Deep Science Publishing
https://doi.org/10.70593/978-93-49910-62-1 _4

DeepScience
Open Access Books

Chapter 4: Using adaptive machine learning systems to detect, predict, and prevent fraud in a hyper-connected banking environment

4.1. Introduction

In today's financial system, fraud is a leading and serious problem. Fraudulent behavior can significantly harm financial institutions and is difficult to differentiate from legitimate transactions. The increasing amount of data allows the establishment of different models to detect fraud in banking. Various approaches have been proposed to tackle this issue, and several factors have drawn increasing attention towards the detection of fraud in financial institutions. Recent findings show that the use of new accounts is the main point for such actions, as they allow implementing borrowing and withdrawing behaviors, which are unfavorable for the banking system (Balasubramaniam et al., 2024; De Luzi, 2024; Jones & Tyson, 2025). Although there have been many studies on bank fraud detection, still there are many things required to be put in place. Once a pattern for a certain kind of fraud is established, some banking institutions implement the same pattern for a long time, which reduces its detection rate. Another problem is the class imbalance; there are only a few fraud actions compared to the huge amount of transactions with no fraud. When it comes to equipping a bank or other financial systems with an appropriate tool to automatically detect fraud actions, a static solution is not optimal: it would solve the problem in small periods of time and not in the long timescales, hence leading to non-optimal decisions. However, the rational behavior of fraudsters goes beyond the bank's fraud space, which makes it quite difficult to build static solutions. An ideal system would dynamically learn the behavioral space of the fraudsters: a transparent and self-adaptive model fitting for any kind of potential patterns that has high predictive performance on the data available, boosting the clear exposure of discrepancies for the fraud actions. The problem seems to be time-evolving in terms of both the data showing changes over time and the predictable solution evolving accordingly. This clearly indicates a solution that could quickly deal with

abrupt changes of the data moving away from its predefined distribution (Micheal, 2025; Nelson et al., 2025; Singh, 2025).

4.1.1. Background and Significance

Fraud management in banking is a critical, common, and persistent financial security problem, resulting in significant financial losses for the financial institutions. Fraud detection has existed for years, dating back to the 19th Century. As technology has developed, so have the methods that fraudsters have devised to commit their unethical acts. The banking sector has had to innovate and keep pace with the evolving schemes designed to deceive financial institutions and their customers. However, with the financial and non-financial consequences of frauds on the rise and the escalating sophistication of criminals, the design and implementation of an adaptive model for fraud detection systems is important.

Fig 4.1: Using Adaptive Machine Learning Systems to Detect

Fraud management has been defined as a complex set of technologies, applications, and agreements that work together to manage the issues associated with differing fraud types,

fraud attempts, and fraud prevention, detection, and investigation processes at all levels of the payment process. It is evident that collaboration, external analytics, and machine learning have the potential to enhance the current fraud detection capabilities. While all banks have existing fraud detection operations, the increasing impact and breadth of fraud combined with the demand placed on banks to respond quickly has led the banking industry to explore more destructive fraud detection techniques.

4.2. Understanding Fraud in Banking

Fraud is defined as any intentional deception intended to give someone a benefit in order to benefit the perpetrator. It can be seen in numerous fields and sectors including agriculture, business, cyber, education, elder, healthcare, insurance, investment, securities, mortgage, public benefits, tax, telemarketing and banking. It is unlawful. The perpetrator could be a single or group of people. Fraud losses for financial institutions are estimated to be more than tens of billions of U.S. dollars annually. The number of fraudulent financial transactions exceeding $15,000 has been steadily increasing.

In this section, we will focus on fraud in the banking system. After the emergence of online banking, investments and credit cards, criminals started committing fraud on a vast scale from any part of the world. Additionally, these actions negatively affect financial institutions, their customers and the overall economy. Criminals engaged in payment and card fraud transaction types exploit banking policies, operational procedures, and security systems. Fraud schemes include money laundering, bank false statements, using false passports, checks, wire transfers, or account access devices, proscribed moneylending, and banks doing manual balance discrepancies. The impact of fraud on the financial institutions is much greater than the actual financial amount lost. Financial institutions can experience: significant transaction costs and loss of infrastructure security that are associated with purchasing, deploying and maintaining new physical infrastructure; inability to scale their infrastructure and services to meet business growth opportunities; inability to comply with Critical Infrastructure Protection and regulatory data security requirements; and inability to provide continuous availability and recovery of services due to hardware and human errors, cyber, natural and unintentional disasters.

