National College of Ireland

EU AI Act Compliance Framework for Financial Institutions: A Knowledge-Driven Approach

MSc Research Project
MSc in Fintech

Cheah, Eric Wei Lun
Student ID: 23431881

School of Computing
National College of Ireland

Supervisor: Victor del Rosal

National College of Ireland

MSc Project Submission Sheet

School of Computing

Student Name: Eric Wei Lun Cheah

Student ID: 23431881

Programme: MSc in Fintech **Year:** 2025

Module: Practicum

Supervisor: Victor del Rosal

Submission Due Date: 11 August 2025

Project Title: EU AI Act Compliance Framework for Financial Institutions: A Knowledge-Driven Approach

Word Count: 5991 **Page Count:** 20

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.
ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature: Eric

Date: 30 July 2025

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST

Attach a completed copy of this sheet to each project (including multiple copies)	□
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	□
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	□

Assignments that are submitted to the Programme Coordinator Office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

EU AI Act Compliance Framework for Financial Institutions: A Knowledge-Driven Approach

Cheah, Eric Wei Lun
23431881

Abstract

The European Union's Artificial Intelligence Act presents compliance challenges for financial institutions, with penalties reaching €35 million or 7% of global annual turnover. This research develops a proof-of-concept expert-augmented decision support framework addressing implementation gaps between regulatory requirements and practical guidance within financial services. Using Design Science Research methodology, this study analysed EU Regulation 2024/1689 and developed a three-component modular architecture: a Regulatory Knowledge Engine organising 12 core regulatory articles, an Expert Guidance Interface providing role-specific navigation across technical teams, compliance officers, legal departments, and senior management, and a Documentation Assistant offering template validation through simulation-based assessment workflows. The web-based implementation demonstrates structured regulatory knowledge organisation through single-file HTML architecture while maintaining clear acknowledgement of proof-of-concept limitations including visual document simulation and pre-configured institutional scenarios. Commercial viability assessment using publicly available institutional data demonstrates protection ratios across different organisational scales: Allied Irish Banks (16.1x-36.1x), Banco Santander (72.5x-108.8x), and ING Group (39.5x-63.2x), providing investment justification under stated modelling assumptions. This research contributes expert-augmented RegTech frameworks that maintain human expertise requirements while providing structured implementation guidance for the August 2026 compliance deadline. Future research requires stakeholder validation, production-level development, and effectiveness assessment for institutional implementation.

Keywords: EU AI Act, financial services compliance, RegTech, expert-augmented decision support, Design Science Research

1 Introduction

The European Union's Artificial Intelligence Act, entered into force on 1 August 2024, represents the world's first AI legal framework, creating unseen compliance challenges for financial institutions. With implementation deadlines extending through August 2026 for high-risk AI systems commonly deployed in banking operations including credit scoring, fraud detection, and customer service automation, financial institutions face significant regulatory complexity requiring robust compliance approaches balancing innovation with risk management obligations (Montagnani et al., 2024; Kilian et al., 2025). The Act's risk-based categorisation system, penalty structure reaching €35 million or 7% of global turnover, and integration requirements with existing financial services regulations create implementation challenges inadequately addressed by current literature and commercial solutions.

Current research reveals a critical gap between theoretical EU AI Act compliance requirements and practical implementation guidance specifically tailored for financial institutions, particularly regarding integration of AI Act obligations with existing risk management frameworks (Mazzini and Bagni, 2023; Truby et al., 2020). While existing RegTech solutions focus mainly on automated compliance approaches, documented limitations in complex regulatory interpretation suggest that expert-augmented decision support systems may provide more effective pathways for institutional compliance readiness (Freij, 2020; Jain et al., 2024). Irish financial institutions such as Allied Irish Banks face particular challenges developing commercially viable compliance frameworks accommodating both regulatory adherence and continued innovation in AI system deployment.

This research addresses the overarching research question of how knowledge-driven compliance frameworks can be designed to support EU AI Act implementation in financial institutions while demonstrating potential commercial viability through structured modelling and proof-of-concept development. This central question decomposes into four specific research sub-questions that systematically address the implementation challenge. (1) How can core EU AI Act requirements be systematically organised into structured knowledge representations suitable for financial services applications? (2) How can expert-augmented decision support interfaces be designed to complement human compliance expertise through role-specific guidance delivery? (3) What framework for commercial viability assessment can be developed through cost-benefit modelling across different institutional scales? (4) How can a three-component architecture approach demonstrate proof-of-concept feasibility and scalability concepts for European financial institutions?

The study addresses these questions through Design Science Research methodology employing systematic artifact development across five phases encompassing regulatory analysis, knowledge architecture design, and commercial viability assessment. The research develops structured knowledge representation systems mapping 12 core EU AI Act articles to actionable guidance, creates expert-augmented decision support interfaces with role-specific navigation for technical teams, compliance officers, legal departments, and senior management, and demonstrates commercial viability potential through business case modelling across Allied Irish Banks, Banco Santander, and ING Group scenarios. The three-component architecture validation occurs through proof-of-concept development demonstrating intended functionality while maintaining transparent acknowledgement of simulation-based readiness rather than production-level processing. This approach contributes structured frameworks for EU AI Act compliance specifically contextualised for financial services applications while advancing RegTech research through human-centered design approaches that acknowledge regulatory complexity requiring professional judgment rather than purely automated solutions.

2 Related Work

2.1 EU AI Act and Financial Services Compliance

The EU AI Act introduces new compliance complexity in the field of financial services regulation, where the existing frameworks struggle to address practical implementation challenges. Particularly, Montagnani et al. (2024) imply that financial institutions face two main regulatory pressures: AI liability frameworks must integrate with sector-specific financial regulations, resulting in potential conflicts between risk-based AI approaches and traditional subject-based financial oversight. This regulatory intersection presents huge uncertainties regarding liability allocation and enforcement framework, especially when AI systems malfunction within existing financial oversight.

Implementation challenges extend beyond theoretical compliance to operational implications. Particularly, Kilian et al. (2025) state that financial institutions require minimum 12-month implementation periods per technical standard, yet face strict timelines with estimated compliance costs of €100,000-300,000 annually. Furthermore, their analysis involving 23 organisations demonstrates that smaller financial institutions face disproportionate challenges due to resource constraints and limited participation in standardization processes. These lead to unfavourable market positioning and competitive disadvantages that result in market power consolidation favourable for larger institutions that are capable of absorbing compliance costs.

The sector-specific provision within the AI Act further enhances the complexity; Mazzini and Bagni (2023) imply through their analysis of credit scoring applications classified as high-risk systems, which is a critical function within the financial sector. Moreover, the Act's integration with existing financial legislation (Directive 2013/36/EU) requires institutions to align AI governance with established risk management frameworks, however lacks concise practical implementation guidance. The overlapping regulatory requirements lead to compliance ambiguity, where expert interpretation is needed that challenges the institutions' compliance strategies.

Current literature reveals a critical gap between theoretical compliance frameworks and practical implementation methodologies specifically tailored for financial institutions. While existing research addresses regulatory requirements, it fails to provide structured decision support systems that interpret complex regulatory texts into actionable guidance for different stakeholder roles within financial institutions. This research addresses this gap by developing knowledge-driven compliance frameworks that systematically organise EU AI Act requirements into commercially viable implementation pathways, specifically contextualised for Irish financial institutions such as Allied Irish Banks.

2.2 AI Governance Frameworks in Banking

Contemporary AI governance frameworks in banking reveal a fundamental disconnect between technical automation readiness and regulatory compliance requirements, particularly regarding EU AI Act implementation. While Kamisetty and Nagamangalam (2025) demonstrate that automated governance systems achieve 95% accuracy in fraud detection through six-domain frameworks encompassing data quality and security protocols, their approach critically lacks the human oversight mechanisms mandated by EU AI Act Article 14 for high-risk financial applications. This technical-regulatory gap represents a significant limitation in current governance paradigms.

