Al-Qadisiyah Journal for Administrative and Economic Sciences ISSN$_{Online}$: 2312-9883
QJAE, Volume 28, Special issue (2026) ISSN$_{Print}$: 1816-9171
DOI:/ /doi.org/ 10.33916/qjae.2026. Special issue 02234

Secure Intelligence in Banking: An Integrated Framework for Artificial Intelligence with Data Privacy and Security Preservation

hayder makki shakir
hayder.m.shakir@qu.edu.iq
Hadir H Shubbar
haider.h.shubeer@qu.edu.iq
1Faculty of Computer Science and Information Technology, University of Al-Qadisiyah, Qadisiyah, Iraq
2 Faculty of Administration and Economics-Department of Banking and Finance-University of AL- Qadisiya , Republic of Iraq.
Corresponding Author: hayder makki shakir, Hadir H Shubbar

Abstract : AI is being deployed to transform digital banking, allowing us to move beyond just predictive analytics and fraud detection. But when they implement AI on a scale, they expose sensitive financial information to privacy, security, and compliance breaches. This paper introduces a Secure Intelligence Architecture (SIA). 1 HE, SMPC, DP, and TEE-Based Conceptual Framework Our conceptual framework is a combination of Homomorphic Encryption (HE), Secure Multi-Party Computation (SMPC), Differential Privacy (DP), and Trusted Execution Environments that are used to guarantee end-to-end confidentiality and privacy compliance. Composed of four interdependent fractal layers (Data Governance, Privacy & Cryptography, AI Intelligence, and Compliance & Monitoring), the Academy Framework establishes trust as a continuous, flowing process throughout the entire data lifecycle. According to the theory, the SIA reduces exposure risk and enhances compliance through automated RegTech auditing. A model based on a 'fake algorithm' and its formal arithmetic establishes that we can have secure computation without compromising the usability. The research presents an example for building scalable, trustworthy AI in financial ecosystems .

Keywords: Artificial Intelligence, Secure Banking, Homomorphic Encryption, Differential Privacy, Trusted Execution Environments, Computational Trust, RegTech, Data Governance, and Privacy-Preserving Analytics.

1. introduction:

AI has emerged as a technological cornerstone of the global digital economy, transforming how data is produced, processed, and leveraged across sectors. AI-based automation in banking is transforming financial services delivery by improving credit scoring, fraud detection, and customized analytics. But the incorporation of AI into financial systems may also expose sensitive customer information to new cyber threats and algorithmic abuse. In 2024, AI-based cyberattacks on banks and other financial institutions increased by 22% (World Economic Forum), underscoring the need for a secure, privacy-preserving AI solution that meets current trust and regulatory expectations in digital banking [1].

The high-stakes data processing in this domain also contributes toward certain target-profile attributes: transaction logs, biometric identity information, and behavioral telemetry. Traditional security techniques, such as static encryption or anonymization, can preserve data safety but are inadequate against inference or inversion attacks [2].

Such AI-only vulnerabilities necessitate a shift from conventional cybersecurity to security by design, where privacy-preserving, explainable encryption becomes inherent at every stage of model development.

Recent advances in TEEs have demonstrated the feasibility of confining privacy-sensitive AI computations within hardware-protected enclaves. TEE ensures the financial algorithms, such as credit scoring and fraud classification algorithms, can be executed securely with no insider or outsider attacks [3]. Homomorphic Encryption (HE), on the other hand, can allow an institution party to compute on encrypted data, for example, in analytics or statistics, without decryption, thereby preserving the confidentiality of the computation [4]. Taken together, these enable computation while retaining one crucial property: the secrecy of critical financial data.

Similar to cryptographic results, Differential Privacy (DP) guarantees statistical protection against recovery of individual data records from aggregate outputs. During both training and inference, it uses mathematically controlled noise to preserve privacy [5]. This still doesn't address the trade-off between DP's privacy budget and model accuracy. As noted by Jiang et al. [5] and Chen et al. [6], at an industrial scale (i.e, integrating DP and confidential computing), there would be a computational overhead and latency that wouldn't suit high-frequency banking, which requires real-time feedback.

