Deep Science Publishing
https://doi.org/10.70593/978-93-49910-78-2_10

DeepScience
Open Access Books

Chapter 10: Ensuring compliance with regulatory standards in artificial intelligence-driven financial systems

10.1. Introduction

The financial sector is pervaded by high uncertainty and at a constant risk of damage to multiple stakeholders, either voluntarily or involuntarily. The highly unpredictable, multi-stakeholder, and multi-dimensional implications of machine learning assurances in finance have caused regulators around the world to impose strict regulations on their use. The heavy documentation requirements imposed on AI systems primarily aim to increase transparency through collaborative scrutiny of different stakeholders by allowing audits to be performed. This auditability requirement raises additional challenges for the implementation of distributed ledger-based systems and can discourage companies from utilizing the advantages such technologies convey. Nonetheless, operating in a system lacking collaborative transparency can pose even higher risks. Hence, the use of AI systems in finance needs to be adequately scrutinized in a manner that maintains the advantages of decentralization while ensuring the maintenance of internal and external compliance (Addy et al., 2024; Balakrishnan, 2024; Garud, 2025).

This chapter examines the regulatory landscape concerning the adoption of AI by financial companies, including criticisms of regulators' current approaches and proposed improvements. In the latter part, it focuses on accountability as it pertains to internal and external compliance and discusses the bridges between AI and DLT technologies that can ensure that implementation of such mechanisms does not pose excessive additional burden on companies' operations. The chapter ultimately urges the use and further development of AI and DLT in conjunction to ensure that the technology and techniques developed by regulators in the search for compliance are made of the same core components used in the systems for which those compliance requirements have been developed (Kothandapani, 2024; Kothandapani, 2025).

10.2. Overview of AI in Financial Systems

The increasing complexities in financial systems, their constant state of evolution, or the ever-growing need for monetary security are several factors that have led to researchers and industry experts investing extensively in finding solutions to address these issues. Technology has stepped up to this need and has brought forth several solutions that are being readily adopted. Artificial intelligence is emerging as one such robust solution that is being implemented at various levels in financial systems. The advent of artificial intelligence and machine learning has introduced the ability for computer systems to learn from the techniques and applied knowledge of highly skilled financial experts and apply learned information to monetary systems that could not have been addressed previously if at all.

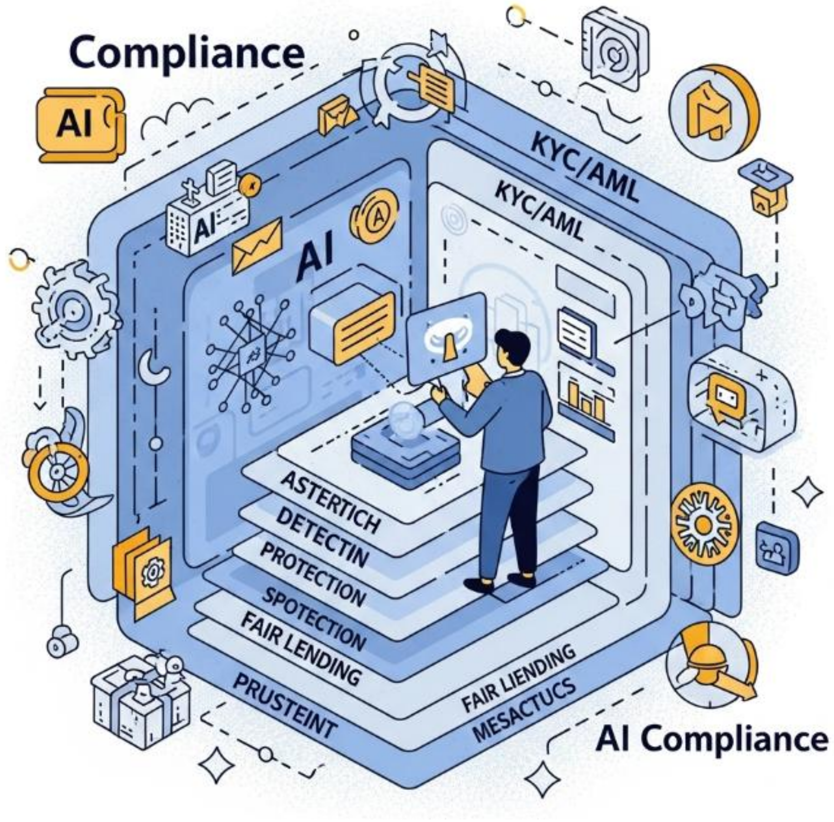

Fig 10.1: AI in Financial Regulatory Compliance

The outstanding performance of AI in automated trading systems has made it considerably popular as a means in creating business earnings. Several financial institutions are extensively relying on machine learning to forecast various problems. Financial experts leverage AI-neural network models to predict probability estimations in a financial decision-making setup. Artificial intelligence also has the potential to assist in protecting highly structured financial networks. The fast-paced requirement of companies to collect data regarding customer interaction patterns has opened doors for customer data modeling – an area which AI could significantly contribute towards in the coming years. Natural Language Processing procedures and techniques will soon be available to extract monetary data from various resources. The next wave of technological innovation will be AI-deployed digital labor. Financial companies will have access to digital advisors that are capable of scrutinizing thousands of rules and thus processes, allowing companies to streamline their processes and reduction of manpower costs as well.