4.2.1. Types of Fraud

As the products and services of financial operations grow, there is a matching increase of interest by fraudsters in executing fraud schemes using fraudulent procedures. Indications are that the total cost of anticipation, detection, termination, and recovery from fraud is anywhere between 1% and 5% of sales for financial institutions. Financial

fraud losses are not generally covered by insurance, and in the case of banks, can worsen the existing relationship with clients through the letdown of trust. For these reasons, financial institutions are on the constant lookout for fraud detection technologies, and this is where the field of adaptive machine learning is poised to support solutions, helping to reduce reporting costs and decrease customer service effort during returns handling for both merchants and issuers. There are several types of fraud relating to banking. A culprit might assemble various strategies by which he/she exploits weak spots in operations of financial institutions. We present the main fraud schemes hereafter: • Identity Theft. Fraud stems from obtaining sensitive characteristics from one customer and exploiting them for illegally performing transactions on the customer's account. • Application / Account Misuse. Fraud stems from submitting an application to a bank using a person's data to open a fraudulent account for the purpose of getting loans, collecting credits, or laundering money. • Payment / Transaction Fraud. Fraud stems from attempting to use stolen access data for credit/debit cards to execute a fraudulent transaction in a bank. • Payment Accepting Fraud. Fraud stems from funds that should not have been accepted or have to be returned by the merchant, who doesn't conduct his/her business with due diligence. • Check Fraud. Criminals who commit check fraud are usually sending or receiving checks that are subsequently returned by the bank. This is not a complete list but these schemes are often handled by financial institutions by means of AI technologies.

4.2.2. Impact of Fraud on Financial Institutions

Fraud is one of the biggest challenges banks face today. Banks spend billions of dollars each year on anti-fraud measures, and despite these expenses, losses due to fraud continue to rise every year. The estimated fraud loss was over $16 billion. The cost to a bank of processing fraudulent transactions far exceeds the value of the loss. Additional costs include lag costs, chargeback costs, and distribution costs. But banks can suffer losses not only in terms of money but also in terms of reputation with their customers. The perceived risk associated with a transaction will be very high when trust in banks is low.

The first consideration of a bank when making decisions about anti-fraud measures is their costs. Financial institutions typically choose from the following three methods for imposing anti-fraud measures: Accepting the transaction with no restrictions; rechecking the transaction in real-time but with a slight delay; performing a retrospective analysis with no real-time merchant input. The first two alternatives can impose costs on banks by allowing transactions to be authorized that are actually fraudulent. This method normally results in the lowest customer compliance costs but the highest fraud costs for banks. The third option is the least expensive for banks, although it imposes the highest

cost for customers. Compared to the other two options, it provides banks with the best results in both detecting fraud and preventing it. It also does not discriminate against genuine customers.

4.3. The Role of Technology in Fraud Prevention

Banking became one of the first sectors to receive significant benefits from technology implementation. Customers were from the beginning attracted by the possibility of remotely interacting with their bank, executing operations like checking their balance or transferring money. The digital revolution of the late 90s and early 2000s accelerated the use of electronic money and digital banking. With the aim of improving user experience, banks started offering personalized services and facilitating transactions, which attracted even more customers. This technological evolution, however, did not come without costs. The introduction of new IT tools increased the number of channels through which customers could interact with the banks, but also the number of exposures to external attackers. Banks have been remorselessly attacked by criminals interested in stealing information for illegal purposes. Surprisingly enough, this immediate consequence of IT adoption was not enough to deter banks from implementing innovative technologies into their daily operation.

Fig 4.2: Role of Technology in Fraud Prevention

Nonetheless, the explosion of mobile banking services occurred in the last 15 years, rapidly promoted by the rise in use of smartphones and other portable devices, which

hugely intensified the cybercrime activity against banking institutions. Fraudsters found it easier, faster, and less expensive to attack bank systems through the Internet than any other methods that they had used until then. As a consequence, not only banks, but also regulators were forced to dedicate resources to fraud detection systems, which computationally verified the legitimacy of each transaction to avoid money loss by both banks and customers. CISO teams were created in the public and private sectors in order to safeguard critical assets against cyber risks. The same happened in banks, which intensified the use of cybersecurity tools and cyber threat intelligence systems that helped them in identifying probable victimization attempts through data analysis and behavior identification techniques.

