, human monitoring should not be considered redundant. As the involved processes in alert generation need to be monitored and require a high level of diligence due to regulatory requirements, changes in the system behaviour should be frequently reported to ensure transparency and clarity.
- The advanced methodologies involved might weaken the explainability of alert outcomes. In response, there should be a clear and detailed documentation of each AI-based solution. Explainable AI can also contribute to increased transparency.
- Similarly with the aforementioned AI-based approaches, these processes need to be well taught to staff and should not be blindly relied upon.

— **Business Benefits of an AI-based Alert Generation**

- AI-based algorithms integrate historical information analysis, thus allowing for adaptive learning from past trading. This information provides evidence that is beneficial not only to compliance officers but also to regulators.
- The number of false positives can be reduced, as with increasing experience, the system learns to suppress alerts that exhibit the same patterns as alerts that have been classified as false positives in the past.
- An AI-based alert generation approach allows for an entity-centric instead of an event-centric view. This means that the focus of surveillance can be

set on entity-specific abnormal behaviour instead of bundling several entities together and parameterising group thresholds.

- Newly evolving abusive patterns might be detected without explicitly programming a new rule, therefore reducing the cost of maintenance, and reducing the risk of false negatives.
- The identification of superior patterns instead of triggering alerts as soon as specific rule-based conditions are met allows for more meaningful alerts and less double alert triggering.
- Cases of collusion can be identified more effectively, as dependencies and correlations between data can be taken into consideration.

2.4 Integration of Unstructured News Data

Unstructured data (e.g., text or audio) typically make up for the majority of all available business data. At the same time, they are the most difficult to process and to analyse. However, this kind of data provides valuable information for trade surveillance purposes. A typical example would be news data that can be useful for alert generation, e.g., to detect cases of insider trading. In addition, it can provide insights for investigation purposes as it allows the evaluation of whether certain transactions and market movements are related to external factors.

— Challenges related to Data Integration

Unstructured data can vary greatly in their format, source, and layout. Consequently, for a seamless integration into a trade surveillance tool, the data must be made available in a structured form. For the specific case of news data this might mean that it must be mapped to instruments that are (potentially) affected by the news event. While for some asset classes (e.g., equities) structured news data is often readily available from data vendors, in most asset classes, such as FX or interest rates, mostly unstructured news data can be obtained which cannot be directly processed in a surveillance tool. Moreover, information about the expected direction in which a price will change based on a news event can be beneficial – but requires extra effort in the analyses.

— AI-based News Data Integration

With the goal to transform the initially unstructured news data into formatted, usable information that can further be processed by trade surveillance systems, Natural Language Processing (NLP) techniques can be applied. The following three steps are relevant:

- 1 Data extraction: NLP allows for text recognition of the relevant information and extraction into a standardised format for further analysis.
- 2 Classification: The derived formatted information can be further processed and classified into the respective asset class and instrument type.

- 3 Sentiment analysis: In addition to the information directly included in a news article, such as timestamp, source, and involved entities, NLP, through the use of keywords and n-grams, allows for an analysis of the soft facts included, such as the tone and language used. This way, additional information about the expected direction of a price movement after the news information has been released can be attained.

— **Challenges & Risks of an AI-based News Data Integration**

- The (potential) universe of products that might be affected by a news event can be large. Therefore, only a rough mapping of the news data to instruments might be possible.
- Depending on the context of the news event, sentiment information might be difficult to obtain or unreliable.

— **Business Benefits of an AI-based News Data Integration**

- News data can be beneficial for alert generation, manual investigations, and parametrisation through the following mechanisms:
 - The risk of false negatives can be mitigated because attempted market abuse will be easier to detect. A typical example where this is the case would be a perpetrator who is building up a position based on insider information that is expected to cause a price change after the information is released to the market. In case it turns out that the information has already been priced into the market price or the price effect is much smaller than expected, this is considered an attempt of market abuse that could not have been detected when a rule uses only simple information such as actual price movement.
 - False positives can be reduced as more (relevant) information is used in the alert generating process.
 - A more effective parametrisation can be realised as news data allows to filter noise before statistical analyses are applied to the data.
- When compared to a manual search, a direct integration of news data into the trade surveillance system, including a correct mapping to potentially affected instruments, can be time and cost efficient.