10.3. Regulatory Landscape

There is a sizable and growing body of legislation, the regulations of which in many cases pre-existed AI. Financial Systems that employ AI necessarily have to comply with current regulatory guidelines. In addition to sectoral regulation, non-sectoral regulation on the handling of information about individuals also has to be taken into consideration for the handling of data which is required to train the AI systems. Regulations specifically affecting AI technology are at present relatively limited in scope. However, several jurisdictions are formulating and debating guidelines and framework legislation that restrict particular use cases of AI.

The regulatory landscape affecting the use of AI in the financial sector is summarized in relevant sections. Non-sectoral regulations affecting the use of AI are discussed in relevant sections. Key sectoral regulations affecting AI in Finance are summarized in relevant sections. These would include legislation which require disclosures regarding algorithmic bias in cases of algorithmic decision-making affecting consumers in protected classes. A particularly salient issue that is raising red flags among consumer advocates is so-called “black box” decision-making systems. The European Union is also reportedly preparing proposals related to AI that would require systems for certain high-risk-based AI decisions, including credit scoring, to be transparent and capable of being audited.

10.3.1. Key Regulations Affecting AI in Finance

As increasing numbers of financial firms use AI capabilities such as machine learning and data articulation, multiple financial services regulators around the world are considering whether – and if so, how – to regulate the development and use of such technology in the provision of products and services in the financial services space. Diverse regulatory proposals are on the table, with some agencies considering issuing principles and guidelines while others are focused on risk-based frameworks that govern certain aspects or elements of AI implementation. Despite the patchwork of regulatory proposals that could create confusion for financial institutions, it is essential for the industry players to stay on top of evolving regulations and adjust their AI-related activities according to these policies to encourage responsible innovation and clearly lay out penalties for noncompliance.

While many of the current regulatory proposals do not mention AI at all, it is difficult to escape the reality that their design and intent are in some ways shaped by the actual and potential harms AI is capable of doing. AI is covered indirectly in numerous pieces of legislation, such as provisions governing automated decision-making activities. However, there are a number of regulatory policy proposals that engage AI more directly, either in the form of sectoral oversight specific to financial services or broad frameworks that deal with AI being used in any industry. What impact the proposed legislation could have on the financial industry is still uncertain at this point, although each of these proposed regulations will likely put added pressure on financial services firms to consider the ethical implications of the services they provide, including how those services will be designed and supervised.

10.3.2. Global Regulatory Differences

Compared to the United States, there are greater regulatory differences across jurisdictions in a variety of areas defining the regulatory landscape such as consumer protection, risk management, privilege banking, and anti-money laundering. The regulatory environment is considerably more uniform in the U.S., where the regulatory framework across states and agencies is designed around shared principles for most provisions and risks. Moreover, there is no distinct specialized regulatory authority that defines the rules and standards for the use of AI in Finance. Instead, there are agencies which regulate financial entities for consumer fairness in transactions at various institutions. These rules have generally been equivalent to DEI, but may require amendment to meet the heightened scrutiny near the forthcoming U.S. Presidential election.

The variation in pace of regulatory response to the emergence of AI in Finance creates a potentially attractive competitive advantage for foreign firms over domestic firms. The absence of a distinct regulatory authority crafting specific rules, in the manner of financial regulation or prudential risk, could lead to significant regulatory arbitrage. A possible solution is for a consortium of regulatory authorities to establish a task force to engage foreign regulators, industry professionals, and relevant stakeholders on the issue of AI regulatory mapping. The goal would be the identification of specific work-streams where the domestic regulatory standards, derived from principles of fairness and risk management, can be mapped to similar foreign standards that can be harmonized at scale through reciprocity.

10.4. Compliance Challenges

The amount and nature of data collected and processed by AI systems generate a number of compliance challenges as new challenges are presented to existing regulation and new regulatory proposals. Existing regulation can be seen as problematizing the relationship between the user and the system since it is only partially able to protect the interests of users. However, regulatory proposals also generate compliance difficulties, as the suggestions proposed by regulators are rarely simple or straightforward. This inevitably leads to users being faced with a dilemma: comply or reject the service. The burden of complexity weighs heavily on companies that build and deploy AI technologies as requests for compliance deflect attention from the fundamental alignment problem of AI technologies. Companies will inevitably be forced to weigh-up costs and perceived benefits when weighing up regulatory compliance.

Data regulations impose strict guidance on the ways in which data is collected and processed and stored for future use. They stipulate that an organization must obtain the right consent from users; users must be informed of how data is going to be processed, the purpose it serves for accessing the service, how long it is going to be stored, if it's going to be transferred, and whether it's going to be used for automated decision making beyond and well in advance. For AI systems, organizations are faced with questions of how consent is obtained for the collection and sharing of data and what to do when the data is gathered from a data broker. And here we encounter real issues for AI systems: if the data is collected without consent, is the algorithm still valid? If the algorithm looks into non-consenting individuals, is that an ethical breach?

