AI-Enhanced Sanctions Screening A risk-controlled predictive compliance Framework

Pratik Chawande

Independent Research
Dallas, Texas
chawandepratik00@gmail.com

Abstract:
The current global financial system is under strain due to the growing rate and complexity of increasing sanctions regimes, where the screening systems of the past, based on rules, produce false positives in unsustainable levels of over 95%. Such operational overload burns compliance resources, causes alert fatigue and blurs indications of real danger caused by advanced evasion strategy. It is in this paper that I suggest a new, risk-managed framework in which artificial intelligence (AI) and machine learning (ML) are merged to transform the sanctions compliance paradigm into a reactivity pattern-matching system into a proactive, predictive approach to risk. The multi-layered architecture of the framework takes advantage of advanced natural language processing (NLP) to resolve entities, network analytics to map relationships and a collection of supervised and unsupervised machine learning models to contextually score risk. Most importantly, it includes explainable AI (XAI) to enable transparency in decisions, a well-developed model risk management layer of governance, and a non-alterable audit trail to achieve regulatory traceability. This prospective framework offers a scaled solution to the problem of false positives to a significantly lesser degree, elevated true positive detection rates, and a defensible, efficient and adaptive sanctions screening program to meet contemporary regulatory standards by making use of the critical triad of accuracy, explainability and auditability.

Keywords: AI-Powered Sanctions Screening, Reduction of false positive, Explainable AI, Predictive Compliance, Model Risk Management, Financial Crime Technology, Regulatory Technology.

1. INTRODUCTION

The effectiveness of the implementation of economic and trade sanctions forms the basis of the integrity of the international financial system. Financial institutions play the most significant role as gatekeepers; they are required to vet various transactions of millions of daily transactions against dynamically growing lists provided by the Office of Foreign Assets Control (OFAC), the United Nations, and the European Union among others. The difficulty has taken a point of critical incidence. The traditional screening paradigm, based on deterministic, rule-based systems, which operate on lexical matching algorithms, is essentially ill-prepared to the modern era. Such systems are not contextually intelligent, which results in an epidemic of false positives, which is usually 95-99% of all alerts [1]. This does not only inflict disabling operation costs of about 25-150 per alert investigation, but also creates alert fatigue to the analysts, thus raising the risk of overlooking real matches (false negatives) against the noise.

At the same time, attackers use more advanced evasion techniques, such as obfuscating names, multi-level corporate hierarchies, and structure of transactions, that are subject to easy circumventing of simple string-matching reasoning. There are also changes in regulatory expectations that focus on the necessity of risk-based programs that are more efficient in allocating resources to the areas that have the highest risks- a concept that is contrary to the concept of systems that squander the majority of resources on low-risk noise.
This is the gap in technology and operation that is imperative in this paper. It announces a system of AI-enhanced sanctions screening that is aimed at changing compliance into a transactional checklist that is

reactive into an AI-driven, risk-controlled, and intelligent process. The suggested architecture does not just represent a model but a developed system that includes the data fusion, advanced analytics, and human-focused decision support with strict governance. It discusses the primary pain points of the industry directly such as reducing the false positive burden, identifying complex evasion typologies, and deliver explainability and auditability to allow a regulatory acceptance [2]. The sections below describe the architectural elements, theoretical performance advantages, implementation aspects and the necessary governance structures that enable this framework to be innovative and at the same time deployable in the tight-knit niche of financial regulation.

2. THE AI IN SANCTIONS SCREENING IS IN DIRE NEED OF IMPROVEMENT

2.1. The Growing Nuisance of the Screening Problem

Financial institutions now have found themselves in a regime of sanctions that is of spectacular complexity. The Office of Foreign Assets Control (OFAC) alone upholds more than a dozen sanctions programs comprising of about 12000 (12,000) entries, and world institutions are required to simultaneously oversee United Nations, European Union, and hundreds of national lists of restricted parties comprising of hundreds of thousands of parties. The number of designations in this volume is expanding exponentially over 1,700 new designations were introduced by OFAC in 2022 alone, which is nearly 35 percent more than the prior year.

Besides list volume, the screening burden is also seen to include:

- Geographic sanctions involve screening based on the location.
- Sectoral sanctions that require code analysis in the industry.
- Ship and ship registries requiring special identification.
- Real time dynamic lists to be updated and implemented.