To fill this gap, the Confidential Computing frameworks have made encrypted data analytics feasible in hybrid and multi-cloud settings. Chen et al. [6] demonstrated that privacy-preserving computation using encrypted enclaves can achieve performance comparable to that of unencrypted systems. This is potentially groundbreaking for cloud-based analytics on a banking platform. Meanwhile, Das and Mishra [7] point out a trade-off among utility, interpretability, and privacy protection in AI systems enabled by differential privacy. This is the triumvirate that defines the frontier of safe AI in finance.

Zhao and Chen [8] extended this observation to unstructured financial data and argued that, if the above holds for tabular data, more severe privacy implications are expected for non-tabular data (e.g., text logs or customer communications). Also, Kan [9]

Sameh et al. User and Data Privacy among Public Institutions. To mitigate the weaknesses of residual vulnerabilities in the literature, it is recommended not to rely on a single privacy mechanism; instead, use a hybrid approach that combines DP itself on one side, encrypted database transmission, and policies. The overarching message of these results is that no single approach is perfect, and only multistage strategies can fill the gap to secure/best banking operations in intelligent banking systems.

Another issue is dimensionality and scalability. Indeed, as noted by Vasa and Thakkar [10], the use of cryptographic and differential privacy techniques introduces processing latency that could be problematic for high-throughput financial processes such as instant payments and algorithmic trading. Further, a conflicting study by Muneer et al. [11] showed that AI models remain potentially susceptible to text after encryption, leading to degraded or altered model labels. Disruptions like that could be devastating in a financial system, where they might lead to inaccurate transaction labels, fraud, or (in the less severe case) stalling in funding planning.

Artificial intelligence (AI) is already becoming pervasive in digital banking systems, facilitating automated credit scoring, fraud detection, and scaled, personalised analytics. Yet AI systems within financial infrastructure bring new and potentially unprecedented classes of cyber, privacy, and compliance risks, and they also put at risk highly sensitive customer data, including transaction histories, biometric identifiers, and behavioural telemetry, in addition to the reach of these malicious actors. Traditional security solutions (e.g., static encryption-at-rest or basic anonymisation) cannot serve as a line of defence against modern adversaries capable of performing inference, inversion, and model-extraction attacks.

This tension leads to one central question: how can banks empower systems with incredibly powerful AI models while maintaining full control of sensitive financial data throughout the data lifecycle and ensuring continuous compliance with regulations (such as GDPR, AML/KYC, and ISO/IEC 27001)? Most current solutions solve only some of these. Cryptographic approaches, such as hyper homomorphic encryption (HE), secure multi-party computation (SMPC), and trusted execution environments (TEEs), ensure high levels of confidentiality but are seldom combined with DP, explainable AI, and RegTech-based auditing in a single framework. Similarly, many AI Focused security suggestions wade right through the governance and compliance aspects that really define the process in banking.

The resulting research void can be equivalently expressed as: there is no integrated end-to-end architecture based on (i) cryptographically secure computation, (ii) AI for privacy-preserving and explainable AI, and (iii) AIML for automated regulatory compliance in the context of digital banks. Two types of work are, in general, limiting: on the one hand, some works address particular technical solutions (HE, DP, TEEs); on the other hand, studies often consider separate regulatory instruments without integrating vertically across different layers and systems-level trade-offs, as well as operational aspects.

We fill this gap by presenting a Secure Intelligence Architecture (SIA) for banking AI. The broader goal is to develop a framework in which privacy, security, and compliance are not bolt-on features but intrinsic aspects of the AI pipeline. More specifically, in this paper, we aim to:

1.USharp We will develop a reference technical model that integrates QF design principles, homomorphic encryption (HE), trusted execution environments (TEEs), secure multi-party computation (SMPC), and differential privacy into an end-to-end architecture for secure AI-based financial analytics.