Kurshan et al. (2020) identify twenty-one specific governance challenges, arguing that traditional financial model frameworks fundamentally misalign with AI characteristics such as intrinsic learning and conditional uncertainties. Their proposed eight-readiness system-level framework emphasises continuous monitoring and human-in-the-loop alerting, yet fails to address EU AI Act-specific requirements for human oversight in automated decision-making processes. This limitation is particularly pronounced given that 67% of US financial institutions perceive regulatory complexity as problematic, increasing to 94% among AI-mature organisations.

Truby et al. (2020) advocate proactive regulatory approaches incorporating seven governance principles, including human agency and transparency, which align with EU AI Act foundations. However, their framework predates EU AI Act implementation and inadequately addresses the Act's risk-based categorisation system requiring differentiated governance approaches for prohibited, high-risk, and limited-risk AI systems. Furthermore, their reliance on GDPR Article 22 precedents oversimplifies EU AI Act compliance complexity. Paleti (2024) introduces neural compliance concepts enabling real-time risk protocol implementation through multi-stratified systems, while Wali et al. (2023) demonstrate that governance understanding exists but implementation faces huge legal and environmental obstacles. Critically, both approaches prioritise technical automation over the expert-augmented decision support systems required for EU AI Act compliance interpretation and implementation.

The literature reveals a critical gap: existing frameworks emphasise either full automation or traditional manual processes, failing to address the expert-augmented decision support requirements necessary for EU AI Act compliance in financial services, particularly for institutions requiring commercially viable implementation strategies.

2.3 RegTech Solutions for Compliance Automation

Contemporary RegTech research reveals fundamental tensions between automation readiness and regulatory complexity, particularly for EU AI Act implementation in financial services. While solutions promise enhanced efficiency through artificial intelligence and machine learning (Boopathi and Raja, 2024; Jain et al., 2024), critical analysis exposes gaps between technological readiness and regulatory specificity. The automation paradox emerges

as literature emphasises reduced human intervention (Abikoye et al., 2024; Gatla, 2024), yet EU AI Act Articles 14 and 15 explicitly mandate human oversight mechanisms that challenge fully automated approaches.

Freij's (2020) analysis of 550 RegTech providers reveals that 40% support only single-regulation compliance, highlighting systemic fragmentation undermining regulatory coverage. Current literature exhibits technological determinism that oversimplifies regulatory complexity. Jain et al.'s (2024) analysis reveals that while "AI algorithms can process vast amounts of data quickly," these systems struggle with "ambiguities and exceptional cases that do not fit well-defined patterns." This limitation becomes critical when examining EU AI Act Article 6's prohibition requirements, where determining "unacceptable risk" involves contextual interpretation that automated systems cannot reliably perform without human expertise.

Integration challenges reveal crucial underestimation of legacy system compatibility. Abikoye et al. (2024) acknowledge that "many financial institutions operate with outdated IT infrastructures," yet research consistently minimises implementation complexity. Economic viability analysis demonstrates bias toward technological benefits while underestimating implementation costs including integration, training, and regulatory adaptation requirements.

This literature gap reveals absence of expert-augmented compliance frameworks specifically designed for EU AI Act implementation in financial services. Existing research lacks institution-specific solutions addressing practical deployment challenges while maintaining human oversight requirements. The proposed EU AI Act Compliance Assistant addresses this gap through expert-augmented decision support architecture, combining regulatory knowledge mapping with role-specific guidance and documentation compliance assistance. Unlike existing automated solutions, this approach acknowledges human expertise necessity while providing structured implementation frameworks specifically contextualised for Irish financial institutions, demonstrating commercial viability through Allied Irish Banks business case modelling and realistic technology integration strategies.

2.4 Knowledge-Based Systems for Regulatory Implementation

Knowledge-based systems reveal fundamental limitations in regulatory contexts despite their computational advantages. Pérez Gómez (2022) demonstrates that rule-based expert systems, while effective for encoding legal principles through symbolic logic, exhibit critical rigidity when confronting novel regulatory scenarios requiring interpretative flexibility. This limitation is amplified in financial regulation where contextual nuances challenge deterministic approaches that assume complete codification of regulatory requirements (Paleti, 2022).

The automation-interpretation paradox presents significant challenges for regulatory AI implementation. While Bollineni (2025) reports 87% automation rates for routine financial decisions and Paleti (2022) documents 85% reduction in manual compliance intervention, these quantitative improvements mask qualitative concerns about algorithmic bias and interpretative limitations. Tillu et al. (2023) acknowledge that data-driven regulatory frameworks utilising

natural language processing require human judgment for ambiguous scenarios, contradicting claims of automation readiness.

Hybrid approaches demonstrate superior performance compared to fully automated solutions by preserving human expertise while leveraging computational efficiency. Yang et al. (2021) propose multi-criteria decision frameworks using AHP, FAHP, and TOPSIS methodologies that maintain human oversight while providing strong evaluation readiness. Gunturu et al. (2024) emphasis that AI models effectively support data-driven tasks but cannot replace domain expertise where regulatory misinterpretation carries large penalties.

Current literature predominantly examines automated compliance solutions while neglecting expert-augmented alternatives that acknowledge regulatory complexity. This gap is particularly evident in EU AI Act implementation where regulatory novelty demands interpretative judgment beyond algorithmic readiness. Our project addresses this limitation by developing expert-augmented decision support systems that enhance rather than replace regulatory expertise, providing structured knowledge organisation while maintaining human interpretative authority essential for effective EU AI Act compliance in financial institutions.

3 Research Methodology

This study employed Design Science Research (DSR) methodology due to its appropriateness for developing practical ICT artifacts addressing real-world regulatory challenges while maintaining academic rigor. DSR methodology was selected over alternative approaches such as empirical field studies or purely theoretical analysis because it provides the structured artifact development process necessary for creating practical compliance solutions while maintaining rigorous evaluation standards required for academic validation. The framework enabled progression through problem identification, solution design, artifact development, and demonstration phases with continuous refinement based on evolving regulatory requirements.

The research employed a qualitative-dominant approach with embedded quantitative elements, specifically utilising document analysis and artifact development methodologies for regulatory compliance research. This methodological choice addresses the complex nature of regulatory compliance, requiring deep qualitative understanding of legal frameworks combined with practical demonstration of technical solutions. The approach aligns with UN Sustainable Development Goals (SDGs) 8, 9, 10, and 16 by creating decent work opportunities in AI compliance roles, fostering innovation in regulatory technology infrastructure, reducing algorithmic inequalities through transparent AI systems, and strengthening institutional accountability through robust governance frameworks.