4.3.1. Emerging Technologies in Banking

The banking sector has evolved rapidly thanks to the use of technology, which has created global solutions for business, financial, and payment systems. However, while providing banks with tools capable of high transaction volumes, broad global geographic coverage, user-friendly interfaces, etc., use of this technology has also made fraud easier. For instance, technology has taken online banking from primarily corporate to a new segment of the population, who needs banking services after experiencing a disaster. New payment methods are coming online daily, and the business environment is learning the vulnerabilities of those new methods, as many are based on confidential data, email accounts, and/or mobile devices. As a result, the banking sector has become a large target for fraud, and the cybercriminals have implemented technology to steal funds using phishing, identity theft, and denial-of-service via such methods, as fake international money transfers, ATM skimming, sending computer viruses via emails, and using devices to extract information from mobile phones.

Biometrics has a large application in anti fraud to confirm identity or possession for accessing customer accounts or making transactions; in identity processing to reduce external identity risks, prevent the fraud certificate from being issued, examine teller transactions to identify money laundering, reconcile transactions to detect false alarms, and deploy identity orchestration to reduce customer drop out; and finally, in account opening and reviews to enable fraud prevention and detection for call and contact centers. Another emerging technology in fraud detection is Artificial Intelligence. Acknowledgement experts, researchers, and scientists usually define Artificial Intelligence as a method designed for computers or other devices to be utilized in the same way as human reasoning.

4.3.2. The Importance of Cybersecurity

Cybersecurity breaches can affect banks adversely in many ways. Sometimes, gains by hackers in cyberspace are so huge that they threaten the viability of the banks. A famous case is that of a bank which lost a significant amount due to a lapse in its firewall. With such a significant loss from a single cybersecurity breach, it is no surprise that the global annual loss due to cyberattack is growing at a rapid pace. The global cost of cybercrime is predicted to exceed a substantial amount annually by 2025. The value of cybercrime is set to increase due to various reasons, one of which is the economic growth in both developed as well as developing countries. Many multinational banks operate on legacy systems. Hackers can exploit the vulnerabilities in these legacy systems. Cyber Attackers can also use suppressed countries as a launchpad because they do not have strong legal ramifications. This is the second reason why the value of cybercrime is set to reach such high levels. While technology vendors the world over boldly claim that they can offer the most efficient IT solutions to prevent cybersecurity breaches and offer highest levels of assurance to any customer, hackers too are catching up and devising more and more innovative methods to penetrate the strongest walls erected by banks. One of the surprising methods developed by hackers in the last decade involves planting malware tools in embedded systems used in ATMs and credit card swiping machines. Successful attempts at penetration through the internet of things framework are also being witnessed.

The battle between banks and financial cybercriminals is likely to escalate in the coming years. Banks must find out ways to consistently make investments in innovative technology to fight a war that requires advanced weaponry. An even more important question is whether they are really equipped with the innovative technology or is it just a façade that is going to end up in a war room planning how to reduce the damage to their systems. It is imperative that cyber threat hunting to prevent a successful strike by cybercriminals is a continuous process targeted at each bank's specific vulnerabilities. Banks need to utilize advanced weaponry using the latest methods, tools, and training to ensure that cybercriminals will think twice before embarking on their mission. Banks can collaborate with other banks and organizations to ensure the latest information is available. Instant coordination can minimize loss at organizations when a breach occurs. Security analysts need to be up to date.

4.4. Machine Learning: An Overview

What is Machine Learning? Machine Learning (ML) is a field of computer science which aims at understanding and constructing tools that can make predictions over unobserved data. Although solutions in ML can have extremely different characteristics, a defining feature is that the prediction model is automatically constructed from the observed data,

instead of hand-crafted by the human programmer. This script often contains thousands of parameters, which need to be adjusted according to the characteristics of the observation data. As we shall now discuss, various approaches exist to formulate such a statistical problem.

4.4.2. Types of Machine Learning Algorithms Although machine learning algorithms can be grouped in more or less loose categories according to very diverse criteria, the simplest and most widely used division is into supervised and unsupervised methods. Supervised algorithms construct the prediction model from labeled data, i.e., data to which the desired output is attached. For instance, in a bank surveillance the training set would consist of past customers' transactions, which are labeled as either being "normal" or "fraudulent". The task of supervised learning is to construct a model capable of labeling unseen transactions as either "normal" or "fraudulent". Unsupervised algorithms have no label information attached to the input data and thus the construction of the prediction model can be guided only by its structure. The task of unsupervised learning is to identify any underlying structure that can explain the training data. In the banking example, an unsupervised algorithm can identify a subset of transactions which are clustered in some region of the feature space; such unusual transactions would need to be tagged for further inspection.