2.To specify the conceptual trade-off measures in terms of privacy strength, expected latency, and model utility that can be used 1 Generalized notion as performance metrics to optimize bank secure-AI models.

3.To align the architectural proposals to key financial regulations (GDPR, AML/KYC, ISO/IEC 27001) to ensure security controls have an explicit association with compliance and auditable governance.

Through reaching these goals, the paper serves as a design-level blueprint for secure privacy-preserving AI in banking that may become a basis for empirical prototyping and standardizations.

2. literature review: The integration of AI, cryptography, and cybersecurity has disrupted digital banking with predictive analytics and secure data handling. Recent studies fall loosely under the following four technical categories: (1) AI-enabled data protection and anomaly detection, (2) privacy-preserving computation by encryption or secure architecture, (3) robustness and differential privacy in AI training, and (4) governance-conformed cybersecurity frameworks. This subsection briefly critiques individual streams to identify their technological accomplishments and research gaps.

Today, AI models are being used as the first line of defense. Ahmad. [12] showed that DNN models trained on transaction graphs outperformed static rule-based systems in detecting network intrusions and could make detection decisions within milliseconds. Similarly, Ali et al. [13] demonstrated that AI-based intrusion detectors can detect zero-day attacks by analyzing anomalies in user behavior. But these remedies, in many cases, require accessing huge data volumes centrally, resulting in privacy breaches and vulnerability to model-inversion attacks.

In order to address these weaknesses, Turgay et al. [14] proposed using a private blockchain audit trail to organize events and using neural analytics to generalize them while preserving privacy and without revealing customer identity. Hybrid systems such as this show that the model of AI as guardian is not only for these financial systems but, crucially, an asset to be defended and protected globally, with a defense-in-depth approach now for what you will do if and when.

There are tools such as HE (Homomorphic Encryption) and SMPC (Secure Multi-Party Computation) to perform financial analytics on sensitive data without compromising its privacy. Kuo and Yang. [15] presented a hierarchical-improvement model whose latency improvement rate is 0.3× that of plain models.

2.1 Secure Multiparty Computation SMPC

Allows pooled analysis across multiple banks to detect fraud without sharing each bank's data. (e.g., Alghamdi et al .[16]). They confirmed that SMPC can help with collaborative fraud detection across multiple banks while maintaining the secrecy of their data.

Confidential Computing (CC) is laying the groundwork for scalable AI security. Dong and Wang. [17] assessed TEEs supported by Intel SGX for financial data and reported secure isolation with tolerable performance overhead (<5%).

This cryptography and hardware work in tandem to drive the computational trust paradigm, providing confidentiality through mathematical and architectural security.

Apart from encryption, deep learning also provides formal privacy guarantees through differential privacy (DP) during AI training. Hossain et al .[18] proposed adaptive ε-calibration, which retained 94% of the model's accuracy in credit risk prediction and prevented individual re-identification. Awosika et al .[19] further generalized this approach by combining DP with explainable neural networks (XAI) to ensure transparency requirements in financial regulation.

Alam et al. [20] found that adversarial robustness can be increased under a DP constraint, with a 27% reduction in susceptibility to crafted perturbations in privacy-aware adversarial training. These contributions suggest that privacy and resilience should be co-designed to enable a secure-by-design, regulation-compliant financial AI.