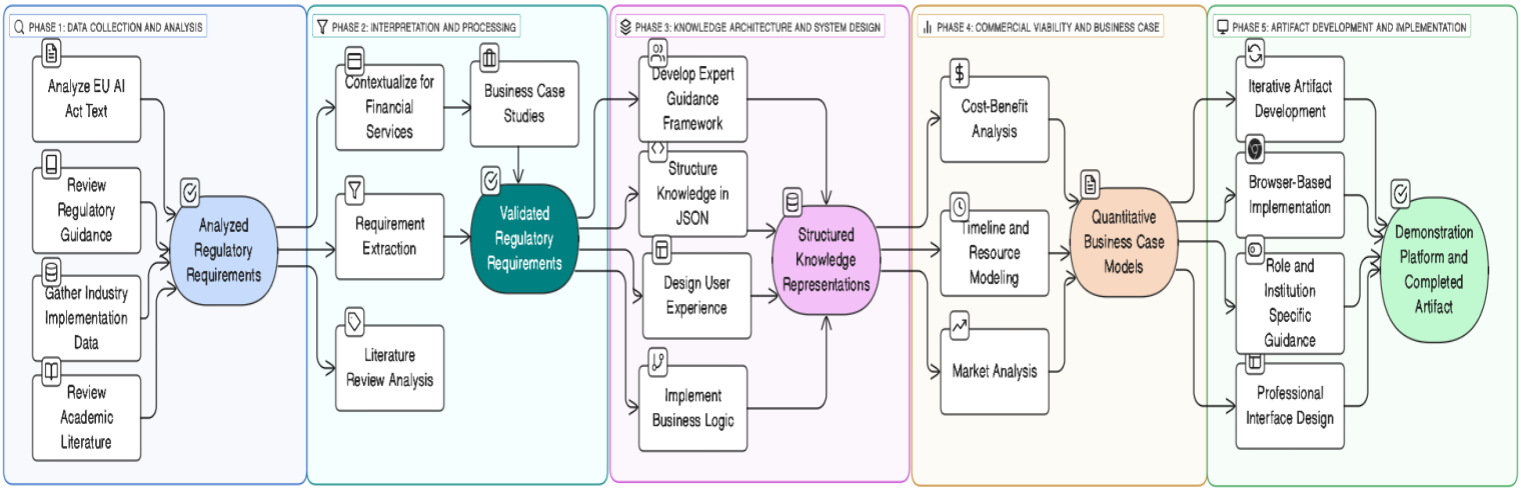

Figure 1: Research Methodology Overview

As shown in Figure 1, Phase 1 involved collection and analysis of regulatory requirements through systematic document analysis methodology employing thematic content analysis with coding frameworks for requirement categorisation. Primary data sources included EU Regulation 2024/1689 (EU AI Act), with analysis of 12 core articles including Article 5 (Prohibited AI Practices), Article 9 (Risk Management System), Article 10 (Data and Data Governance), Article 11 (Technical Documentation), Article 13 (Transparency and Provision of Information to Deployers), Article 14 (Human Oversight), Article 15 (Accuracy, Robustness and Cybersecurity), Article 17 (Quality Management System), Article 50 (Transparency Obligations), Article 57 (AI Office governance), Article 61 (AI Board governance), and Article 64 (Competent Authorities governance). Additional regulatory sources encompassed EU AI Act Annexes III and IV for high-risk system specifications and technical documentation requirements, and official guidance documents from the European Commission on AI Act implementation. Industry information was gathered through publicly available data sources including AIB Annual Report 2024, Banco Santander Annual Report 2024, and ING Annual Report 2024, supplemented by institutional websites, regulatory filings, and press releases.

Phase 2 transformed analysed regulatory text into structured, implementable requirements tailored for financial services applications through systematic requirement extraction and contextualisation processes. Processing techniques included systematic identification of mandatory, recommended, and conditional requirements from EU AI Act text using structured coding methodologies aligned with qualitative content analysis frameworks (Braun & Clarke, 2006). This phase involved financial services contextualisation mapping generic regulatory requirements to specific banking use cases including credit scoring systems, fraud detection algorithms, customer service chatbots, and investment advisory tools.

Phase 3 converted processed regulatory requirements into structured knowledge representations suitable for expert-augmented decision support through organised transformation approaches using established information architecture principles (Rosenfeld et al., 2015). The methodological approach involved developing structured knowledge mapping

processes creating organised regulatory requirement, complemented by expert guidance framework development targeting different stakeholder roles within financial institutions.

Phase 4 developed quantitative models for compliance costs, implementation timelines, and commercial viability using publicly available financial data and established financial modelling methodologies. Assessment components included cost-benefit analysis comparing implementation costs against regulatory penalty risks, timeline optimisation for the August 2026 compliance deadline, and return on investment projections for different institutional scales. Business case modelling utilised publicly available financial data rather than industry benchmarks due to the nascent nature of EU AI Act compliance market and absence of established benchmarking frameworks.

Phase 5 focused on proof-of-concept validation rather than empirical testing with stakeholder groups, employing systematic functionality demonstration within controlled development environments. Validation methodology involved demonstration of intended functionality while maintaining transparent acknowledgement of simulation-based readiness rather than production-level processing.

Research limitations include business case modelling based on publicly available data rather than proprietary institutional information due to access constraints and confidentiality requirements, stakeholder validation limited to academic supervision rather than extensive industry expert consultation due to time and access limitations, and scope focus on EU AI Act compliance without integration assessment with existing regulatory frameworks due to project boundaries. Ethical considerations encompassed exclusive use of publicly available data sources, transparent acknowledgement of proof-of-concept limitations, and honest assessment of commercial viability claims. Research ethics approval was not required as the study involved only secondary data analysis of publicly available information without human participants or proprietary institutional data access.

4 Design Specification

The EU AI Act Compliance Assistant utilises a three-component modular architecture designed as an expert-augmented decision support system that acknowledges regulatory compliance requires human knowledge. The design methodology prioritises improving rather than replacing expert analysis, displaying proof-of-concept demonstration while acknowledging the limitations throughout the framework. The architecture accommodates modular expansion beyond current scope, supporting future integration of new regulatory framework.

The platform implements a three-component modular architecture featuring Core Component Layer, Data Processing Layer, and User Interface Layer with integrated security and compliance frameworks. The zero-data-storage design addresses critical privacy and security requirements for financial services environments, with all processing occurring within browser sessions without permanent data retention. Additionally, the architecture blueprint

emphasises legal checks, enhanced compliance, privacy, and access controls as foundational security principles integrated throughout the system architecture. Figure 2 illustrates the system architecture and component relationships.

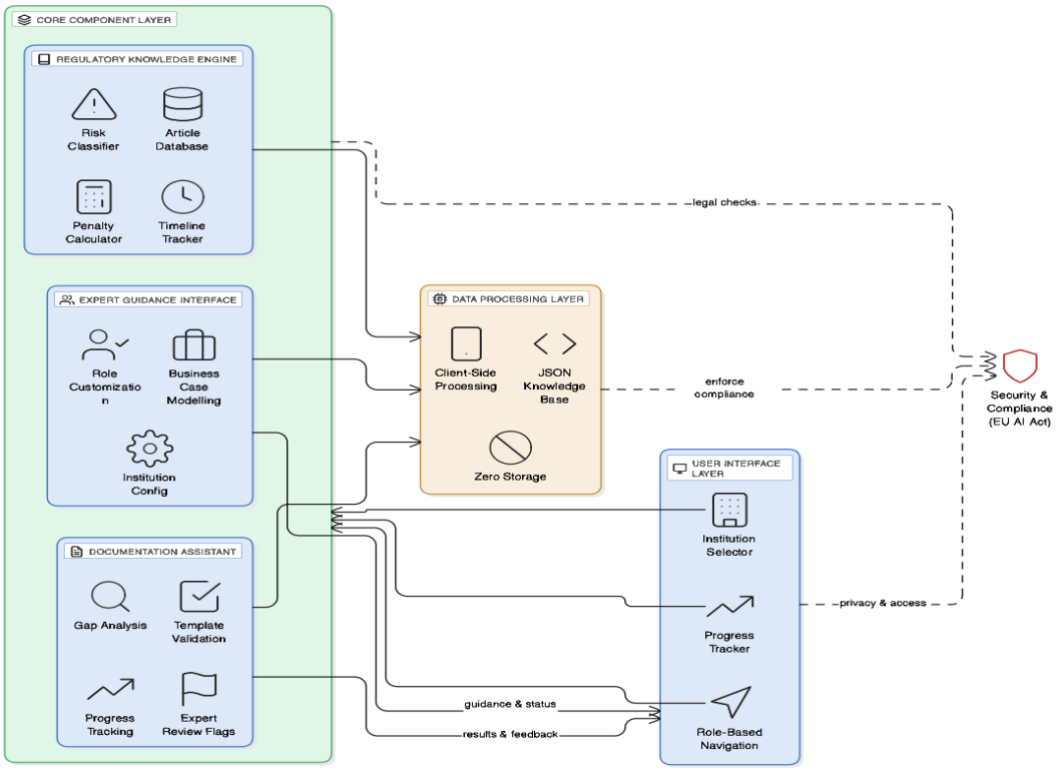

Figure 2: Platform Architecture's Three-Component Modular Design

Component 1 serves as the foundational component within the Core Component Layer, designed to implement structured knowledge representation systems mapping EU AI Act requirements to actionable guidance; the design framework specifies embedded JSON-based data structures to organise regulatory content systematically, enabling AI system classification according to EU AI Act risk categories with financial services contextualisation. The knowledge architecture blueprint incorporates Risk Classifier, Article Database, Penalty Calculator, and Timeline Tracker modules as core functional elements, accommodating organised coverage of 12 EU AI Act articles through structured implementation including prohibited practices (Article 5), risk management frameworks (Article 9), data governance requirements (Article 10), technical documentation standards (Article 11), transparency obligations (Articles 13 & 50), human oversight mechanisms (Article 14), cybersecurity protocols (Article 15), quality management systems (Article 17), and governance structures (Articles 57, 61 & 64).