Besides the above use case of structuring news data with the goal of a direct integration into trade surveillance systems using NLP, similar techniques will be beneficial for other purposes. Particularly in the related area of communications surveillance, NLP is of outstanding importance. Although communication surveillance is beyond the scope of this whitepaper, it should be mentioned that the transfer of communications surveillance data to a trade surveillance system

can be supported by NLP, e.g., to facilitate the mapping of communications alerts to trade surveillance alerts.

2.5 Integration of Communications Surveillance

Communications surveillance, as described in the MAR, refers to the systematic monitoring and analysis of communications within financial entities and markets. The aim is to detect and prevent market abuse by scrutinizing various communication channels used by financial professionals. This surveillance involves the examination of messages, emails, phone calls, and other forms of communication to identify potentially suspicious activities or behaviours that may violate regulatory standards.

— Challenges related to communications surveillance

Acknowledging the challenges and potential conflicts with information privacy regulations, particularly the General Data Protection Regulation (GDPR) in the European jurisdiction, we delve into communication surveillance without constraints. A primary obstacle lies in the acquisition of high-quality, reliable data containing pertinent information for effective surveillance. Communication among salespeople, traders, and counterparts spans various channels, including emails, phone calls, and messaging platforms, transcending regions and languages. Traders frequently shift languages, use abbreviations, trade-specific terms, and employ irony within a single communication thread or message.

Post pre-processing tasks, such as structuring and normalizing communication data, the critical step involves distinguishing incriminating information from non-incriminating content. This categorization process demands accuracy and reliability. Many vendor solutions still rely on simplistic dictionaries, checking communications against predefined words. Such approaches pose challenges to existing solutions due to incomplete dictionaries and linguistic peculiarities such as humour or irony.

— AI-based communications surveillance

Artificial Intelligence holds vast potential in communications surveillance, particularly in uncovering behaviour that may elude classical rules or human investigators. The typical use cases include:

1. Pattern Identification: AI excels at identifying patterns that may evade traditional rules or human scrutiny. This includes the detection of repeated phrases, specific keywords, or other linguistic anomalies.
2. Sentiment Analysis: Utilizing sentiment analysis, AI can discern the mood and tone within text communications. This capability is instrumental in

pinpointing potentially suspicious or aberrant behaviours by capturing shifts in the emotional content of communication.

3. Metadata Analysis: AI can conduct a thorough analysis of metadata associated with communication. This encompasses examining anomalies in texting or calling frequencies or detecting unusual interactions or fluctuations in the number of meetings between specific individuals.

By leveraging these capabilities, AI enhances the precision and efficiency of communications surveillance, providing a robust mechanism for uncovering nuanced insights and potential risks that may elude conventional methods.

— **Challenges & Risks of an AI-based communications surveillance**

- Striking a balance between the necessity for surveillance and the safeguarding of individuals' privacy rights is a paramount challenge. Ensuring compliance with stringent privacy laws, including GDPR, is imperative, necessitating the implementation of robust measures to safeguard sensitive information.
- Conducting communications surveillance often involves monitoring conversations across a diverse array of languages and dialects. To ensure accurate analysis, systems need to demonstrate adaptability to various linguistic variations. An innovative solution could involve creating a company-specific slang dictionary, a task made more feasible by harnessing the capabilities of Generative AI.
- The escalating adoption of encryption technologies introduces complexities for surveillance efforts. Achieving the delicate balance of decrypting and analysing encrypted messages is essential, demanding meticulous adherence to privacy laws while upholding the integrity of surveillance practices.