Algorithmic bias is a research topic that has become of interest to people developing and deploying ML for automated decision systems as well as for regulators. AI systems are known to express human biases that they learn from the corpus of data they are fed. The question of whether or not a regulator can legislate against algorithmic biases thus becomes pressing since these biases occur due to the inherent complexity of the algorithms themselves. Embedding fairness into the algorithms would require designing them in such a way that they do not learn, or allow users to learn sensitive attributes describing a user or take into consideration sensitive attributes.

The topic of transparency and accountability has also become a major concern for both AI researchers and developers and regulators for different reasons. Firstly, there is the paradox of AI system that are opaque by their very nature. In this case, it could be argued that regulators should only push for explainability in such cases of automated decision systems that are being trained for long-term deployment and not for short term needs or training. However, there seems to be pressure on all players: developers to develop AI systems that explicate their operation; end-users to seek counter-explanations from organizations; and organizations to provide them in case of a query.

10.4.1. Data Privacy Issues

AI adoption carries an intrinsic exposure to more stringent regulatory scrutiny. Indeed, regulators make organizations adopting AI responsible for crafting methodologies based on industry and service-specific risks and opportunities. Consequently, organizations must prepare for multiple possible regulatory compliance challenges. The first major layer of regulatory scrutiny comes from regulators' controls aimed at data privacy issues. Financial institutions have long been required to protect client data, and the rise of AI will make it more important than ever to ensure that no client has their privacy negatively affected by AI-driven decisions.

In fact, in the financial sector, AI is usually deployed to automate the execution of high-stakes decisions, potentially negatively affecting the consumer's privacy. A significant regulatory choice in this direction is the General Data Protection Regulation, which governs the capture of data around AI-based decision frameworks. The regulation states that consumers have the right not to be subject to significant automated decisions. However, parties can excuse themselves from this requirement if the decision derives from "appropriate safeguards," such as being controlled by a person and allowing consumers to express their opinion before the decision is made. A key question is whether AI models are capable of being compliant with this requirement while still being optimized to minimize distortion and maximize profit.

10.4.2. Algorithmic Bias and Fairness

Algorithmic models in AI-driven financial systems are increasingly at risk of uncovering, worsening, or perpetuating algorithmic bias and harm through their decision-making processes. Such discrimination could arise because of the original data or the algorithm's interaction with humans using the model. Derogatory bias effects using non-affected groups to mitigate group differences is a classic model analysis with a long history in social science and other fields. Choosing an algorithm that is likely to create inequality-based harms, or using a model for a task that is a priori not reliable for a specific demographic group, can lead to method-based discriminatory consequences. In many contexts, these algorithmic bias effects can arise regardless of how frequent the identification of specific groups at risk of algorithmic bias-based harms is.

With increasing deployments of predictive algorithmic models, it is critical for the developers to use domain expertise to ascertain if and when their models are useful, reliable, and fair in practice for specific affected groups. A linear model that assumes homogeneity cannot predict risk reliably across disparate subpopulations. When addressing a domain with disparate populations, it is critical to work closely with domain experts to ascertain model risk property and fairness for sub-k. Should the creators of the algorithmic models suspect any issues of trust, risk, and fairness in the final deployed models and their applications, additional post hoc model adjustments using domain-specific risk assessments and expertise may improve the model's acceptable risk deployments. Such clarifying procedures can lead model use to distance trust effects from known effects of disparities on outcome harmonies and possible heterogeneous harms.

10.4.3. Transparency and Explainability

The opaque nature of their learning mechanisms poses a threat to the safe deployment of AI in various circumstances, creating challenges for the areas of AI compliance. AI models are often inconceivable for their users, and the intricacy of their models raises doubts about their propriety. Why do some AI models produce results that people regard as correct when their reasoning is beyond understanding? Why do the models suddenly degenerate in performance when subjected to slight perturbations in the input part? Such confusion calls for the interpretable models to establish a certain level of transparency and faith.

Regulations and stakeholders have long requested a degree of transparency and explainability regarding algorithmic outputs, noting that a mere focus on performance does not justify non-transparent decision-making. One of the critical regulatory mandates in various new laws is the inalienable right to human intervention, which prohibits the deployment of AI systems aimed at human well-being without pre-established indices to rely on or human supervision on certain decision-making processes. In addition to the prohibitory stance, AI stakeholders are also growingly involved in designing and deploying AI systems in an ethical manner, advocating for the development of transparent and explainable AI systems. Not only for regulatory compliance or ethical fulfillment, transparency and explainability are also practically important for the AI research industry-center symbiosis.

10.5. Risk Management Frameworks

This section introduces risk management frameworks and describes how they can help financial services organizations identify the risks specific to their AI systems and workflows, as well as examples of tools and mitigation strategies for dealing with the risks that these frameworks help identify.

Without adequate transparency, interpretability, and expectations for correct functioning, it is very challenging for a financial system to be compliant with laws and regulations designed to protect consumers and set the standards for accuracy and equity in financial decision-making. However, these concepts may not fit well with AI systems in general, or may be justified or augmented in the context of AI-powered decision-making that is very prone to errors and bias, is not easily human-directed, or cannot be supervised. Using risk management frameworks can help banks better understand the risks and risk-mitigation strategies to employ for their specific payoffs, as well as for the specific operational models they use. Adopting a risk management framework should be a well-defined process that starts early in a product's lifecycle, is cross-functionally inclusive, and iterative throughout the lifecycle.