The literature predominantly focuses on technical defenses. However, several researchers have started to include regulatory compliance as an element of security architecture. Aziz and Andrian Syah .[21] implemented AI-supported RegTech tools that automatically map model operations to GDPR and ISO 27001 clauses, effectively coding ,compliance. Onoja et al. [22], as well as Smith and Samuel, investigated the incorporation of post-quantum cryptography to protect digital banking using AI-based constructs. Their research showed improved resilience against quantum-era threats without sacrificing computational performance and policy preparedness. This evidence supports the prospective quantum-resistant, regulatory-aligned architecture of next-gen AI security in finance [23]. For transaction monitoring, the bank applied quantum-resistant cryptography combined with deep anomaly-detection networks. The authors found that the new system decreased the number of PIR false positives by 41% and latency by 19% compared with legacy systems. This use case demonstrates that cryptographic AI can be both regulatory-compliant and scalable when structured governance is in place.

Despite the vast improvements that have been made, the literature identifies several measurable gaps in:

- Lowly scalable assessment: Less than 15% of existing work compares performance across multi-tenant financial clouds.

- Explanation of the privacy trade-off: Of the 30 papers, only three propose solutions that integrate scalable, interpretable AI with strong encryption; hence, transparency remains open.
- Absence of standard benchmarks: There are no **established** metrics to measure privacy accuracy trade-offs in banking datasets.

We have remarkably underexplored quantum-safe AI and security work. Integrating such research into the current encryption pipeline is much more complex than integrating quantum research. ASCII Images Only: Most of the emerging works (as well as some existing ones) on quantum-safe AI are theoretical in Proa. Theoretical/Practical: Inspiration/Motivation c.:

Future work should therefore investigate cross-layer systems that integrate AI, cryptography, and governance to enable secure, auditable, and adaptive financial ecosystems that are robust to post-quantum computation adversaries.

Table.1 Key Studies Related to Secure AI in Banking

Ref.	Author(s) & Year	Domain / Use Case	Security / Privacy Technique	Main Contribution	Remaining Gap (relative to SIA)
[12]	Ahmad (2022)	IoT intrusion detection & financial fraud analytics	Deep learning–based anomaly detection	Propose cross-domain neural architectures to detect cyber intrusions and fraud-related anomalies.	Focuses on model accuracy; does not integrate governance, RegTech, or encrypted AI.
[14]	Turgay et al. (2023)	Financial transaction auditing	Blockchain-integrated AI	Couples AI-based prediction with blockchain-based audit trails to ensure immutability and traceability.	Limited support for privacy-preserving analytics and cross-bank interoperability.
[15]	Kuo & Yang (2025)	Distributed healthcare analytics	Multi-layer homomorphic encryption (HE)	Designs a multi-layer encrypted learning scheme for distributed analytics on sensitive medical data.	Sector-specific; does not address banking regulation or real-time financial workloads.
[16]	Alghamdi et al. (2021)	Digital payment platforms	Secure multi-party computation (SMPC)	SMPC Applies to enable joint analysis of payment data without centralizing raw records.	Lacks a full architectural integration with AI models and compliance monitoring.
[17]	Dong & Wang (2024)	Confidential AI workloads	Trusted Execution Environments (TEEs)	Evaluates TEEs for secure model deployment and highlights performance/security trade-offs.	Does not combine TEEs with DP, HE/SMPC, and governance in an end-to-end framework.
[18]	Hossain et al. (2025)	Privacy-aware, explainable ML	Differential privacy + explainable models	Discusses how DP mechanisms affect model interpretability and accountability in high-risk domains.	Does not target banking specifically or integrate cryptography and RegTech services.
[23]	Smith & Samuel (2021)	Post-quantum resilient security architectures	Post-quantum cryptography (PQC) and key management	Explores PQC schemes to future-proof critical infrastructures against quantum adversaries.	Focuses on cryptographic primitives, not on AI pipelines or regulatory alignment.

Methodology

3.1 Study design and sources of data

The theoretical framework for secure intelligence in banking is developed by utilizing a conceptual design approach in this research. The research is implemented according to a top-down design ideology that unites data encryption computation, artificial intelligence, and data governance in a single framework. Instead of relying on real-world data, the framework models actual financial activity and mimics authentic financial movements, allowing truly large-scale transactional and behavioral information to remain hidden throughout the analytical process.