— **Business Benefits of an AI-based communications surveillance**

- AI demonstrates exceptional proficiency in discerning intricate patterns, anomalies, and subtle deviations within communication data. This capability enhances the precision and reliability of detecting potential risks or suspicious activities.
- AI facilitates real-time monitoring of communication channels, empowering organizations to swiftly address emerging risks and potential compliance violations.
- Text analytics plays a pivotal role in establishing connections between communications and associated order and trade data across diverse market participants. This interconnected analysis provides a holistic view of activities within the market landscape.

- Streamlining the investigation process is achievable by integrating communication data with a language model that enables querying the text. This connection enhances the efficiency of investigations by allowing investigators to pose targeted questions based on the content of the communication.

2.6 Data Augmentation

Data augmentation refers to the process of creating additional data by using existing data and modifying it slightly.

— Challenges with current surveillance data

Data quality and availability are among the most important factors to implement an effective trade surveillance system.

Regarding data quality, trade surveillance vendors typically implement data management features and sophisticated processes for extracting, transferring, and loading data as part of an ETL process. Additionally, most financial institutions have developed processes to identify and correct data issues.

However, regarding availability and the lack thereof, solutions are less tangible. A common challenge in trade surveillance is data scarcity. This poses barriers to effectively derive alert rule parameters with a quantitative model. Especially data for true positive cases are a rare occasion when compared with the massive amounts of false positives. This leads to difficulties at configuring and especially calibrating surveillance models in a quantitative manner, because this process requires a certain amount of data to be reliable. On top of that, financial institutions with that lack in terms of data quantity might not be able to implement most state-of-the-art features such as customising surveillance models on entity level (e.g., trading desk or client) or analytics of individual behaviour.

As a conclusion, the availability of sufficient data is a key consideration for alert quality. Especially for financial institutions that want to utilise AI in the future to identify more complex market abuse patterns or to reduce false positives, it is important to know that there is a positive relationship between the complexity of the used AI models and the required amount of training data.

— AI-based Data Augmentation

In order to address the issue of insufficient data, the usage of synthetic data can be an essential tool for developing and validating state-of-the-art quantitative models.

When it comes to trade surveillance, the AI-based augmentation of existing data with realistic artificial data can increase the firms' capabilities to effectively calibrate alert rule parameters and configure alert generation. Especially the

generation of synthetic data for true positives can be helpful, due to the scarcity of real world true positive examples.

Creating synthetic data involves utilizing various models, with two standout deep learning approaches being generative adversarial networks (GAN) and variational autoencoders (VAE). Both models operate on the premise of learning from the joint probability distribution present in real data samples, ultimately generating a new dataset with a matching distribution. This ensures that synthetic data maintains the same statistical properties as the original. These models excel at constructing balanced datasets, distributing observations evenly across each class of interest, such as false positives and true positives.

A cutting-edge advancement in synthetic data creation involves leveraging Generative AI. This approach allows for the articulation of a desired market abuse pattern in natural language, empowering AI to generate synthetic order and trade data that accurately reflects this pattern.

As of October 2023, no vendors seem to incorporate a data augmentation module into their trade surveillance solutions. Nevertheless, this task can be efficiently executed in a distinct system, and the generated synthetic data can be seamlessly integrated into the trade surveillance solution for training or validation. Numerous open-source tools are accessible for both deep learning models, GAN and VAE, with many coded in Python or easily integratable into it.

— **Challenges & Risks of an AI-based Data Augmentation**

- For cases where there is no data available at all, the generation of synthetic data becomes more difficult. However, if it is known what the data should look like with regards to data types, format, and underlying abusive pattern, it is possible to generate data from scratch.
- The calibration and configuration of trade surveillance systems must remain explainable and auditable when utilising synthetic data.

— **Business Benefits of an AI-based Data Augmentation**

- Machine Learning algorithms are currently being used for producing realistic synthetic data in several industries and can create competitive advantages.
- Synthetic data safeguards client privacy by enabling the creation of fully synthesized datasets, offering a robust privacy solution.
- Synthetic data can tremendously help when it comes to parametrisation of alert rules or training an AI-based trade surveillance system which leads to improved overall alert quality.
- Utilizing synthetic data facilitates the attainment of a balanced dataset through effective over- or under-sampling techniques.