Conventional systems are hampered by this size and complexity especially when sanctions shift away from countrywide embargoes to specific and narrow sanctions against individuals, organizations and particular industries.

2.2. Measuring the False Positive Crisis

False positive issue occurs in various dimensions that are operationally and financially serious:

Volume Analysis of alerts: an average mid-sized international bank that manages 15 million transactions per day will typically generate 75,000-150,000 screening alerts with only 300-750 being real matches. This 0.4-0.5% true positive indicates that compliance analysts are wasting 99.5 percent of their time on investigating non-issues.

Resource Allocation Impact: Every alert takes into consideration 15-40 minutes of an investigation that includes the investigation of several systems, external research, and documentation. This translates to 1.5-4.2 million in monthly labor costs of false positive investigation by itself at an average fully-loaded cost of 85 per analyst hour.

Quality Degradation Risks: Alert environments with high volume and low yield are characterized by a number of risk factors [3].

Alert Fatigue: Researchers become biased in their thinking and shortcut procedures.

Threshold Manipulation: When there is pressure to decrease volumes, appropriate matching rule adjustments are not done.

False Detections: Alarms are lost in noise, especially when peak values occur.
Repetitive and Unrewarding Work: Staff attrition is enhanced by repetitive, unrewarding work.

Effectiveness of Compliance Programs: Regulatory expectations- Our expectations of compliance programs keep increasing, Framework of Overseeing Sanctions by the Office of Foreign Assets Control highlighting the concept of risk-based programs developed based on risk profile of an organization [4]. The classic high-false-positive systems are not able to illustrate the risk-based approaches with the majority of resources working on non-risks.

Figure no 1.1: The Sanction Screening Alert Overload

2.3. Advancing Evasion Strategies Needing Sophisticated Detection

The bypassing of traditional screening is accomplished by highly advanced ways of modern sanctions evasion

Name Obfuscation Strategies

- Intentional cases of misspelling, transliteration differences (Putin/Putyn/Putyin).
- Inconsistent use of middle names, initials and honorives.
- Use of non-Latin scripts whose transliterations were more than one.
- Formal and informal name usage in various transactions.

Structural Evasion Methods:

- Location to location multi-layering of corporations.
- Nominee directors Shell companies.
- Tactical application of more lenient jurisdictions of less rigorous screening.
- Misuse of international alignment and enforcement of international lists.

Transaction-Based Techniques:

- Organization of payments that are lower than threshold.
- Alternative value transfer procedures (cryptocurrency, trade-based) are used.
- Stratified correspondence banking relationships.
- Time zone abuse in the screening updates.

These methods make lexical matching of the conventional approach more and more inefficient, leaving false negatives that are a risk of great compliance [5]. Contextual, intelligent screening was required as never before.

3. POSTULATED ARCHITECTURE FRAMEWORK

The AI-based sanctions screening model is the system that redesigned compliance using integrated machine intelligence and strict governance checkpoints. It is an architecture comprising of five linked layers which operate in a synergistic manner to enhance accurate detection and operational efficiency.

3.1. Data Layer and Feature Engineering Layer

This base layer converts various data streams to the structured features to analyse:

3.1.1. Multi-Source Data

The framework consumes and integrates data of various internal and external sources:

Sanctions Lists: List of sanctions in various jurisdictions, received by OFAC, UN, EU, and HM Treasury with metadata such as date of designation, purpose and pseudonyms.

Customer Data: KYC information, customer risk rating, transaction history and behavioral pattern.
Transaction Data: SWIFT MT/MX/SEPA/domestic format payment messages, remittance information, and supplementary documentation.
External Intelligence: Unfavorable media, politically exposed individual databases, corporate registry, and shipping/aviation registry.
Network Data: Relations Counterparty relationships, ownerships and previous patterns of transactions.