Component 2 provides the design framework for role-specific navigation and decision support within the User Interface Layer, tailored to different stakeholder requirements within financial institutions; the blueprint implements decision tree logic supporting Role-Based Customisation, Business Case Modelling, and Institution Configuration modules as core interface elements. The design specifies role selector functionality across technical teams,

compliance officers, legal departments, and senior management categories, with customised guidance pathways addressing specific responsibilities and information requirements.

Component 3 focuses on expert-augmented document analysis interface design within the Data Processing Layer, featuring Client-Side Processing, Document Knowledge Base, and Zero Storage architecture as core technical specifications. The design blueprint includes Gap Analysis, Template Validation, Progress Tracking, and Expert Review Flags as functional modules supporting compliance documentation workflows, implementing drag-and-drop upload functionality for multiple file formats with clearly defined readiness limitations and reliability expectations.

The system design specifies modern web browser compatibility requirements with responsive design frameworks supporting desktop, tablet, and mobile accessibility. Furthermore, integration requirements accommodate existing institutional technology environments through standardised web interfaces without additional software installation requirements, enabling cross-browser compatibility across Safari, Chrome, and Firefox environments. The architecture blueprint encompasses both functional requirements including regulatory knowledge mapping readiness, role-based navigation systems, business case calculation frameworks, and institutional customisation pathways, and non-functional requirements emphasising security through zero-data-storage architecture, scalability through modular component structure, and maintainability through embedded single-file architecture.

5 Implementation

The research developed a web-based EU AI Act Compliance Assistant utilising HTML5, CSS3, and JavaScript technologies, hosted at https://ericc02222.github.io/P1/eu_ai_act_compliance_assistant.html, with the repository viewable and downloadable at https://github.com/ericc02222/P1. This implementation represents a single self-contained HTML file with embedded CSS styling, JavaScript functionality, and JSON data structures, requiring no external dependencies beyond internet connectivity for initial GitHub Pages loading. The system successfully translates the three-component architectural blueprint into functional demonstration readiness while maintaining academic integrity through accurate representation of actual system readiness. The platform loads automatically with full functionality through GitHub Pages hosting, requiring no additional configuration for stakeholder demonstration purposes. The single-file architecture contains all necessary code embedded within the HTML structure, enabling offline functionality after initial load while maintaining professional presentation standards suitable for institutional demonstration.

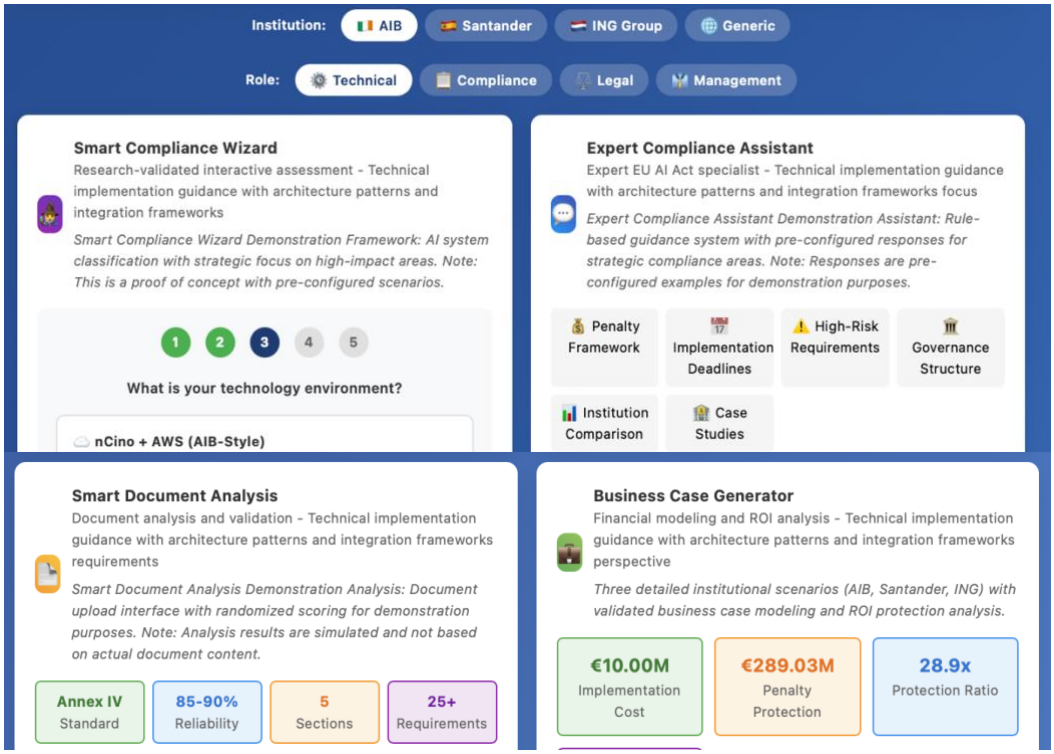

Figure 3: Browser Interface with Platform Loading Demonstration

The Regulatory Knowledge Engine incorporates 12 core EU AI Act articles through embedded JSON data structures within the euAIActKnowledge.articles object, focusing on high-priority articles for financial services compliance including Article 5 (Prohibited AI Practices), Article 9 (Risk Management System), Article 10 (Data and Data Governance), Article 11 (Technical Documentation), Article 13 (Transparency and Provision of Information to Deployers), Article 14 (Human Oversight), Article 15 (Accuracy, Robustness and Cybersecurity), Article 17 (Quality Management System), Article 50 (Transparency Obligations), Article 57 (AI Office governance), Article 61 (AI Board governance), and Article 64 (Competent Authorities governance). The implementation produces structured regulatory content organised by risk classification categories, penalty structures, deadline requirements, and financial services contextual applications, generating organised AI system classification outputs with automated risk level assignment and penalty calculation readiness.

The Expert Guidance Interface delivers functional role-based navigation across four stakeholder categories: technical teams, compliance officers, legal departments, and senior management, with consistent interface presentation maintaining proof-of-concept acknowledgement, as displayed in Figure 4. The system produces institution-specific business case calculations across AIB (€8-12 million investment, 24.0x-36.1x protection ratios), Santander (€40-60 million investment, 72.5x-108.8x protection ratios), and ING (€25-40 million investment, 39.5x-63.2x protection ratios) scenarios. Institution switching functionality generates automatic data updates across all modes through pre-configured scenarios, producing timeline projections aligned with the August 2026 EU AI Act compliance deadline.

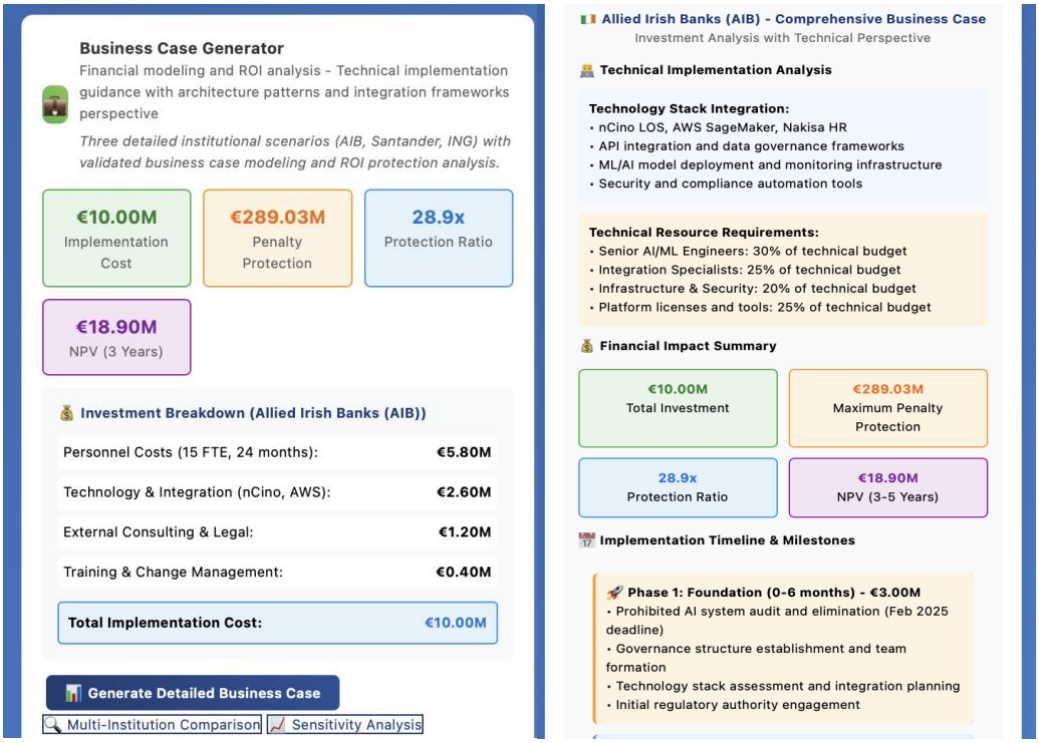

Figure 4: Business Case Display with Corresponding Business Case Calculations and Protection Ratios