Risk management frameworks outline a set of generally accepted tenets for handling risk in an organization, typically in a cross-functional way that engages business decision-makers, risk managers, and internal auditors or assessors, among others. In the context of machine learning, only a few existing frameworks devote focus to the different functional areas an organization must engage when deploying AI systems into production and how to do so and iterate on it throughout a product's lifecycle. The latter are essential to not just detect unsafe systems, but also anticipate, mitigate, monitor, and validate systems that, when made available, can have potentially high social impact. They have general impact across all areas of society.

10.5.1. Identifying Risks in AI Systems

Before discussing how organizations can systematically manage the AI risks and leverage its benefits, it is first useful to point out the areas that the gaps in the existing risk management regimes leave open when asking the question—how might an AI-centric future differ from a future where AI does not play a decisive role? These gaps shed light on what the broader areas are that need identifying risks in AI systems. While there has been a great deal of fruitful research into risks associated with the individual steps in the AI lifecycle—and especially on the sort of problems that occur at the deployment and monitoring stages of that lifecycle, several of the discussions on gaps argued that risk management must do so in a more comprehensive way.

The limitations are twofold. Firstly, at least until more research is done, it is difficult to identify the deeper roots of the shortcomings. Secondly, because the AI mill has only recently come into operation, because empirical studies that document the special features of AI systems are in short supply, and because we are only just beginning to get a grasp of the philosophical and conceptual implications of the new environment, being blind to how such systems might differ in fundamental ways gets in the way of policymakers using available risk management guidelines. We might keep on returning to the AI lifecycle and its parts. But we believe it makes sense to ask deeper questions too. As a rough guidance, we might want to ask how AI systems differ in kind from classical systems in a way that affects the kind of risks that exist and/or the techniques that can be applied to identify, assess, mitigate, and monitor those risks.

10.5.2. Mitigation Strategies

As part of the development and implementation of a risk management framework, it is essential to identify a portfolio of suitable mitigation strategies for addressing the risks identified. These strategies are likely to vary depending on the nature, effects and degree of the risk posed. Mitigation strategies for risks associated with AI systems can be placed into one of three categories: prevention, reduction and review.

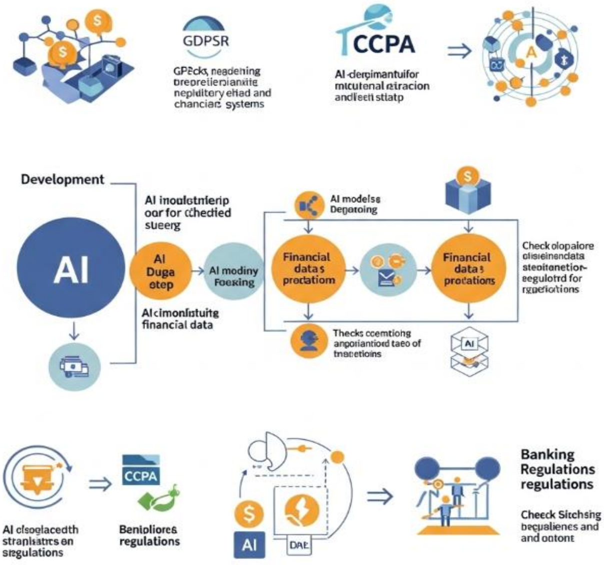

Fig 10.2: AI for financial compliance

Prevention strategies are actions that can be taken to reduce the likelihood of a negative consequence occurring. Developing robust and explainable AI models reduces the probability of negative consequences occurring significantly. And in doing so these models can also reduce the impact associated with a negative consequence. Failure to adequately do so can significantly increase the probability of an adverse outcome occurring. The implication here is that if a certain category of risk involves a high probability of a negative consequence occurring, it is better for firms to avoid using AI to associated processes or activities.

Risk transfer strategies essentially try to transfer the risk and liability associated with a possible negative consequence to a third party. Typically, this is done through insurance in cases where the public is acting as a third party, or by establishing contracts with other businesses or suppliers. Risk transfer strategies cannot eliminate risks associated with the use of AI systems; they merely move the burden to compensate and remedy the consequence. In addition, many negative consequences associated with the use of AI systems are often difficult, if not impossible, to insure.

Risk planning strategies are actions that firms can take to better prepare for a possible negative consequence. As with prevention strategies, by implementing and developing good risk planning strategies, firms can also help reduce the impact associated with a negative consequence. In a practical sense, scenario planning is about thinking through what could happen and proactively making plans that can minimize the damage if an AI system underperforms.

10.6. Best Practices for Compliance

In order to better ensure compliance with regulatory standards in AI-Driven Financial Systems, it would be beneficial for firms to adopt several industry-standard best practices. Following such compliance best practices often helps speed up the processes required to demonstrate compliance and engage with supervised regulators since a demonstrated commitment to compliance has been shown to help achieve a more beneficial working rapport with regulators. In an environment as complex and fast-changing as financial services, compliance must come to mean more than implementing check-list regulatory requirements. Organizations must move toward the creation of a proactive compliance culture throughout the organization. This culture will be the key to building a compliance organization that provides decision-makers with the clear and actionable recommendations on how to push toward the leading edge and reduce risk to a manageable level. To do that, financial services organizations must clearly signal to all stakeholders the priority they place on the compliance function.