3.1.2. Entity Resolution Advanced Natural Language Processing

The classical fuzzy matching models the names as character sequences without semantic knowledge. The pipeline that is established by the framework is an advanced NLP pipeline:
Transliteration Normalization: It is based on Unicode normalization and context-sensitive models of transliteration with linguistic origin [6]. As an example, the Arabic word "محمد" can be translated in many ways with several valid transliterations (Mohammed, Muhammad, Mohamed) and the system considers them to be of the same meaning.
Named Entity Recognition and Classification: Recognizes and classifies types of entities in unstructured text with transformer models trained on financial documents. The system identifies names of persons and company names, name of vessel and geographic location in payment messages.
Contextual Embedding Generation: It uses BERT-based models that have been specifically trained on financial corpus data to generate semantics. These embeddings are mini-understandings, as opposed to character similarity- the knowledge that the use of the word Apple in the context of technology is not the same as the use of the word apple in the context of farming.
Cultural/Linguistic Adaptation: Embodies an understanding of naming patterns in different cultures, such as the use of patronymics in Russian names, the use of family names in East Asian names, and the use of honorifics in the Middle East.

3.1.3. Network Graph Construction

The framework creates dynamic knowledge graphs that are relationships between objects:
Ownership Networks: Multi-hop ownership networks obtained through corporate registries, KYC data and self-declarations, weighted by the level of confidence and verification.
Transaction Networks: Past payment trends in terms of frequency, amount, jurisdiction and use of transactions between parties.
Social and Professional Networks: The connections made based on a similarity of addresses, director roles, educational backgrounds, and work experience with the aid of the graph inference algorithms.
Temporal Network Analysis: Change of relationship over time and detecting abrupt shifts or abnormal connection structures.

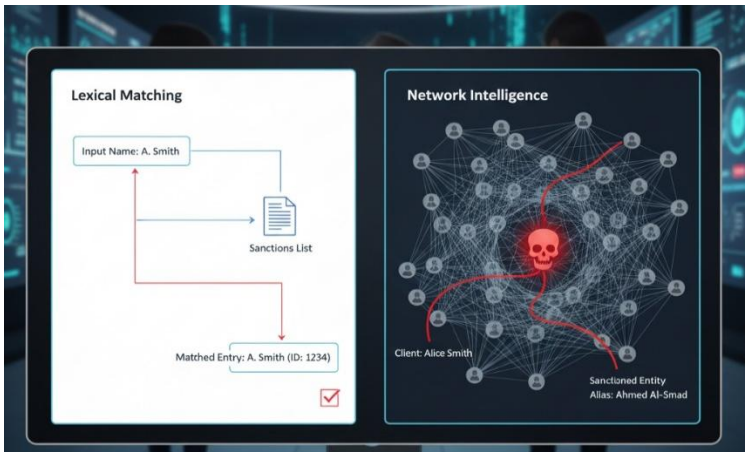

Figure 1.2: Network Intelligence

3.1.4. End to End Feature Engineering

The system builds a rich feature set of every screening process over a variety of dimensions:

Match Quality Features:

- Lexical similarity scores (Jaro-Winkler, Levenshtein)
- Phonetic similarity with language specific algorithm.
- Contextual Embedding Semantic similarity.
- Historical correlation probability of alias matches.

Contextual Risk Features:

- Rating and history of customer risk.
- Jurisdiction geographic risk scores.
- Risk factors related to products/ services.
- Trends and performance of alert signals in the past.

Action and Time:

- Abnormality in common patterns of transactions.
- Time lag with regard to lists of sanctions.
- Geopolitical situation and the correlation of events.
- Cyclical pattern analysis Seasonal pattern analysis

Network Features:

- Connection strength and degree centrality.
- Transaction networks centrality of betweenness.
- Clustering and membership of communities.
- Path analysis to identified high-risk entities.

3.2. Machine Learning Model Ensemble

The model has used a diversified ensemble method that involves a combination of several specialized models:

3.2.1. Managed Classification Models

These models are based on past performances of alert investigation:

Gradient Boosting Machines (XGBoost/LightGBM): Major classification models that deal with the imbalanced data with weighted training and with the sophisticated forms of regularization.

Deep Neural Networks: Multilevel models that learn non-linear and complex features in high-dimensional features spaces especially in integrating various types of data [7].

Addressing Class Imbalance: The imbalance is extreme (the number of true positives is less than half a percent), so the system has special methods:

- With special care and validation Synthetic minority oversampling (SMOTE)
- Asymmetric misclassification costs Cost-sensitive learning.
- Ensemble techniques that are optimized in terms of imbalanced data.