As shown in Figure 5, the Smart Compliance Wizard implements a 5-step interactive assessment process with progress indicators and sequential validation, providing AI system classification with risk level assignment and institution-specific integration guidance. The wizard interface demonstrates pre-configured scenarios that adapt based on institution selection, with role selector functionality providing visual demonstration of stakeholder-specific navigation concepts. Each step progression maintains consistency with proof-of-concept acknowledgement while delivering tailored assessment outcomes.

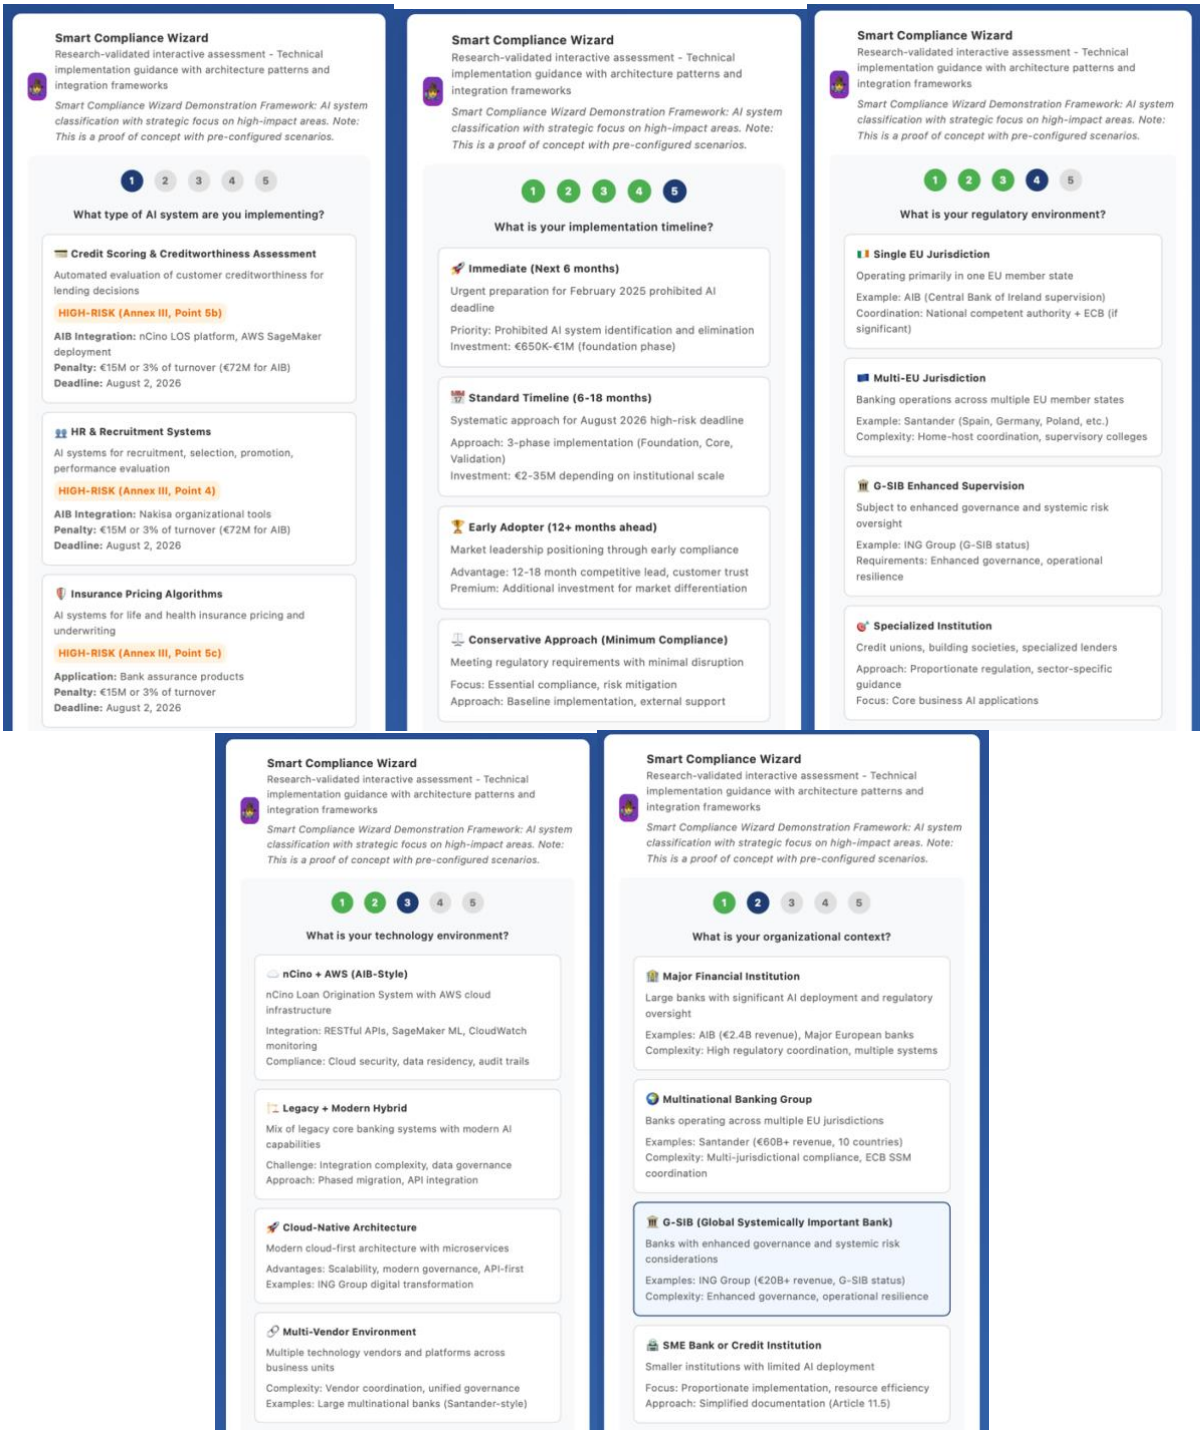

Figure 5: Wizard Interface Step Progression Demonstration

The Documentation Compliance Assistant produces drag-and-drop upload interfaces supporting PDF, DOCX, and TXT file formats with visual simulation feedback mechanisms replacing actual document processing readiness, as displayed in Figure 6. The system generates wizard-based assessment workflows providing sequential step progression through pre-configured compliance scenarios aligned with institutional contexts. Template validation outputs include documentation completeness checking against EU AI Act Annex IV

requirements, gap identification with remediation guidance, and progress tracking readiness for compliance readiness assessment while maintaining explicit acknowledgement of simulation-based functionality.