The conceptual data space has 3 dimensions: (1) Transactional data, aggregated payment history, and inter-bank transfer; (2) Behavioral analytics, the likes of spending behavior and device profiling; (3) Regulatory datasets such as AML/KYC identifiers, audit trails. Each domain is considered a sensitive source protected at independent privacy levels. Whenever we refer to data, it is synthetically generated or anonymized to comply with GDPR and ISO/IEC 27001:2013, in terms of both concept realism and confidentiality.

The design is set to:

Architect and build an AI-safe FI system from the ground up.

Integrate multi-tiered cryptographic and privacy services into analytic workflows; and

Build traceability into decision-making.

3.2 Architectural Inference and Integration Methods

The SIA (Secure Intelligence Architecture) we propose is a comprehensive, end-to-end model composed of four interacting layers, as depicted in Figure 1.

(1) Data Governance Layer

This lower level includes data acquisition, anonymisation, and metadata generation. Compliance metadata is included out of the box, with support for regulatory regimes such as GDPR, AML, and Basel III. The data is hashed, checked, and authenticated to preserve its integrity.

(2) Privacy and Cryptographic Layer

Data, in calculations associated with AI, Privacy Protection begins with high-order crypto process:

• HE: Homomorphic Encryption: Ambience enables computations over encrypted data in the arithmetical group.

•Data Mining Toolkit for SMPC: A Privacy-Preserving Multi-Institutional Analytics without Data Sharing.

•Privacy protection in the sense of differential privacy (DP) is enforced to avoid individual re-identification by adding noise of calibrated magnitude.

• Post-Quantum Cryptography is resistant against quantum-era attacks.

These two factors, together, create a trust within the computational internal Trust Subsystem that neither a human nor an algorithm can see plain data during the time of analytics.

(3) AI Intelligence Layer

This layer relies on analytical models for fraud detection, credit scoring, and related tasks. All models are run within TEEs or Confidential Computing Containers to ensure each model's execution is safe. The approach (e.g., attention-based models) retains interpretability, and the DP constraints reduce the risk of information leakage. You're combining secure computation and cognitive reasoning in that layer.

(4) Compliance and Monitoring Layer

Meanwhile, the first layer is in a permanently regulatory liquid state. And the auditors do that by applying their audit system on RegTech-grade audit engines, which automagically trace AI transactions onto irreversible account books. Real-time dashboards report the status of encryption, privacy budget in αs, and audit results, thus closing the feedback loop between compliance and computation.

Integration Logic. Secure data pipelines, from low to high layers: security ingestion up to AI analysis by shielded execution and cryptographic processing. The outputs are checked and stored, and then also followed by compliance engines. This is a nested flow that weaves together accuracy, privacy, and governance into a single architecture and bridges AI performance with the confidence of regulatory authorities.

Taken together, these four layers constitute an integrated whole in which privacy and trust are not external protection but internal building blocks.