3.2.2. Anomaly Detection with No Supervision

These models define atypical trends without historic labels:

Isolation Forests: The isolation forests are used to efficiently detect anomalies in high dimensional feature spaces with random selection of features.

Autoencoder Networks: Architects Learn Compressed representations of normal patterns In encoder-decoder networks, learn normal patterns and flag anomalous events where the reconstruction error is large.

One-Class Support Vector Machines: The boundaries are around normal observations, and anything that falls outside the boundaries is an anomaly.

Clustering-Based Anomaly Detection: Find the points that do not fit a cluster or belong to small and sparse clusters.

3.2.3. Network Analysis Models

Graphic algorithms determine patterns of relationships:

Graph Neural Networks: Train node representations that use a combination of node features and network topology via message-passing.

Community Detection Algorithms: Find groups of nodes that are densely connected and use the Louvien modularity optimization or Leiden as the techniques.

Risk Propagation Models: This software will apply label propagation and belief propagation algorithms which propagate risk scores across networks using their strength and directionality of connection.

Structural Similarity Analysis: Find subgraphs to provide similarity to known evasion patterns through graph embedding and similarity (Orekha, 2025).

3.2.4. Aggregation and Calibration of Ensembles

Stacking of individual model outputs with meta-learners that optimize on the aggregate is used to combine them. The ensemble is extensively calibrated in terms of probability:

Calibration Methods: Platt scaling and isotonic regression are used to make risk scores look like true probability not arbitrary scores.

Uncertainty Quantification: Bayesian methods and ensemble variance give confidence intervals and uncertainty estimates of every prediction.

Threshold Optimization: There are various thresholds which are set in regards to various categories of risks on the basis of cost benefit analysis and regulatory demands.

3.3. Explainable Artificial Intelligence and Decision Transparency Layer

This is a critical layer which makes sure that model outputs are readable and that decisions can be traced:

3.3.1. Importance and Attribution of Features

In case of every alert, the system measures the contribution of each characteristic to the original risk score:

SHAP (Shapley Additive exPlanations): It gives sophisticated, theoretically-supported values of feature importance that add up to the model deviation to baseline.

Local Interpretable Model-agnostic Explanations (LIME): Generates local faithful approximations of the complex models by the use of interpretable surrogate models.

Integrated Gradients: In the case of neural networks the calculation of the gradient of a path between a baseline and input [8].

3.3.2. Generation of Natural Language Explanations

The framework converts outputs of technical models into simple language descriptions:

Template-Based Generation: Relies on a set of predefined templates containing a set of predetermined feature values and weights of the attribution.

Neural Text Generation: Sequence-to-sequence models that achieve consistent coherent explanations on feature importance data.

Multi-Level Explanations: This has both high-level descriptions of these to the investigators and detailed technical descriptions of them to model validators.

3.3.3. Word Processing and Presentation

Interactive interfaces give the investigator a chance to experiment with model reasoning:

Feature Importance Charts: Top factors represented in a chart.

Counterfactual Explanations: Displays that which would have to vary in order to arrive at an alternate classification of the alert.

Similar Case Analysis: Gives historically similar cases and their conclusion.

3.3.4. Trust and Hesitation Communication

The system expresses confidence of prediction using:

- Calibrated probability scores with confidence interval.
- The indicators of uncertainty that are founded on ensemble variance.

- The flags of data quality which indicate missing or unreliable input data.

3.4. Workflow and Decision Layer-Risk based

The framework enforces the intelligent alert management regarding the risk scores calibration:

3.4.1. Multi-Level Risk Categorization

Alerts fall into 4 risk bands with their respective handling measures (Riti, 2025):

- **Critical Risk (Top 0.1%):** A sudden escalation, possible payment suspension, 24/7 review, and notification to the senior management.
- **High Risk (Next 1.9%):** Priority Review within 2 hours, Upgraded Due diligence, may be a relationship review.
- **Medium Risk (Next 8%):** The normal review of 24 hours later, simplified investigation processes.
- **Low Risk (Bottom 90%):** Automated disposition with 5-10% quality sample assurance (random), periodical pattern analysis on system problems.