4.4.1. What is Machine Learning?

Machine Learning (ML) is a research field that offers cross-disciplinary methods and techniques, able to extract patterns and knowledge from data. The utmost goal of ML is to study and develop algorithms that enable computers to learn from experience, expressed as the training data available to model on and provide a prediction on subsequently encountered data. Researchers in ML are interested in understanding how to create computer systems that automatically improve themselves with experience. Any academic or commercial domain can benefit from ML, since it enables the automatic construction of models based on data, such models being either mainly used to predict the values of new data, or to infer important characteristics of the data, patterns in the data being the important representatives of such characteristics. Models can be of many different forms, statistical models dealing with data as numeric tables or time series, data mining models, or logic-based models, able to represent relations among the data features, for example.

In essence, ML is a combination of two different areas: Artificial Intelligence (AI) and Statistics. ML appeared at first as statistical principles and concepts imported and applied in the research field of AI. Then appeared the inspiration mechanism of AI, applied to the domain of Statistics. The intersection of the two fields incorporates a wide variety of statistical and heuristic algorithms, each with different features, strengths and

weaknesses. These algorithms have all been designed to solve different problems, aiming to approach them in a different way, and often with different results too. The level of success is highly application dependent and it is very frequent that the same problem can be posed and solved using different methods. In recent years, both AI and Statistics have incorporated ideas from each other. Though the field started in AI, it is now broadly recognized as a sub-field of Statistics.

4.4.2. Types of Machine Learning Algorithms

What are the different machine learning algorithms? Due to the fact that the concept of machine learning is broad, it can be sub-categorized in different ways. One way of classifying machine learning is based on the supervised and unsupervised learning algorithms. In supervised learning, we have input variables (often called a feature) that describe the example instances and each instance has a response (often called a label) that is used as an example for the model. Supervised learning is a problem and the example of supervised learning is credit scoring. In unsupervised learning, we have only input variables, and we do not have responses for these variables. When applying unsupervised learning algorithms, we try to find subgroups (or groups) in the data.

Which algorithms are supervised learning algorithms? And what are the well-known supervised learning algorithms? Do supervised learning algorithms have to be general for all problems? The random forest, neural networks and logistic regression are most known supervised algorithms in practical cases. In this section we describe a few commonly used machine learning techniques. We also discussed some problems for the classification and regression problem. After that, this area is broad and there are many algorithms for each of these categories. Some classifiers, like support vector machines, are effective for classification only. So, they are specific to the classification problem. Other classifiers, like decision trees, are used for both the classification and regression are called general methods.

A common supervised problem can deal with both the classification and regression task. A classification task would be, recognizing if an email is a spam email, while the regression task would be, categorizing the email as a 10% spam email and a 90% non-spam email. Many of the supervised algorithms that we mentioned before can also tackle the classification task. For example, decision trees, random forests, artificial neural networks, naive bayes, and vector machines. All these classifiers can be used in the classification case.

4.5. Adaptive Machine Learning Systems

Real-world systems that leverage predictive models implement feedback loops, where the output of the model affects the actual system targeted by the model. This is notably the case for the decisions made by fraud detection systems in banking. The objective of these systems is not to maximize predictive performance. Instead, these systems should maximize long term profitability by balancing the different consequences of model errors. Thus, even if some predictions are sent for human inspection, alerts might still not be followed-up, and fraud inquiries might not be successful. Adapting an ML algorithm to the real-world system it targets is a much more difficult challenge than adapting a model to the real-world observed data, but it is also the most important challenge for success.

Adapting a system to the fraud detection problem at hand is essential. However, once deployed, an ML fraud detection model will likely not perform optimally. It is a consequence of several interrelated reasons: First, the distribution of inputs may evolve over time due to external factors, becoming very different from the training distribution. Second, the consequences of past predictions for the actual system might differ from those of current and future predictions. For instance, it might be very desirable to reduce the number of successful frauds targeted to high-net-worth clients, since the potential impact of such transactions for the bank is large. Third, the manner in which predictions and model outputs are leveraged in practice might change. For example, a bank might drop the majority of human review for the alerts generated by a specific type of model. The adaptivity challenges associated with these three types of identified changes make the challenge of deploying ML models in operational settings all the more complex.