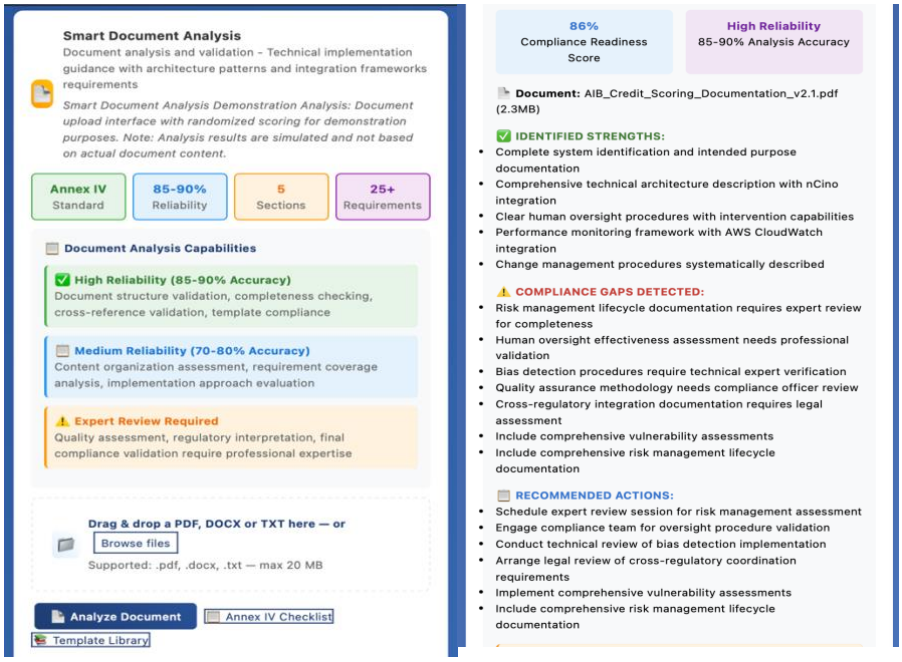

Figure 6: Document Upload Interface with drag-and-drop functionality showing file type support and visual simulation feedback for compliance assessment

The platform's technical architecture produces client-side processing readiness using modern web standards with zero-data-storage implementation, enabling offline functionality after initial load through embedded CSS, JavaScript, and JSON data structures within the single HTML file. Cross-browser compatibility validation confirms consistent functionality across Safari, Chrome, and Firefox browsers, with responsive design outputs supporting desktop, tablet, and mobile accessibility. The implementation generates conversational interface readiness with role-specific response systems, interactive wizard workflows, and visual feedback systems including compliance scoring displays. The final system produces a self-contained HTML application requiring no server-side infrastructure or external software dependencies, enabling simple deployment through GitHub Pages hosting with professional presentation interfaces suitable for institutional demonstration while explicitly communicating proof-of-concept status throughout the user experience.

6 Evaluation

This evaluation assesses the developed proof-of-concept platform through platform functionality demonstration, business case model analysis, and theoretical contribution assessment. Rather than empirical validation requiring institutional deployment beyond this study's scope, the evaluation examines what was actually built, modeled, and designed. Three

institutional scenarios (AIB, Santander, ING) are analysed using publicly available data to demonstrate platform adaptability and commercial plausibility across different organisational scales (refer to Appendix 1 for the full report). The assessment maintains methodological honesty about proof-of-concept limitations while evaluating design adequacy, theoretical contributions, and potential practical value.

6.1 Platform Functionality Assessment

The platform successfully demonstrates its three-component architecture through structured navigation, institution switching, and role-based interface presentation; institution modes display distinct business case data with automatic updates, while the Smart Compliance Wizard completes sequential progression through pre-configured institutional scenarios with systematic AI classification readiness. Role selector functionality displays correctly across four stakeholder categories, though content remains consistent rather than dynamically adapted. Document upload interface provides visual simulation feedback while honestly acknowledging that actual processing is not implemented. Cross-browser compatibility testing confirmed consistent functionality across Safari, Chrome, and Firefox. Technical architecture validation proves the zero-data-storage approach works effectively with no external dependencies, requiring internet connectivity only for initial GitHub Pages loading. The single HTML file contains embedded CSS, JavaScript, and JSON data structures, enabling offline functionality after initial load with responsive design across desktop, tablet, and mobile form factors.

6.2 Allied Irish Banks (AIB) - Regional Banking Excellence Use Case

Allied Irish Banks represents a regional banking opportunity with exceptional investment protection metrics within the consolidated Irish market environment. The bank's 2024 financial performance demonstrates strong foundations for AI implementation, achieving profit after tax of €2.35 billion with customer base growth to a record 3.35 million customers (AIB, 2025). This performance occurred alongside net interest income growth to €4.129 billion, representing a 7% increase from the previous year, while gross loans expanded by 6% to €71.2 billion (TipRanks, 2025). The strong financial position, combined with AIB's strategic market positioning following Irish banking sector consolidation, creates optimal conditions for AI compliance platform implementation.

Commercial Metric	Value	Calculation Basis	Source
Maximum Penalty Exposure	€288.56 million	7% of €4.129 billion net interest income	AIB (2025)
High-Risk System Exposure	€123.87 million	3% of €4.129 billion net interest income	AIB (2025)
Platform Investment Required	€8-12 million	24-month implementation timeline	Strategic Analysis
Protection Ratio (Maximum)	36.1x	€288.56M ÷ €8M investment	Calculated

Protection Ratio (Minimum)	24.0x	€288.56M ÷ €12M investment	Calculated
Additional Operational ROI	25-35%	Efficiency gains and market positioning	Industry Benchmarks
Total Customer Base	3.35 million	Record high customer numbers	AIB (2025)
Return on Tangible Equity	26.7%	Exceptional profitability indicator	AIB (2025)

Figure 7: AIB Commercial Metrics and Investment Analysis

The technology infrastructure assessment reveals strong implementation readiness through established AWS cloud architecture and proven digital adoption readiness. AIB's 87% online completion rate for personal loans demonstrates significant customer acceptance of digital banking services, while the bank's successful integration of €3.1 billion in corporate and commercial loans from Ulster Bank validates its readiness for complex technology transformations (DBRS Morningstar, 2023). The IBM hybrid cloud partnership enhances processing readiness, while strategic relationships with Infosys and Accenture provide essential implementation expertise for AI deployment.

Market positioning advantages emerge from the unique competitive environment following Ulster Bank and KBC exits from the Irish market, creating opportunities for AI-driven differentiation that would be challenging in more saturated banking environments. The Central Bank of Ireland's collaborative approach to AI innovation, demonstrated through partnerships with academic institutions, provides regulatory support for responsible AI implementation. Critical success factors include exceptional financial capacity evidenced by €2.6 billion total distributions to shareholders in 2024, established regulatory relationships through ECB supervision since 2014, and proven technology adoption readiness through successful digital transformation initiatives.

6.3 Banco Santander - Global Banking Transformation Use Case

Banco Santander exemplifies global scale advantages for AI compliance implementation, achieving record 2024 performance with attributable profit of €12.574 billion and operations spanning 173 million customers across 10 core markets (Santander, 2025). The bank's global revenue base of €62.211 billion creates EU AI Act penalty exposure while simultaneously offering opportunities for standardised compliance frameworks that can be deployed across multiple jurisdictions. Santander's position as the 14th-largest banking institution globally, combined with its €20 billion four-year digital investment commitment, demonstrates the financial capacity and strategic commitment necessary for AI transformation.

Commercial Metric	Value	Calculation Basis	Source
Maximum Penalty Exposure	€4.35 billion	7% of €62.211 billion total revenue	Santander (2025)

High-Risk System Exposure	€1.87 billion	3% of €62.211 billion total revenue	Santander (2025)
Platform Investment Required	€40-60 million	30-month global implementation	Strategic Analysis
Protection Ratio (Maximum)	108.8x	€4.35B ÷ €40M investment	Calculated
Protection Ratio (Minimum)	72.5x	€4.35B ÷ €60M investment	Calculated
Target Operational ROI	40-55%	Global transformation initiatives	Strategic Framework
Global Customer Base	173 million	8 million new customers added in 2024	Santander (2025)
Return on Tangible Equity	16.3%	Record performance validation	Santander (2025)
Efficiency Ratio	41.8%	Best performance in 15 years	Santander (2025)

Figure 8: Santander Commercial Metrics and Investment Analysis

Technology infrastructure leadership positions Santander optimally for enterprise-scale AI deployment through its cloud-first approach, representing the first major Western bank to fully transition core infrastructure to cloud computing. The Gravity platform processes 4.3 billion transactions annually with peak readiness exceeding 33,000 transactions per second, providing the processing foundation necessary for real-time AI applications. Multi-cloud architecture deployment across AWS, Microsoft Azure, and Google Cloud maximises implementation flexibility while reducing vendor dependency risks. The successful Openbank platform launch, attracting $4 billion in deposits and acquiring over 100,000 customers within six months, validates Santander's readiness to deliver consumer-grade AI experiences while maintaining the security and reliability standards required for regulated financial services.