Developing a Compliance Culture. Developing a proactive compliance culture is the first task for a compliance department in a fast-changing environment. An effective compliance organization will keep compliance on the agenda of senior management and boards, appropriately champion regulatory constraints and pressures, encourage an open dialogue about problems and encourage the proper representations of risk issues throughout the organization. It will be responsive to management's needs and offer an appropriate perspective when exploring new initiatives. A successful culture generally starts from the top of an organization led by a commitment from the CEO and the board. It may be helpful to assign a senior executive to ensure compliance is taken seriously and to oversee a program of training and reinforcing compliance objectives.

10.6.1. Developing a Compliance Culture

A solid compliance culture usually stems from the values of the top management and the Board of Directors of the company. These entities must believe in the integrity of the company and promote a transparent and clear corporate behavior aimed at ethical values. The Financial System is a position of great trust placed in the Financial Capital by its investors and those who analyze, evaluate and give an opinion about the company's assets or facilitate their transparency, either directly from the entity through the auditors or indirectly through regulatory entities and development agencies.

The compliance department prevents illicit practices and generates internal and external trust in the organization, but its work must be solidly sustained not only by control processes but also by training, knowledge, and the awareness of the people who work in the organization about the importance of prevention in order not to be accomplices or unaware of the risks assumed. Compliance placed people as active subjects of prevention and urged them to take actions such as whistleblowing that help to prove whether there are facts in the company that endanger the organization's integrity. A huge responsibility in the culture of the entity is attributed to the compliance department, which defines the policies and procedures that regulate how virtual assets can be bought and sold, how these transactions are recorded in the company's books, how the internal control processes function, and how the risk of a given operation turns into an operative risk for the company. It is important that all members of the organization know their responsibilities and duties about regulatory compliance, that they be aware of their importance, and that they know what kind of behavior or action can generate problems and risks for the company and themselves.

10.6.2. Training and Awareness Programs

Training and awareness programs are an effective way to boost compliance with laws, regulations, and internal policy requirements. Financial institutions should tailor their programs to meet the needs of diverse employee roles. In addition to ensuring that staff members are aware of their obligations, institutions should keep employees' attention and reduce information overload. Training should be delivered at different levels, be relevant to particular groups, and be designed for people with different learning styles using a mix of classroom training, computer-based training, and distance learning. Financial institutions should explore nontraditional training methods. For example, sessions may be popular with employees and allow for the sharing of experiences. Interactive training with quizzes can provide incentives to employees to understand and absorb training content. Clearinghouses can serve as referral resources for compliance staff.

Companies and firms can also use reminders in employee paycheck envelopes, compliance-related tips in the monthly employee corporate newsletter, and information in the company intranet to keep compliance issues fresh in employees' minds. Institutions should integrate compliance training into existing training programs. For instance, human resources staff can deliver compliance content during new employee orientation and internal audit staff during training on internal controls. Further, regulatory compliance should be part of formalized managerial assessments. For example, there are controls in the performance appraisal system that measure how each employee complies with regulations.

Training is especially important for maintaining compliance in a rapidly changing regulatory environment. The pace of change may be greater now than when the current compliance programs were designed. The trauma, uncertainty, and mystery surrounding recent bank failures have caused employees to be reluctant to breach the boundaries set by their institutions for compliance with laws and regulations. Bank regulators must work together and openly discuss any issues that arise from implementing training and awareness programs.

10.6.3. Continuous Monitoring and Auditing

Like code, at every step, digital financial systems share a lot of real-time information with various stakeholders. However, constantly auditing and monitoring that information requires significant human, technical, and financial capital, which startups cannot afford to be compliant at every step. More importantly, regulators are still refining their understanding of how to keep these systems in check without hampering innovation. They are now packing their toolkits of data-sharing platforms, auditors, and least-privilege access, for example, to perform continuous monitoring, which may become the norm in the near future. Currently, governments cannot rely solely on compliance through product testing at launch. Particularly as many financial services go unregulated currently, auditing and monitoring to create a fool-proof layer of trust takes the onus away from citizens. The solution lies in continuous conversations between companies, industry associations, governments, and the companies' stakeholders, which include citizens. Regulators need powers to periodically ask questions of system processes and system-created decisions because companies, even if compliant, may not be able to recall the specific data because it is not stored always. Service providers' reports to industry associations, who should act as a bridge between them and the regulators, should allow the associations to highlight common themes, like the nature of disallowed recurring transactions in crediting.

10.7. Case Studies

This chapter presents examples of financial institutions, ecosystem actors, and regulatory agencies that have successfully implemented compliance management systems for their AI-driven systems, as well as others that have earned severe sanctions for not doing so. In relation to the former, the implementation can be seen on a two-front basis. Initially, it represents successful implementation for an AI-driven service, which therefore does not harm anyone. One can think of fraud prediction, creditworthiness assessment, and money laundering detection systems, for example. Subsequently, we discuss the compliance management system for implementing compliance management in a company that provides AI-driven services, instead of providing services monitored by AI-driven systems. In the case of the European Union, the goal of requiring transparency and accountability towards users of any high-risk AI-driven service is enshrined in the proposed regulation for Artificial Intelligence. Thus, the first two come from the provision of AI-driven financial services, while the last two come from the proposed AI regulation.