3.4.2. Human-in-the-Loop Interface Design

Full decision support: Investigator interface gives detailed support:
Consolidated Alert View: All the information displayed in one dashboard such as customer profile, transaction details, match information, and AI insights.
Smart Workflow: Free-step investigation instructions according to the type of alerts and risk factors.
External Research Integration: One-Click access to third-party database systems, news items and research resources.
Documentation Templates: Construction of the investigation summaries and regulatory documentation with the help of automated programs.

3.4.3. Disposition Rules that are automated

In low-risk alerts, the system is automatized on the basis of rules:
Confidence Thresholds: Alerts that have a high risk score which is accompanied by a high level of confidence are automatically closed.
Pattern-Based Rules: Validated recurrent patterns of false positives are then automated.
Escalation Triggers: Automatic Notification of certain conditions that need human attention even though the scores are low.

3.4.4. The Feedback Loop of Continuous Learning

The improvement of the model is fed by every decision made by an investigator:
Active Learning: System is used to identify uncertain cases that are supposed to be given priority human review to enhance efficiency in learning.
Concept Drift Detection: Can be used to detect the change in patterns as time goes by and implements model retraining in case the performance reduces.
False Negative Analysis: The missed detections are reviewed periodically to determine the weaknesses of the model.

3.5. Governance, Compliance and Audit Layer

This base layer provides regulatory conformity and control of operations:

3.5.1. Risk Management Framework of Model R

Installs effective governance that is in line with regulatory advice:
Model Inventory and Classification: The full list of all the models and risk classifications in terms of materiality and complexity.
Development Requirements Rigorous documentation requirements including intended use, theoretical rationale, data provenance and implementation details.

Independent Validation: Three levels of defense that are development team validation, independent risk team assessment, and internal audit review [9].
Performance Monitoring: This involves constant monitoring of discrimination metrics, the quality of calibration, measures of fairness, and measures of stability.
Change Management: There are formal processes which are used in the model changes, version control and deployment approvals.

3.5.2. Prejudice Detection and Intervention

I am going to take proactive steps to guarantee the fairness of the algorithm
Pre-Processing Techniques: Representation Imbalances Training data balancing and reweighting.
In-Processing Methods: Constraints of fairness in training adversarial-debiased and regularized models.
Post Processing Adjustments: Thresholds should be altered according to various demographic groups in order to attain fair results.
Constant Checking: Frequent disparity auditing by nationality, geography as well as other features that are safeguarded with corrective measures.

3.5.3. Immutable Adit Trail System

Line by line recording of any screening activities:
Cryptographic Signing: All the audit records are cryptographically signed and timed with the help of blockchain methods.
Complete Logging: There are input data, versions of the model, values of features, risk scores, decisions, and justifications.
Tamper-Evident Design: Architecture is used to ensure that historical records are not altered and yet remain reachable.
Regulatory Reporting: Reporting which is pre-configured and meets certain regulatory requirement and generation with a single click.

3.5.4. Framework of regulatory engagement

Formal attitude to regulatory communication:
Documentation Transparency: Accessibility of system capabilities and limitations.
Performance Reporting: Frequent reporting of the key metrics and effectiveness measures.
Examination Readiness: This is through constant preparation of regulatory reviews using structured packages of evidence.
Change Notification: Preactive reporting of change in material systems to the concerned regulators.

4. THEORETICAL APPLICATION AND PERFORMANCE EXPECTATIONS

4.1. Complicated Detection Scenario Analysis

Scenario: Authorized personal utilizes obfuscated corporate hierarchy and transactions:
Background: Company A (registered in Cyprus), owned by designated person of the company "Ivan Petrov" has an interest in ownership in Company B (UAE) which in turn contracts with Company C (Switzerland). Company C pays Company D (Turkey) on services, the money is eventually paid to the associate of Petrov. Conventionally, the system performance is reflected through the following:

- At any point of transaction, no name matching is given directly.
- All payments are screened without any warnings.
- Whole avoidance triumphant.

Capability: It can identify the various elements within a framework.<|human|>Functionality: It is able to detect the different components in a framework.
Network Analysis: Finding ownership chain by cross-referencing corporate registry links Company D with Petrov over three degrees of separation with high confidence.
Behavioral Anomaly: Indicates an abnormally complicated payment chain of normal services as being statistically abnormal according to industry patterns.