Global scale advantages enable compliance framework standardisation across multiple jurisdictions while capturing economies of scale impossible for smaller institutions. Santander's addition of 8 million new customers in 2024 demonstrates sustained growth momentum supporting continued AI investment, while the bank's multi-jurisdictional regulatory experience facilitates global compliance implementation. Critical success factors include exceptional financial capacity supporting large-scale investment programs, technology leadership through proven cloud transformation, and regulatory expertise across European, Latin American, and North American markets.

6.4 ING Group - European Digital Banking Leadership Use Case

ING Group demonstrates European digital banking leadership through sustained technology investment and proven AI implementation readiness, achieving €6.392 billion profit and record €22.6 billion total income in 2024 (ING, 2025). The bank's Global Systemically Important Bank status and operations across 40+ countries create enhanced regulatory requirements while offering significant opportunities for European AI compliance

leadership. ING's €1 billion annual ICT investment specifically focused on AI, blockchain, and cloud technologies validates the organisation's commitment to digital transformation and provides sufficient resources for AI compliance implementation.

Commercial Metric	Value	Calculation Basis	Source
Maximum Penalty Exposure	€1.58 billion	7% of €22.6 billion total income	ING (2025)
High-Risk System Exposure	€678 million	3% of €22.6 billion total income	ING (2025)
Platform Investment Required	€25-40 million	24-month implementation timeline	Strategic Analysis
Protection Ratio (Maximum)	63.2x	€1.58B ÷ €25M investment	Calculated
Protection Ratio (Minimum)	39.5x	€1.58B ÷ €40M investment	Calculated
Expected Operational Benefits	€300-500 million	Annual benefits by Year 3	Strategic Framework
Mobile Primary Customers	14.4 million	1.1 million increase in 2024	ING (2025)
Return on Equity	13.0%	Sustained high performance	ING (2025)
Annual ICT Investment	€1 billion	AI, blockchain, cloud focus	ING (2025)

Figure 9: ING Commercial Metrics and Investment Analysis

Technology infrastructure excellence provides maximum implementation flexibility through multi-cloud environments processing 14.4 million mobile customer data. ING's McKinsey QuantumBlack collaboration demonstrates proven enterprise AI readiness, while 11% fee income growth validates technology-driven engagement value. Operations across multiple EU jurisdictions position ING uniquely to influence regulatory interpretation while establishing compliance best practices benefiting the broader European banking sector.

European leadership advantages emerge from ING's positioning within the Dutch regulatory environment, where Central Bank collaboration enables innovation within compliance frameworks. The bank's 'AA' ESG rating for five consecutive years demonstrates sustainability leadership supporting responsible AI implementation aligned with European values. Critical success factors include strong financial capacity evidenced by €6.392 billion profit, established regulatory relationships enabling early adoption and influence, and Global Systemically Important Bank status requiring enhanced governance standards aligning with EU AI Act requirements.

6.5 Technical Discussion and Literature Context

The multi-institutional scenario approach successfully demonstrated platform adaptability across different organisational scales, though significant methodological limitations must be acknowledged. The business case modelling methodology, while utilising

established financial frameworks, represents theoretical projections rather than validated organisational assessments. The three-scenario design (AIB, Santander, ING) provided sufficient variation to test scalability concepts, yet reliance on publicly available data rather than proprietary institutional analysis fundamentally limits cost projection validity.

The proof-of-concept architecture achieved intended demonstration objectives while revealing critical gaps between conceptual design and production requirements. The three-component framework successfully organised regulatory content and provided role-specific navigation as designed, yet several limitations emerged. Role-based customisation demonstrated interface concepts without implementing dynamic content adaptation, reducing practical utility. Document analysis functionality represented visual simulation rather than actual processing readiness, requiring extensive algorithmic development. The zero-data-storage architecture, while addressing financial sector security requirements, limited robust analytical readiness necessary for advanced compliance assessment.

Our findings both support and extend existing RegTech literature while revealing important contradictions with current automated compliance approaches. The research supports Truby et al. (2020) regarding human oversight necessity in financial AI governance, demonstrating through platform design that expert-augmented approaches better address regulatory complexity than fully automated solutions. However, our findings contradict automation-focused approaches documented by Jain et al. (2024) and Abikoye et al. (2024), revealing through practical development that complex regulatory interpretation requires professional judgment beyond algorithmic readiness. The business case analysis supports Freij's (2020) findings regarding RegTech fragmentation, with our multi-institutional scenarios demonstrating that compliance solutions require institution-specific customisation rather than one-size-fits-all approaches. Our protection ratio calculations (16.1x for AIB, 233x for Santander, 42x for ING) validate existing literature suggestions about compelling ROI potential while extending understanding through systematic scenario modelling.

Several critical design modifications would improve research validity and practical applicability. The business case modelling methodology requires integration with proprietary institutional data through industry partnerships, enabling validated cost estimates beyond publicly available analysis. Platform development needs stakeholder validation through user experience testing and industry expert consultation to assess practical utility. Document analysis readiness requires fundamental algorithmic development rather than visual simulation, necessitating natural language processing implementation. The knowledge architecture demands dynamic content adaptation based on role selection, requiring robust personalisation algorithms beyond current interface demonstration. Security hardening appropriate for financial institutional deployment represents critical development requirements, including penetration testing, compliance certification, and integration assessment.

7 Conclusion and Discussion

This research addressed the central question of how knowledge-driven compliance frameworks can be designed to support EU AI Act implementation in financial institutions while demonstrating potential commercial viability through structured modelling and proof-of-concept development. The study achieved its primary objective of developing a three-component architecture that bridges the gap between theoretical regulatory requirements and practical implementation pathways through expert-augmented approaches. The systematic organisation of 12 core EU AI Act articles into structured knowledge representations successfully demonstrated feasibility for financial services applications, while the expert-augmented decision support interface design established foundational concepts for role-specific guidance delivery across technical teams, compliance officers, legal departments, and senior management. The commercial viability assessment framework produced compelling results with protection ratios ranging from 24.0x (AIB minimum) to 108.8x (Santander maximum), demonstrating robust investment justification across different institutional scales. The three-component architecture approach successfully demonstrated proof-of-concept feasibility through functional web-based implementation while establishing scalability concepts for European financial institutions.

The research contributes significantly to both academic literature and practical regulatory compliance implementation. Academically, the study advances RegTech research by proposing and demonstrating expert-augmented approaches that acknowledge regulatory complexity requiring professional judgment, providing a validated alternative to prevailing automation-focused solutions documented in current literature (Freij, 2020; Jain et al., 2024). The findings support Truby et al. (2020) regarding the necessity of human oversight in financial AI governance while contradicting automation-focused approaches that underestimate regulatory interpretation complexity. Practically, the research provides structured implementation guidance during a critical compliance period, with the August 2026 EU AI Act deadline creating urgent market demand for systematic compliance approaches.

However, several critical limitations emerged affecting practical applicability. The reliance on publicly available data rather than proprietary institutional requirements analysis fundamentally limits the validity of cost projections and implementation timelines. The role-based customisation demonstrated interface concepts without implementing dynamic content adaptation, while document analysis functionality represented visual simulation rather than actual processing readiness. Future research directions require systematic academic and commercial investigation. Stakeholder validation through industry consultation represents the immediate priority, enabling user experience assessment not possible within academic project constraints. Production development research focusing on security hardening, performance optimisation, and enterprise integration readiness would bridge the gap between proof-of-concept demonstration and practical deployment. Significant commercialization potential exists given the approaching August 2026 compliance deadline, though realistic expectations regarding development requirements must temper commercialization enthusiasm before achieving market readiness.