In 2021, a financial institution paid a penalty to settle charges that its subsidiary failed to comply with rules designed to prevent its investment advisers from relying on inaccurate and incomplete data when voting fund proxies. The regulatory body found that the subsidiary did not adopt, implement, or maintain policies and procedures that were reasonably designed to ensure that proxy votes that were to be made using a proprietary model were in the best interest of funds and their shareholders. In particular, the subsidiary could, however, not demonstrate that it had the capacity and resources to objectively and accurately assess and select the appropriate algorithm.

10.7.1. Successful Compliance Implementations

Over the past few decades, several organizations have engaged Artificial Intelligence and Deep Learning in supporting and strengthening their compliance processes. The rush to employ machine learning models for this objective comes mainly from the fact that these types of models commonly perform considerably better than the models conventionally used. These usually are probability threshold classifiers, where the final decision is taken on the probabilities of the final output layer, and have been trained for decades on collective decisions made by armies of compliance experts with years of experience. With the advent of Deep Learning, there have been Master and Ph.D. theses where deep neural networks have been applied on compliance problems and achieved accuracy results considerably better than the conventional rules.

Despite this boost in compliance processes' performance, there is one area usually where rules are still used: The explanation of why a certain transaction has been flagged for further investigation. Innovation agencies of the different countries around the globe have repeatedly highlighted as a priority the development of algorithms accepted as "valid", in cases of non-combined systems. Non-combined systems are those systems where Deep Learning is used only for the detection part of the pipeline. There are proprietary AI-combined models for which by browsing through the values available, the score associated with the transaction that caused the penalty can be identified, along with some other information that can be useful in understanding if involved parties have any previous criminal record.

10.7.2. Lessons from Non-compliance

Introduction While the appointment of specialists tasked with monitoring regulations is a recommended practice for increasing compliance, it may not be easy considering the current skills shortage. Lack of organizational awareness related to compliance is a major cause of infractions, as are penalties for errors in good faith and lack of anonymity mechanisms. In the case of smaller institutions lacking a clear culture of compliance, the identification and appointment of a compliance officer are strongly recommended, personally reporting to the institution's board. In addition, internal information and training campaigns explicitly including all employees is key to ensuring smooth communication channels and proper checkups throughout the entire financial cycle. Many financial institutions have been found guilty of non-compliance in the past. These facts emphasize the importance of ensuring the proper implementation and configuration of AI tools for regulatory control these institutions are developing and adopting. The failure to comply, or to properly comply in a timely manner, can bring along devastating consequences as moneys involved in the knowledge of offense become enmassed, possibly inflating the fine due to lack of proactive action on the bank's side in identifying pertinent suspicious transference.

10.8. Future Trends in AI Regulation

In exploring future trends in AI regulation, it is necessary to consider that AI development will advance and therefore the system cannot be static. Additionally, AI development's influence on all aspects of our lives may beg a holistic approach to regulation that focuses not only on what AI is doing in finance or other industries but also how it influences society as a whole. Since it is likely that this holistic regulation will involve more than one sectoral regulator, then the holistic approach to regulation will need to ensure cooperation and information sharing between regulators, while also ensuring that the AI itself can respond to requests for updates that prove compliance. We have already hinted at the importance of collaboration in regulatory approach and this holistic collaboration is of the greatest importance. If ethical principles are provided by the holistic regulation, specific technical standards must be created to facilitate compliance with the ethical principles. Given the increasing connectivity of AI systems in processing information, technological innovation to foster compliance is crucial. Tech-enabled compliance can help companies who wish to consider compliance as a business priority. However, companies wishing to use the tech-driven compliance tools need to devote sufficient resources to training the tools in a way that facilitates compliance. This requires a team of diverse employees to ensure that the tech-driven compliance is holistic in nature as well. These considerations become more pressing when the internal and third-party tools develop in different directions. If there is a disconnect between various AI tools that a business is using, then reconciliation and monitoring of that AI tools is an ongoing and resource-extensive process.

10.8.1. Emerging Regulatory Frameworks

Increased public interest surrounding the development and deployment of deeply impactful artificial intelligence, driven by the integration of new generation systems such as large language models in commercial applications, has inspired governments and regulatory bodies around the world to create and implement regulatory frameworks governing its usage across sectors, particularly those most vulnerable to risks common to the technology in their processes, such as clear risks of errors, bias, and discrimination, privacy violations, security vulnerabilities, and reputational harm. These include financial services, where AIs are being used as tools for or drivers of anti-money laundering, credit decisioning and scoring, customer support and engagement, fraud detection, investment management, risk compliance, and trade surveillance. Given their usage and the vulnerability of their operations, financial business and institution must comply with increasing scrutiny from both consumers and regulators, and exploring these trends towards review and regulation of AI in finance is important to understand the implications of what is to come for their compliance efforts.