4.5.1. Definition and Characteristics

Adaptive machine learning can be defined as a modification of the original design of a machine learning model intended to support the practical needs of the model users, mostly derived from the type of data and business, work relationships, technology, and user interactions with the machine learning model. An adaptive machine learning model development process is an ongoing process that allows for regular updates of all stages of the machine learning model development summarized in a general way in a business science model development funnel. For many tasks of data analysis, it is common to have classical machine learning models developed, tested, and operationalized at a scale that allows a diverse and demanding user base to consume the insights they deliver. Furthermore, the above work on classical machine learning models is usually based on a relatively small data sample of data relevant for the task.

A machine learning model is constructed using a selected portion of features and a size allotted to the model that allows it to resolve the task in the most efficient way. Typically, machine learning methods capitalize on more or less sophisticated technology to promote the efficiency of a developed machine learning model. Naturally, users should not be left to fend for themselves, as significant help is usually needed regarding how best to adopt the insights delivered by the model, in what way they can be fitted into a broader analytical operation, and how often the results can be expected to be refreshed as new data becomes available. However, although useful from a technical specification perspective, and certainly entertaining for users or prospects of using the machinery, machine learning is first and foremost a business operation preoccupied with delivering insights that help authorities and investigators use resources to better effect: minimize losses, optimize the growth of blocks of business that have smooth fraud detection parameters.

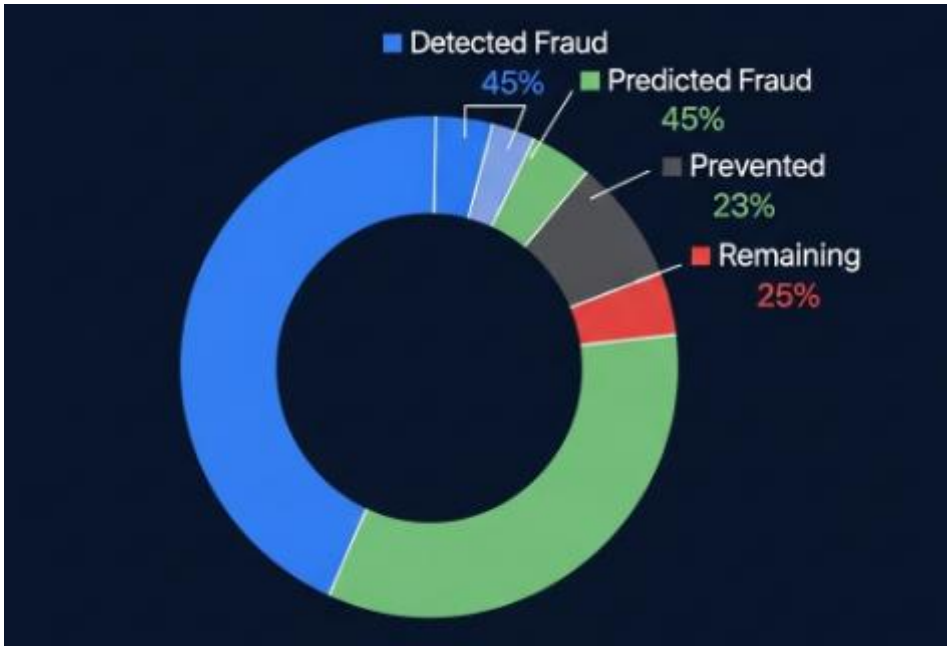

Fig : Prevent Fraud in a Hyper-Connected Banking Environment

4.5.2. Benefits of Adaptability

A typical machine learning system, once in use, will produce a relatively fixed collection of predictions for any given input; only the model weights change as they are updated in prediction errors. In contrast, an adaptive system chooses a different prediction function for each input, for example by initializing the weights of the base prediction model for new inputs using the base model's predictions from previous similar inputs. The

advantage of this simple approach to adaptation is that the system can use a fixed collection of prediction functions generated by the team using all the data; thus it can benefit from the lessons learned by all the personnel who are attempting to detect fraudulent transactions. Prediction functions will differ based on input characteristics. For example, a bank's data may show that fraudulent transactions on young adult customers typically take place in the early-evening hours of Saturday, while fraudulent transactions on retired customers typically take place during 9 AM to noon on a Tuesday. A better fraud detection system dynamically adapts the direction towards which the model is being trained based on the type of transactions – legitimate versus fraudulent – associated with each customer's account.