- Synthetic data does not lack diversity of asset classes or observation count so that financial institutions do not have to rely on expert estimates for the configuration of their alert rules. Besides ensuring better auditability, this also can lead to significant savings in the (re-)parametrisation process.

3 Generative AI: Redefining Trade Surveillance Dynamics

Generative AI (GenAI) refers to a category of artificial intelligence that involves computers creating or generating new content, such as text, images, program code or other forms of data. Instead of relying solely on explicit programming, Generative AI models are trained on large datasets and learn patterns and structures within this data. This allows them to generate novel and coherent output that often mimics human-created content. As such, its application also extends to various facets of trade surveillance.

— **Generative AI Models and their current and future possibilities for trade surveillance**

As of today, numerous firms express dissatisfaction with off-the-shelf vendor surveillance models, citing inflexibility or misalignment with their business models. While many vendors offer their customers the option to program custom models within their solution, this capability often sees limited use due to a lack of expertise or capacity in many Compliance departments. Generative AI addresses this gap by enabling code generation in nearly any prevalent programming language based on a verbal description of desired models and outcomes. This empowers companies to independently develop, implement, and extend trade surveillance with no or minimal coding experience. For those unsatisfied with current models, GenAI provides an excellent opportunity to enhance their surveillance landscape significantly.

This advantage extends to the creation of synthetic data. When calibrated correctly, GenAI can directly synthesize data or provide coding examples for synthesizing data with custom rules, distributions, and variables. For example, users can instruct GenAI to generate targeted abusive behaviours such as "Pump and Dump" or "Layering & Spoofing," mirrored in synthetic order and trade data, seamlessly aligning with the company's data model.

Moreover, harnessing GenAI for communications surveillance brings advantages in crafting company or department-specific language dictionaries. These features elevate the efficiency, security, and overall effectiveness of communications surveillance. Additionally, after alert generation, GenAI aids investigations by swiftly processing communications data. With prompts from investigators, large language models can assist in presenting potentially incriminating text, telephone calls, or trading patterns. This allows investigators to focus on the substance of the investigation rather than dedicating time to scouring for relevant text, emails, or phone calls.

— Challenges in the usage of Generative AI for Trade Surveillance and how to overcome them

Effectively implementing the outlined use cases encounters a substantial hurdle in precisely configuring and applying the models. The clarity of task definition plays a pivotal role, as the current models may fall short of delivering a satisfactory response when faced with ambiguous tasks or questions. Even with accurate phrasing, suboptimal solutions may arise, necessitating additional adjustments. This challenge is particularly pronounced in critical tasks like compliance topics. Addressing this issue involves ensuring optimal tuning and calibration of models through rigorous testing and validation processes. Regular updates and fine-tuning based on real-world performance data are crucial for achieving more accurate and reliable outcomes.

In addition, establishing trust in the solutions provided by these models remains a significant challenge. Currently perceived as opaque black boxes, these models generate intricate outputs that pose challenges for human interpretation. To overcome this hurdle, investing efforts in enhancing the interpretability and explainability of these models can significantly increase transparency for human users.

Furthermore, fostering collaboration between AI systems and human experts is essential. Integrating human expertise into the loop enables professionals to validate, interpret, and refine results generated by Generative AI. This synergy ensures that the investigation process remains comprehensible and aligned with human understanding.

In addition, providing comprehensive education and training on low-code or no-code platforms can empower individuals across departments to actively participate in Generative AI endeavours. This democratization of coding skills broadens user engagement, allowing a diverse range of individuals to pose pertinent questions and actively contribute to the development and maintenance of AI models.