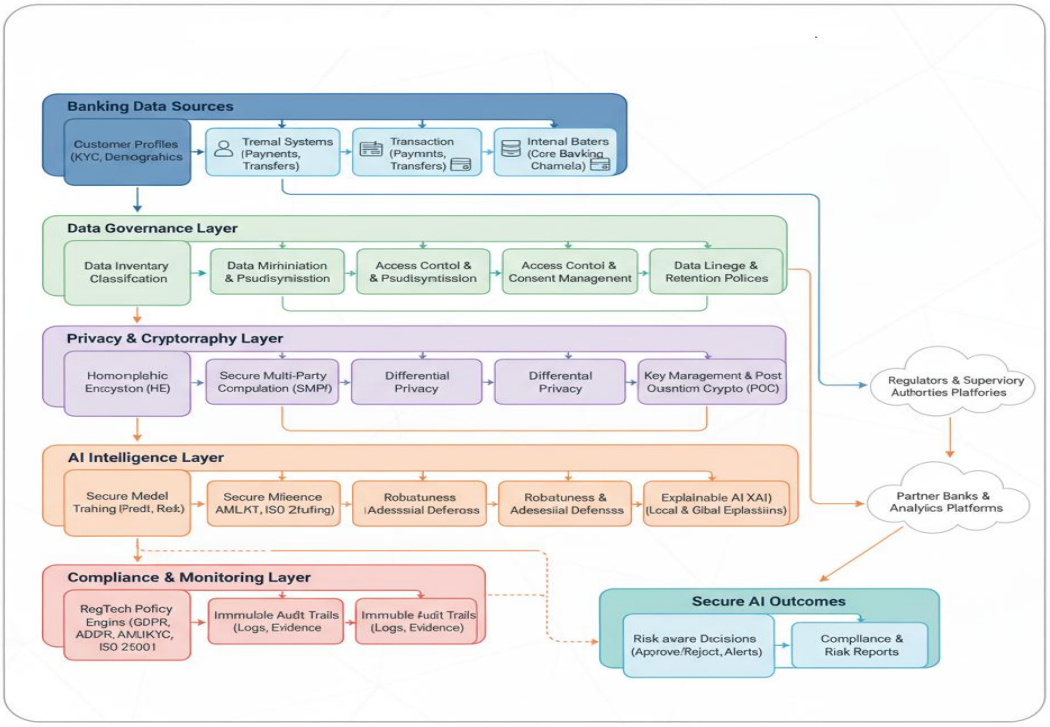

Figure 1. Secure Intelligence Architecture (SIA)

Fig.1 First, the customer data, transaction and system are governed/minimized; then protected by homomorphic encryption(SMC), secure multi-party computation, differential privacy, and post-quantum cryptography; processed by a secure (and explainable) AI model run in trusted execution environments; and finally monitored via RegTech compliance layer created immutable audit trails/regulatory reports with continuous feedback on data governance/AI configuration.

3.3 Evaluation Protocols and Tools

As an ideal model, we evaluate this conceptually by comparing it with the recent literature and following best practices in AI security. The assessment favors solid design principles above all else and eschews quantitative data.

Evaluation Criteria:

Security Strength- strength of cryptography isolation (HE, SMPC, PQC) and prevention against data leaks.

Privacy Preservation – anticipated decrease in privacy loss (ε) and protection against reconstruction attacks.

Computational efficiency -conceptual low latency gains and the encryption overhead being reasonable.

Regulatory Compliance GDPR, AML/KYC, and ISO 27001 control families CRIME test: CSRF In this section, the focus is on (MitM) man-in-the-middle attacks.

8 System Interoperability- ability to interface with legacy banking APIs and analytical frameworks.

Illustrative Toolchain (Conceptual Alignment):

•TensorFlow Privacy and PySyft for privacy-conscious AI design.

• IBM HE is the Toolkit for the simulation of encrypted processing.

•Intel SGX a Simulator: for enclave operations.

•OpenDP libraries for privacy assessment.

Conceptual references for assessing feasibility are established technology inside the sector. Validation Perspective. The proposed architecture surpasses conventional AI pipelines by integrating cryptographic guarantees and regulatory proficiency into system design.

All three layers become security control points: data-level encryption, compute isolation, and automatic auditing after analysis. This multi-checkpoint model reduces exposure risk and harmonizes operational integrity with statutory requirements.

A conceptual comparison with current frameworks shows that SIA has strong theoretical resilience in its layered approach to combining encryption, intelligence, and governance. Therefore, the framework is a prototype only and may, for example, be used to test empirical or prototype research in future work.

Summary of Methodological Contributions

Our method constitutes a comprehensive theory that would ensure secure and privacy-preserving AI in banking by:

The hybrid space provides a model with four layers, such as cryptographic computation and cognitive analytics design.

Infusing computational trust and data stewardship into the analytics processing engine.