Model adaptation can also add significant representational power to a bank's system by enabling it to transition smoothly between performing the complex classifications that will likely be necessary during some transactions for the institution to predict fraud accurately, with the need to perform the easier classifications that will allow it to give legitimate customer transactions speedy responses. For example, on-line customers who normally do not download their bank balances, who suddenly do so after a major deposit has been made, and who simultaneously initiate several large wire transfers, may be committing fraud during a few transactions when it may be hard for the bank to task customers and thereby add delays to the processing of these transactions.

4.6. Conclusion

Adaptive machine learning systems are a new class of intelligent systems built on a foundation of intelligent software systems and adaptive agents. Although the building blocks that generate actionable intelligence for such systems have been developed individually over many years, it is only recently that the development of adaptive intelligent software systems has become an area of active investigation, and only in the past few years that the intelligent agent paradigm has truly emerged as an integrated framework for the development of intelligent software systems and systems-of-systems. The establishment of self-* attributes and other desirable characteristics for single or multiple, heterogeneous, co-operating systems, within a multi-agent paradigm, has led to new solutions for an entire range of open problems. Self-* intelligent agent systems embody the best of the autonomous agent paradigm combined with the desirable global feedback control from the intelligent software systems development.

The field is still at a very nascent stage in the sense that while most branches of general adaptive and intelligent software systems have received considerable attention for individual or multiple cooperating or competing systems, the concepts of intelligent software systems are not implemented as evolving planning or scheduling for single or multiple systems. There clearly exist challenges within the context that are particularly

pertinent, exciting, and will be the focus of research well into the future. These challenges include service-oriented reference model composition, self-* architectures and behavior, asynchronous design for evolvable agents and agent societies, self-* adaptive agent societies, and validation and verification of agents and agent societies.

4.6.1. Future Trends

The rapid growth of transaction data and the complexity of data in a digital economy have hastened advances in fraud analytics capabilities and performance. Transaction data is by nature large, online and irregular in structure; hence, efficient analytical approaches are favored by financial institutions when tackling the timely challenges of fraud detection. The rationale for the addition of more sophisticated techniques to the fraud fighting arsenal is that fraudsters themselves continue to innovate in tandem with technological advances and that they are cannibalizing one another's crime products and services. As machine learning and adaptive methods are increasingly employed for fraud detection, we also expect the development of ever more adaptable approaches to capturing the changing nature of the fraud detection domain. The principle of employing additional degrees of freedom in the analytical method can be coupled with supervised labels, but these are often sparse and hard to model. As the problem evolves, so does the surrounding environment; unsupervised anomaly detection provides the possibility of continuously adapting to the changing structure of transactions as well as being a record analyst that never forgets, i.e., utilizing a long memory window to discover new patterns of normalcy based on very infrequent events. As the analytics develop, so do the underlying technologies. With the rapid advances in cloud computing infrastructure and storage capacity, we can expect organizations to increasingly demand the availability and easy deployment of the latest algorithms, not just as underlying software services, but also as rapidly deployed expected black boxes so that fraud detection building blocks can be easily and invisibly incorporated within enterprise-wide systems for processing transactions in real-time to support key business processes.

References

Nelson, J., Lansnort, A., & Frank, E. (2025). Advanced Machine Learning Models for Fraud Detection.

Singh, H. (2025). Securing High-Stakes Digital Transactions: A Comprehensive Study on Cybersecurity and Data Privacy in Financial Institutions. Available at SSRN 5267850.

De Luzi, F. (2024). Digital transformation in a hyper-connected world. Transforming legacy systems through blockchains, IoT and AI.

Balasubramaniam, S., Prasanth, A., Kumar, K. S., & Kadry, S. (2024). Artificial Intelligence-Based Hyperautomation for Smart Factory Process Automation. Hyperautomation for Next-Generation Industries, 55-89.

Micheal, D. (2025). Comprehensive Review of Cybersecurity Frameworks: Fusing Machine Learning, Cryptographic Algorithms, and Blockchain for Resilient Digital Infrastructure.

Jones, H., & Tyson, K. (2025). Data-Driven Threat Mitigation: How AI Identifies and Neutralizes Cyber Threats in Real Time.