Jurisdictional Pattern Recognition: Recognizes particular sequence of jurisdiction (Cyprus-UAE-Switzerland-Turkey) to be similar to known sanctions evasion typology with 87% similarity.
Temporal Analysis: This denotes clustering of transactions according to the designation date of Petrov.
Ensemble Scoring: This is a combination of network, behavioral and contextual information into high-risk score (0.92) without name matches.
XAI Explanation: Gives investigator a clear narrative: (1) 3-hop network connection to identified sanctioned person Ivan Petrov (85% confidence), (2) Plot of payment chain which does not match the purpose of business, (3) Pattern of jurisdiction, which fits a known evasion typology.
Result: Check deposit was detained, inquiry commenced, possible regulatory reporting activated.

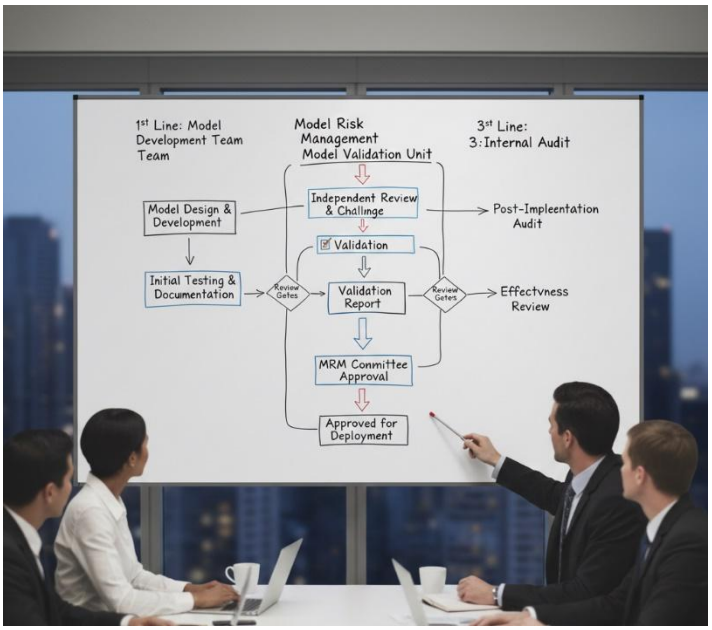

Figure 1.3: Risk Model Management

4.2. Quantitative Performance Forecasts

According to the similar AI applications in fraud prevention and monitoring transactions:

Table 1: Immunization of Expected Performance

Performance Metric	Current Benchmark	AI-Enhanced Target	Improvement
False Positive Rate	95-99%	15-30%	65-84% reduction
True Positive Detection	70-85%*	90-98%	10-28% increase
Investigation Time	20-45 minutes	5-15 minutes	60-75% reduction
Alerts per Investigator	10-20/day	40-60/day	200-300% increase
Cost per Alert	$25-75	$8-20	60-75% reduction
New Typology Detection	3-12 months	1-4 weeks	85-95% reduction

Operation Impact Analysis
Resource Reallocation: 70-80% of the time spent by investigators that is not in false positive review is liberated.
Quality Improvement: Investigators concentrate on high-risk cases that are complex using superior tools.
Cost Savings: $15-50 million of annual savings in case of medium-sized international bank.
Risk Reduction: The advanced evasion patterns are better detected [4].
Confidence of Regulation: Exhibit confidence is obtained through transparent and evidence-based processes.

4.3. Implementation Plan and Reflections

Phase 1: Foundation and Parallel Run

- Implement AI system and combination with current screening.

- Concentrate on low risk alert automation and decision support.
- Instituting governance structure and validation.
- Research new tools and procedures in training.
- Gather performance information to make a comparison.

Phase 2: Expansion and Integration

- Dilute into medium risk categories.
- Adopt feedback and constant learning.
- Integration with case management and reporting systems.
- Formal model validation and regulatory briefing.
- Start gradual phasing off of old rules.

Phase 3: Complete Implementation

- Full adoption to AI-enhanced screening.
- They use legacy system as fall-back.
- Grow to new products and jurisdictions.
- Put up sophisticated network analytics.
- Create ongoing process improvement.