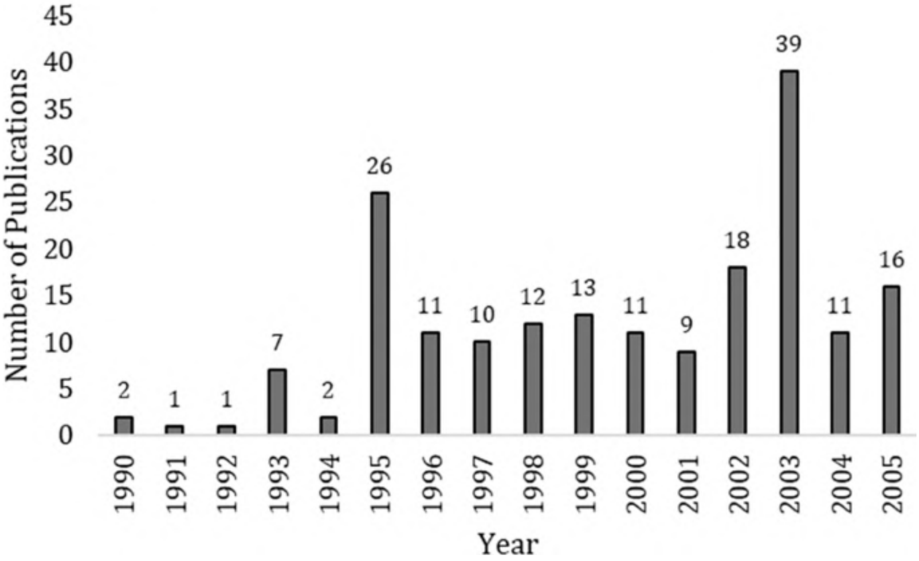

Fig : AI integration in financial services

The European Union has taken an early, unilateral approach to regulation of AI with the proposed framework embodied in the Artificial Intelligence Act released by the European Commission. Dubbed the “AIA”, it creates a regulatory pyramid constructed with four ascending risk levels, from minimal to unacceptable, that dictate regulatory compliance measures and requirements depending on the AI use case and its applicable level. These range from a simple requirement for compliance with a set of 15 EU Trustworthy AI Guidelines for apps and systems that fall into the minimal risk category, to the total banishment of AI systems from the market or their usage for harmful intended purposes, such as social scoring, eye-bio-biometrics and subliminal manipulation. There is an additional high risk category that is justified by these important areas being particularly vulnerable to the aforementioned AI vulnerabilities, within which repeated measures apply to the financial services sector.

10.8.2. The Role of Technology in Compliance

Technological advancements are easing the fulfillment and implementation of compliance efforts. Regulatory technology, or regtech, is a sub-sector of fintech whose services are geared towards providing solutions for the regulatory compliance industry.

With the number of regulations for financial companies increasing and growing complex and diverse on an international level, many companies have turned to regtech firms and services to better cope with their compliance needs. By using state-of-the-art technologies such as distributed ledgers, machine learning, and APIs, regtech services significantly lower the costs of compliance while ensuring better outcomes. In fact, these systems promise to provide regulatory authorities insights into compliance practices that would have previously only been available if the companies going through audits and examinations had opened their doors for the process. The regtech landscape comprises many types of applications ranging from know your customer solutions to transaction monitoring, fraud detection and prevention, reporting, compliance for market abuse, and agent and partner onboarding and due diligence. Outputs from regtechs can be especially useful for banks and regulators. By reducing the compliance burden, regtechs increase the value of regulatory insights and create higher incentive for financial institutions to share data with authorities. Moreover, many regtech applications utilize automated machine learning and are adaptable to the companies and situations they are dealing with. This means that machines can easily pick up on the features they should be paying attention to for a given company, while easing the burdensome aspects of regulatory compliance. With better data and promotion of information sharing across borders and business lines, digital assets can achieve a brighter future of regulatory compliance. By allowing companies to increase their focus on their customers while minimizing costs, regtechs also are creating new opportunities for banks to serve underserved populations.

10.9. Ethical Considerations

Artificial Intelligence (AI) systems are arguably the most transformative technology in contemporary world, bringing with their tremendous technological and economic potential, significant ethical concerns as well. The extent of AI's impact has sparked debates about its alignment with social goals and how errors in AI applications may cause disastrous consequences. Specifically, automated processes in AI systems may reflect existing inequalities of society, hence bias detection statistics have become an important matter. Although media coverage and technology research has grown from scandal centered information into a systematic exploration of ethical AI agendas with recommended practice, much of these explorations remain recommendations without thorough monitoring mechanisms. This raises the question if it is ever possible to humanize AI properly without losing its efficiency potential. Several debates discuss whether technically appropriate solutions align ethically relevant goals. The basic idea is that ethical AI could be achieved through gathering more heterogeneous data, fitness of any AI process must be monitored by testing against ethical goals of a society, and new techniques like explainable AI could be helpful for developers.

We contribute to this discussion by outlining the specific topic of ethical AI application in a regulatory manner, facilitating the exchange of the greater context of ethical AI in the more specific field of AI centric financial system regulation. Balancing between innovation and compliance is only the first aspect. To achieve this balance, the respective stakeholders need to be engaged to find a solution that is constructive for all of them. Important stakeholder groups in that discussion are represented by founders of AI system-enabled companies who are driven by competition and technological progress of their own countries. However, the group of established financial institutions need to be engaged as they represent the regulators' core focus of interest - the long-term stability of the financial system. Finally, politicians and regulators of the countries offering innovation frameworks also need to be included for international alignment.