4 Opportunities and Risks arising from the application of Machine Learning in Trade Surveillance

Opportunities		Risks	
Consistent true positive identification	Actual market abuse cases will be identified more effectively and efficiently. Thus, financial institutions can reduce the risk of regulatory fines and loss of reputation.	Black box	With a fully integrated AI approach, it could be difficult to understand the decision-making process of AI when generating an alert or calculating a parameter for an alert rule.
Less false positives	Less false positive alerts are generated, which reduces the required resources for the expert analysis.	Transition Barriers	The transition should be tested, monitored, and the integral parts of the new implemented systems should be handled with confidentiality and kept under scrutiny.
Reduction of market abuse incentives	An industry-wide application of AI-based trade surveillance will lead to better identification of actual cases of market abuse and thus reduce the incentives for misconduct among market participants.	Monitoring	A fully integrated AI-based surveillance system should not be left unattended as it might develop some unintended behaviour.
Reduction of set-up costs	AI-based systems reduce the initial effort required to set up traditional surveillance systems, e.g., regarding model calibration and parametrisation.	Misclassified Alerts	With the current state of AI, these systems are not supposed to completely substitute human labour. Instead, they should supplement and assist the efforts of compliance officers.
Reduction of maintenance costs	As AI-based systems are highly automated and dynamically adopt to changing environments, maintenance costs will be reduced.	Overfitting	A more elaborate model using a specific sample of orders, trades or alerts can be trained to predict historical market abuse cases very accurately but can perform poorly when confronted with new patterns.
Standardisation	Automated systems based on similar methodologies can produce a more standardised alert generation process and reporting, allowing for more transparent transmission of information between financial institutions and regulators.	Data Scarcity	AI-based trade surveillance systems require an adequate volume of data for training that might only build up over time.
Less room for subjectivity	Automated, AI-based systems reduce human involvement, which in turn lowers the impact of human biases that can prevail when setting parameter thresholds or assessing alerts.	Manipulation	AI-based trade surveillance systems can still be liable to malicious training and data poisoning. A set of rules should be set to ensure compliance and transparency.

5 Conclusion

Following recent technological advancements in AI and Machine Learning, trade surveillance is currently undergoing a digital transformation. Both dissatisfaction with legacy systems and increased regulatory pressure call for an improvement of tool-based trade surveillance. Besides a more effective identification of market abuse practices, increased efficiency and lower cost are the main drivers of those developments. AI-based solutions can improve data availability, structure, and anonymity in sensitive customer data. However, such a transition is not as straightforward as it might sound. Financial institutions can choose from a large universe of vendor solutions offered in the market while there are no specific guidelines available on how to approach and implement modern AI-based trade surveillance systems.

AI-based surveillance systems present a viable alternative to legacy systems, enhancing the detection of malicious market practices and improving efficiency concurrently. Another approach involves optimizing legacy systems by seamlessly integrating new AI-based tools into existing processes and infrastructure. Following the transition from legacy systems to AI-based surveillance, the potential for additional optimization arises through the incorporation of Generative AI. This ongoing development is linked to shorter time-to-market and lower costs when compared to a greenfield implementation.

Considering the opportunities, applications, and associated risks of implementing modern technologies, the effective integration of AI-based trade surveillance systems can not only decrease maintenance, monitoring, and regulatory costs but also diminish incentives for market abuse. This, in turn, enhances the resilience and integrity of financial markets. In this context, a smooth transition to state-of-the-art surveillance systems is crucial and requires a clear understanding of the underlying technical processes, thorough personnel training, and transparent communication of information among system providers, applicants, and regulators.

– **Author**



Philipp Faulstich, CFA
Manager


+49 69 971485-416
philipp.faulstich@l-p-a.com

– **Contact**



Christian Behm
Partner LPA Consulting

+49 69 971485-437
christian.behm@l-p-a.com



Founded in 1999 by Stefan Lucht and Roland Probst to offer solutions as the finance industry evolved, LPA was built on the needs of its clients. Today the company has over 400 employees in 12 cities worldwide, delivering cutting-edge consulting through its specialist teams, and technology through software products.

Defy the processes of the past. Deploy innovation in a way that suits you. Run to the forefront of global innovation.

For more information, please visit www.l-p-a.com.