Phase 4: Optimizing and Evolution (Continuous)

- Frequent training and improvement of the models.
- Increase into other compliance domains.
- Connection to larger financial crime systems.
- Constant regulatory involvement and reporting.

Critical Success Factors:
Executive Sponsorship: C-level investment in change.
Cross-Functional Teams: Co-operation between the compliance, technology, data science, business.
Regulatory Partnership: You can have supervisors on board.
Change Management: All-inclusive training and communication.
Iterative Development The use of Agile and regular delivery.
Performance Measurement: Measures that are clear and frequent.
Resource Allocation: Advancing enough in technology and talent.

5. CONCLUSION

The sanctions compliance environment has become critical to a point where the conventional technology solutions are no longer adequate to counter the advanced dynamic threats. The crisis of the false positives is not merely an operating efficiency issue, but a set of fundamental flaws in methodology of the detection that induce both high costs and operational risks of compliance.

More importantly, it implements explainability and auditability as design principles, which offer the transparency that regulators always need and offer the performance that institutions need. A governance model, easily manageable, model risk management, bias reduction and auditing trails provide a responsible roadmap towards AI implementation in the most regulated settings.

REFERENCES:

[1] Orekha, P. O. AI And Reinforcement Learning In Algorithmic Trading: Optimizing Market Execution, Liquidity, And Risk Exposure. https://www.researchgate.net/profile/Precious-Orekha/publication/389652844_AI_AND_REINFORCEMENT_LEARNING_IN_ALGORITHMIC_TRADING_OPTIMIZING_MARKET_EXECUTION_LIQUIDITY_AND_RISK_EXPOSURE/links/67e21675e2c0ea36cd9df8d8/AI-and-Reinforcement-Learning-in-Algorithmic-Trading-Optimizing-Market-Execution-Liquidity-and-Risk-Exposure.pdf

[2] Akkizidis, N. (2025). Responsible AI Governance Roadmap in Investment Firms From Zero AI to Scalable, Regulated, Real-World Adoption. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5792442

[3] Riti, R. I., Oțel, C. C., & Bacali, L. (2025). Queue-Theoretic Priors Meet Explainable Graph Convolutional Learning: A Risk-Aware Scheduling Framework for Flexible Manufacturing Systems. *Machines*, *13*(9), 796. https://www.mdpi.com/2075-1702/13/9/796
[4] Riberg, J., Selin, L. J., & Monaco, V. (2021). DEFENSE ANALYSIS CAPSTONE REPORT. https://calhoun.nps.edu/server/api/core/bitstreams/5689d8a4-e452-465f-9b53-4968423d0f21/content
[5] Riberg, J., & Selin, L. J. (2021). *Tip of the spear: can special forces lead the way for military applications of AI?* (Doctoral dissertation, Monterey, CA; Naval Postgraduate School). https://calhoun.nps.edu/bitstream/handle/10945/67802/21Jun_Riberg_Selin_Needs%20Supplemental.pdf?sequence=1
[6] Sun, Y. (2025). *A survey of statistical arbitrage pair trading with machine learning, deep learning, and reinforcement learning methods* (No. 2025-22). https://www.wne.uw.edu.pl/application/files/5617/5819/7786/WNE_WP485.pdf
[7] Riti, R. I., Oțel, C. C., & Bacali, L. (2025). Queue-Theoretic Priors Meet Explainable Graph Convolutional Learning: A Risk-Aware Scheduling Framework for Flexible Manufacturing Systems. *Machines*, *13*(9), 796. https://www.mdpi.com/2075-1702/13/9/796
[8] Estievenart, Y., Patra, S., & Ben Taieb, S. (2025, September). Risk-Based Thresholding for Reliable Anomaly Detection in Concentrated Solar Power Plants. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases* (pp. 111-128). Cham: Springer Nature Switzerland. https://link.springer.com/chapter/10.1007/978-3-032-06129-4_7
[9] Tan, B., Wang, Z., Duan, J., Xu, K., Shen, H. T., Shi, X., & Shen, F. (2025). Conformal Lesion Segmentation for 3D Medical Images. *arXiv preprint arXiv:2510.17897*. https://arxiv.org/abs/2510.17897