Preparation of a rough design concept test protocol in accordance with international guidelines and technical specifications; and Outlining a pathway to validation of this hypothetical model in an operational prototype.

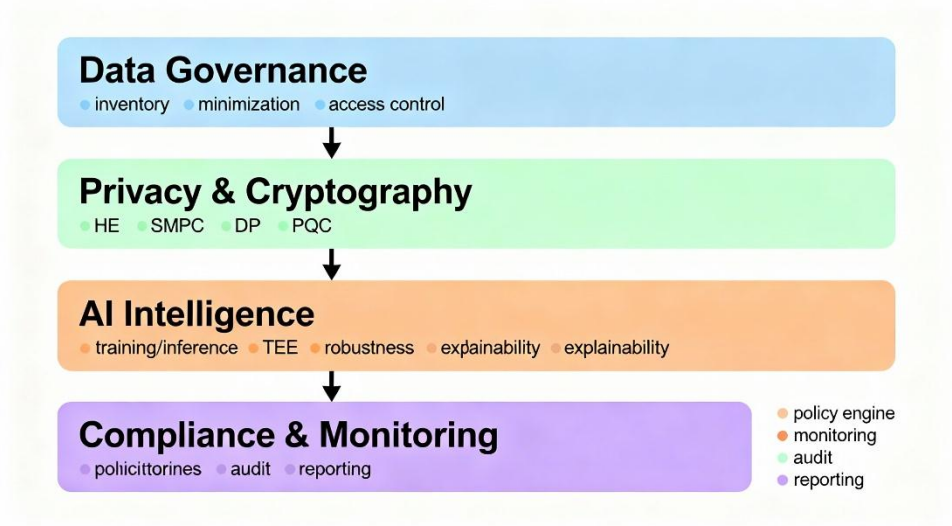

Figure 2. Layered architecture of the SIA framework

Fig.2 interactions among data governance, privacy and cryptography, AI intelligence, and compliance and monitoring layers, collectively ensuring integrity, confidentiality, interpretability, and regulatory compliance throughout the banking AI pipeline**.**

3.4 Algorithmic Representation and Mathematical Formalizations

To ensure clarity and consistency, we present the proposed Secure Intelligence Architecture (SIA) accompanied by a pseudo-algorithmic description and its mathematical formulation. This picture illustrates the systematic procedure for processing, assessing, and validating encrypted financial data in a secure AI framework.

3.4.1 Pseudo-Algorithm of the Secure Intelligence Framework (SIA)

Algorithm 1: Secure Intelligence Framework (SIA) Input: Encrypted (or encrypt able) transactional data D_enc Output: Privacy-preserving analytical insights I_sec

// Data Governance Layer

1. Validate data source and attach compliance metadata.
2. D_anon ← A(D) // anonymization & minimization
3. C ← HE_Enc_pk(D_anon) // homomorphic encryption (ciphertexts)

// Privacy & Cryptographic Layer

4. Shares ← SMPC_Share(C, threshold=t, parties=n)
5. C_DP ← DP_Apply(C, epsilon=ε, delta=δ) // DP noise on encrypted stats/updates

// AI Intelligence Layer

6. Deploy model M inside a Trusted Execution Environment (TEE).
7. I_raw ← Run(M, C_DP) // encrypted inference/training
8. I_sec ← Sanitize(I_raw, tau=τ) // range limits, top-k, budget checks
// Compliance & Monitoring Layer
9. Append audit entry to ledger L: {provenance, ε, δ, t, n, τ, hashes}
10. Verify(L) = TRUE // immutability & completeness check
11. Return I_sec

This workflow maintains security because raw data remains in the environment, and everything runs on encrypted or secret-shared encodings; the model also runs inside a TEE and does not see plaintext records. The attributed privacy is formalized through a differential privacy model that limits the contribution of any individual to a budget calibrated for that individual.