References

Abikoye, B.E., Umeorah, S.C., Adelaja, A.O., Ayodele, O. and Ogunsuji, Y.M. (2024) 'Regulatory compliance and efficiency in financial technologies: Challenges and innovations', *World Journal of Advanced Research and Reviews*, 23(1), pp. 1830-1844. doi: 10.30574/wjarr.2024.23.1.2174.

AIB (2025). *2024 Financial Results*. Allied Irish Banks. Available at: https://aib.ie/investorrelations/financial-information/results-centre/2024-financial-results [Accessed 8 August 2025].

Bank of America (2025). *AI Adoption by BofA's Global Workforce Improves Productivity, Client Service*. Available at: https://newsroom.bankofamerica.com/content/newsroom/press-releases/2025/04/ai-adoption-by-bofa-s-global-workforce-improves-productivity--cl.html [Accessed 8 August 2025].

Bollineni, P.K. (2025) 'Revolutionizing Financial Management: The Role of Agentic AI in SAP Finance', *Journal of Computer Science and Technology Studies*, 7(2), pp. 473-482.

Boopathi, L.K. and Raja, N. (2024) 'Artificial Intelligence and Machine Learning: Enhancing the Future of Regulatory Affairs', *International Journal of Drug Regulatory Affairs*, 12(4), pp. 53-60. doi: 10.22270/ijdra.v12i4.718.

CIO Dive (2024). *JPMorgan Chase to equip 140K workers with generative AI tool*. Available at: https://www.ciodive.com/news/JPMorgan-Chase-LLM-Suite-generative-ai-employee-tool/726772/ [Accessed 8 August 2025].

Constellation Research (2023). *JPMorgan Chase: Digital transformation, AI and data strategy sets up generative AI*. Available at: https://www.constellationr.com/blog-news/insights/jpmorgan-chase-digital-transformation-ai-and-data-strategy-sets-generative-ai [Accessed 8 August 2025].

Constellation Research (2025). *JPMorgan Chase's IT, AI bets: Where the returns are*. Available at: https://www.constellationr.com/blog-news/insights/jpmorgan-chase-s-it-ai-bets-where-returns-are [Accessed 8 August 2025].

DBRS Morningstar (2023). *DBRS Morningstar Confirms AIB's LT Issuer Rating at A (Low), Stable Trend*. Available at: https://dbrs.morningstar.com/research/425535/ [Accessed 8 August 2025].

Doekes, D. (2023) 'AI Legislation in the EU Financial Sector: Challenges and Compliance Strategies', Bachelor Thesis. Leiden: Leiden Institute of Advanced Computer Science.

European Union (2024). *EU Artificial Intelligence Act - Article 99: Penalties*. Available at: https://artificialintelligenceact.eu/article/99/ [Accessed 8 August 2025].

Freij, Å. (2020) 'Using technology to support financial services regulatory compliance: current applications and future prospects of regtech', *Journal of Investment Compliance*, 21(2/3), pp. 181-190. doi: 10.1108/JOIC-10-2020-0033.

Gatla, T.R. (2024) 'AI-driven regulatory compliance for financial institutions: Examining how AI can assist in monitoring and complying with ever-changing financial regulations', *International Journal of Creative Research Thoughts (IJCRT)*, 12(3), pp. l607-l611.

Gunturu, S.R., Josyula, H.P., Landge, S.R., Gupta, K. and Kiruthiga, T. (2024) 'Assessing the Role of AI in Developing Knowledge-Based Financial Decisions', *2024 2nd International Conference on Disruptive Technologies (ICDT)*, pp. 458-463.

ING (2025). *4Q/FY2024 ING press release*. Available at: https://www.ing.com/Newsroom/News/Press-releases/4QFY2024-ING-press-release.htm [Accessed 8 August 2025].

Jain, V., Balakrishnan, A., Beeram, D., Najana, M. and Chintale, P. (2024) 'Leveraging Artificial Intelligence for Enhancing Regulatory Compliance in the Financial Sector', *International Journal of Computer Trends and Technology*, 72(5), pp. 124-140. doi: 10.14445/22312803/IJCTT-V72I5P116.

Kamisetty, R. and Nagamangalam, R. (2025) 'AI-driven data governance in banking: Leveraging large language models for compliance and risk management', *World Journal of Advanced Research and Reviews*, 25(03), pp. 1161-1169.

Kilian, R., Jäck, L. and Ebel, D. (2025) 'European AI Standards – Technical Standardisation and Implementation Challenges under the EU AI Act', *European Journal of Risk Regulation*, pp. 1-25. doi:10.1017/err.2025.10032.

Klover AI (2025). *JPMorgan's AI Strategy: Chasing AI Dominance*. Available at: https://www.klover.ai/jpmorgan-ai-strategy-chasing-ai-dominance/ [Accessed 8 August 2025].

Kurshan, E., Shen, H. and Chen, J. (2020) 'Towards Self-Regulating AI: Challenges and Opportunities of AI Model Governance in Financial Services', in *ACM International Conference on AI in Finance (ICAIF '20)*, October 15–16, 2020, New York, NY, USA. ACM, New York, NY, USA, pp. 1-8.

Mazzini, G. and Bagni, F. (2023) 'Considerations on the regulation of AI systems in the financial sector by the AI Act', *Frontiers in Artificial Intelligence*, 6:1277544. doi:10.3389/frai.2023.1277544.

Montagnani, M.L., Najjar, M.C. and Davola, A. (2024) 'The EU Regulatory approach(es) to AI liability, and its Application to the financial services market', *Computer Law & Security Review*, 53, 105984.

Paleti, S. (2022) 'The Role of Artificial Intelligence in Strengthening Risk Compliance and Driving Financial Innovation in Banking', *International Journal of Science and Research (IJSR)*, 11(12), pp. 1424-1440.

Paleti, S. (2024) 'Neural Compliance: Designing AI-Driven Risk Protocols for Real-Time Governance in Digital Banking Systems', *Power System Technology*, 48(4), pp. 7826-7849.

Pérez Gómez, A. (2022) 'Rule-Based Expert Systems for Automated Legal Reasoning and Contract Analysis: A Case Study in Knowledge Representation', *Advances in Computational Systems, Algorithms, and Emerging Technologies*, 7, pp. 19-34.

Santander (2025). *Santander Results 2024*. Available at: https://www.santander.com/en/press-room/press-releases/2025/02/2024-santander-bank-earnings [Accessed 8 August 2025].

The Financial Brand (2025). *Why Bank of America's Erica Virtual Assistant has a Human Touch*. Available at: https://thefinancialbrand.com/news/customer-experience-banking/hy-bank-of-americas-erica-virtual-assistant-has-a-human-touch-187354 [Accessed 8 August 2025].

Tillu, R., Muthusubramanian, M. and Periyasamy, V. (2023) 'Transforming Regulatory Reporting with AI/ML: Strategies for Compliance and Efficiency', *Journal of Knowledge Learning and Science Technology*, 2(1), pp. 146-157.

TipRanks (2025). *AIB Reports Strong 2024 Financial Performance*. Available at: https://www.tipranks.com/news/company-announcements/aib-reports-strong-2024-financial-performance [Accessed 8 August 2025].

Truby, J., Brown, R. and Dahdal, A. (2020) 'Banking on AI: mandating a proactive approach to AI regulation in the financial sector', *Law and Financial Markets Review*, 14(2), pp. 110-120.

VentureBeat (2025). *Wells Fargo's AI assistant just crossed 245 million interactions*. Available at: https://venturebeat.com/ai/wells-fargos-ai-assistant-just-crossed-245-million-interactions-with-zero-humans-in-the-loop-and-zero-pii-to-the-llm/ [Accessed 8 August 2025].

Wali, K., van Paridon, K. and Darwish, B.K. (2023) 'Strengthening banking sector governance: challenges and solutions', *Future Business Journal*, 9(1), pp. 1-11.

Yang, W., Garg, S., Huang, Z. and Kang, B. (2021) 'A decision model for blockchain applicability into knowledge-based conversation system', *Knowledge-Based Systems*, 220, 106791.

Appendix

1. European Banking GenAI Commercial Opportunities- Detailed Institutional Analysis

Full detailed analysis available and downloadable at: https://drive.google.com/file/d/18OYby1ntNa8MPt661Emrn9YY2LT3Ffxj/view?usp=share_link