10.9.1. Balancing Innovation and Compliance

Technological innovation in the financial sector can be stimulated by deregulation, while regulatory compliance can be seen as a burden by industry players. In designing a regulatory framework, policy makers face a fundamental tension: on the one hand, regulations must provide the right incentive structure to ensure that financial innovation occurs mainly in socially beneficial directions, with monitoring in areas such as identification and mitigation of risk to fundamentals, consumer protection, and long-term value creation. At the same time, regulation must be adequately flexible, to ensure that it does not stifle innovation, which is an important driver of growth and employment. It is often argued that this is more easily achieved when the regulatory framework is kept simple for the wider financial ecosystem than for niche players whose offer is only partially regulated or not regulated at all.

“Sandboxes” for fintech companies, where business models may be tested and temporary exemptions or lighter regulations granted, are one way of trying to strike a balance in this respect. While such initiatives may be an effective way of testing the waters for more disruptive and risky innovations, there remains a number of challenges, in particular: first, what is the right size, and how long should the exemption last, before a decision is made about whether the model should be further supported or normal regulations applied? Second, how is innovation seen from the standpoint of the other stakeholders of the company? Some business models, such as robo-advisers, could create a lot of value for the companies offering their services. However, they do not create any direct value for the economy as a whole since they only represent a substitution of one product for another when viewed from the angle of the advisory role. Hence, this innovation has a negative social impact as it increases the short-term pressure on those asset managers relying on active portfolio management.

10.9.2. Stakeholder Engagement

It is important to note that engagement with stakeholders might happen at any stage of the system lifecycle; however, it is generally recognized that the goal for engagement is different depending on the stage. For example, in the early stages of concept development, engagement helps identify risks and considerations important to stakeholders, which can be used to guide the design of the system. During the development and deployment of a system, stakeholders can be consulted to identify issues of concern that may have a negative impact on such populations. Moreover, different objectives for engagement may depend on the stakeholder group. Engaging with the developer and operator is different than engaging with end users or the subjects that are impacted by the deployment of the AI-enabled system. Obviously, users and subjects of the process are usually not involved in the design of the system.

For certain systems, however, it may be a legal requirement to include some or all stakeholders in the design of the technology. Specific industry standards can define scope and rules for the engagement with stakeholders. Many experts encourage going beyond regulations and promoting open discussion by experts and industry to define best practices for stakeholder engagement. It is commonly believed that many people will be impacted by a system, and each is unique; thus, it would be pointless to try to apply one single consultative methodology for all individuals. There are also frameworks for stakeholder engagement to show the possible options for different sponsors or stakeholders in an effort to make the technology assessment and risk management more understandable; mixed methodologies may be considered. Outlining the approach upfront can clarify the accountability of the developers and operators, making this a more transparent process with a clear governance model.

10.10. Conclusion

The deployment of AI applications in financial services has great potential, but it also implies challenges. Ensuring compliance with regulatory standards and satisfying the new demands of stakeholders is increasingly difficult in these novel market conditions and landscapes. The complexity of these systems may lead to adverse impact on customer safety, privacy, and trust. AI-driven financial services might also increase their vulnerability to cyberattacks, affecting the protection of the data involved. Because of this, the financial industry and regulators are called to work closely together to apply risk-based supervisory expectations that ensure regulatory compliance and the accountability of supervised entities. It is possible to state that it is not possible to provide AI tools and applications with a full set of compliance requirements and ensure that they have the widest possible use in the financial domain while, at the same time, ensuring compliance in the sense of not having any flaws or breaches.

This is true even in the financial services sector where the full combination of regulations does ensure a high level of safeguards, specifically tailored, but very demanding. The financial services industry has made considerable investments in applying and applying the rules with a considerable degree of knowledge gained along the years in sensitive sectors. The implementation of principles, rather than rules, is enabling the use of AI tools and applications at a much wider scale while keeping the prior need for confidentiality. There still is a long list of requirements, operational and otherwise, that must be fulfilled to ensure that the application is compliant, especially if looking at a global level. Compliance is expected to negatively correlate with the number of applications developed. At the same time, this list remains broad, stemming from widespread areas, such as Data Governance, the proposed AI Act, GDPR, Financial Action Task Forces' Guidelines, MiFID and MiFIR regulations, all the regulations on data interpretation and the implementation of International Financial Reporting Standards.

References:

Balakrishnan, A. (2024). Leveraging artificial intelligence for enhancing regulatory compliance in the financial sector. International Journal of Computer Trends and Technology.

Kothandapani, H. P. (2025). AI-Driven Regulatory Compliance: Transforming Financial Oversight through Large Language Models and Automation. Emerging Science Research, 12-24.

Kothandapani, H. P. (2024). Automating financial compliance with AI: A New Era in regulatory technology (RegTech).

Garud, S. (2025). AI-Driven Risk Management in Financial Services From Theory to Practice in Smart Education. In Smart Education and Sustainable Learning Environments in Smart Cities (pp. 77-92). IGI Global Scientific Publishing.

Addy, W. A., Ajayi-Nifise, A. O., Bello, B. G., Tula, S. T., Odeyemi, O., & Falaiye, T. (2024). Transforming financial planning with AI-driven analysis: A review and application insights. World Journal of Advanced Engineering Technology and Sciences, 11(1), 240-257.