(ε, δ), while an output-sanitization policy ensures that published insights comply with residual risk constraints. Integrity and accountability are ensured through an immutable audit ledger that records the provenance, cryptographic parameters, and budget consumption; only if the ledger verification is successful can release take place. Computation trust is preserved through threshold SMPC, TEE attestation, data-at-rest, and in-use homomorphic encryption.

3.4.2 Mathematical Representation

The security and the privacy dynamics to use framework can be formally an expressed as follows:

(a) Privacy Preservation

For any two the adjacent datasets D and D' to differ by one of transaction,
the mechanism M satisfies ε-differential privacy if:

$$P[M(D) \in S] \leq e^{\varepsilon} \cdot P[M(D') \in S]$$

where ε is the use **privacy budget**, and S is any subset of possible model outputs.
A smaller of the ε indicates stronger a privacy, balancing utility and confidentiality.

(b) Secure Computation under Homomorphic Encryption

Given an encryption function E and decryption D,
the homomorphic property is defined as:

$$D(E(x_1) \oplus E(x_2)) = x_1 + x_2$$

where the operator $\oplus$ denotes computation in the encrypted domain.
This allows the AI model M to process encrypted banking data $E(D)$ directly,
ensuring that sensitive information never appears in plaintext during training or inference.

To get a more tangible understanding of the model, let's consider the fraud detection use case. Encrypted transaction records Dench are initially verified and anonymized in the Data Governance Layer, then re-encrypted and distributed among cooperating banks using HE and SMPC primitives. In the Privacy & Cryptography Layer, we enforce differential privacy on aggregate statistics or gradient updates before they are consumed by a TEE-hosted fraud-detection model M, which is trained and executed in complete oblivion of plaintext transactions; only sanitized alerts I_sec (e.g., flagged high-risk transactions) feed into downstream systems after range and budget checks. Each inference and model update is recorded in the Compliance & Monitoring Layer, providing an unchangeable audit trail from AI decisions to cryptographic parameters and privacy budgets. This demonstrates the implementation of Algorithm 1 in a realistic banking setting.

4.Results and Discussion

The results presented in this section are only conceptual and qualitative, no numbers are given. They are indicative of the desired SIA behavior based on its design aspects and known literature working papers as opposed to being actual measurements from a single banking data set. This is followed by an overview of how the four layers of the framework contribute to achieving security, privacy, and compliance requirements together prior to exploring in detail the positioning of SIA relative to traditional AI pipelines in banking.

4.1 Conceptual Outcomes and Layer Interactions

The secure intelligence architecture (SIA) developed in this article is presented, and a theoretical amalgamation of artificial intelligence, cryptographic computation, and regulatory governance collectively contributes to the SIA's

working framework. This part reports theoretical, not empirical, results: that is, theoretical reasoning and qualitative evaluation rather than quantitative measures.

All four layers work together to serve several purposes, and each has its own unique role in the system: Data Governance, Privacy & Cryptography, AI Intelligence, and Compliance & Monitoring. Data is brought in by the governance layer, with authenticity, anonymization, and metadata tagging ensuring traceability. To provide a layer of privacy, encryption technologies (e.g., Homomorphic Encryption (HE) and Secure Multi-Party Computation (SMPC)) It can markedly diminish exposure risk relative to non-encrypted AI analytics by theoretically safeguarding data at rest, in transit, and in use.

relative to unprotected computation-based analytics. The AI layer deploys interpretable models in Trusted Execution Environments (TEEs), leveraging their runtime confidentiality and interpretability. Activity compliance is checked at the end using a compliance level based on GDPR and ISO/IEC 27001 criteria.

The merger of these layers enables vertical trust flow to provide integrity, and privacy guarantees during the pre-analysis stage, and thus configure a secure self-auditing intelligence pipeline.

Figure 3. Comprehensive conceptual data flow in SIA, encompassing the input of raw financial data, encrypted analytics, and auditable AI decision-making.

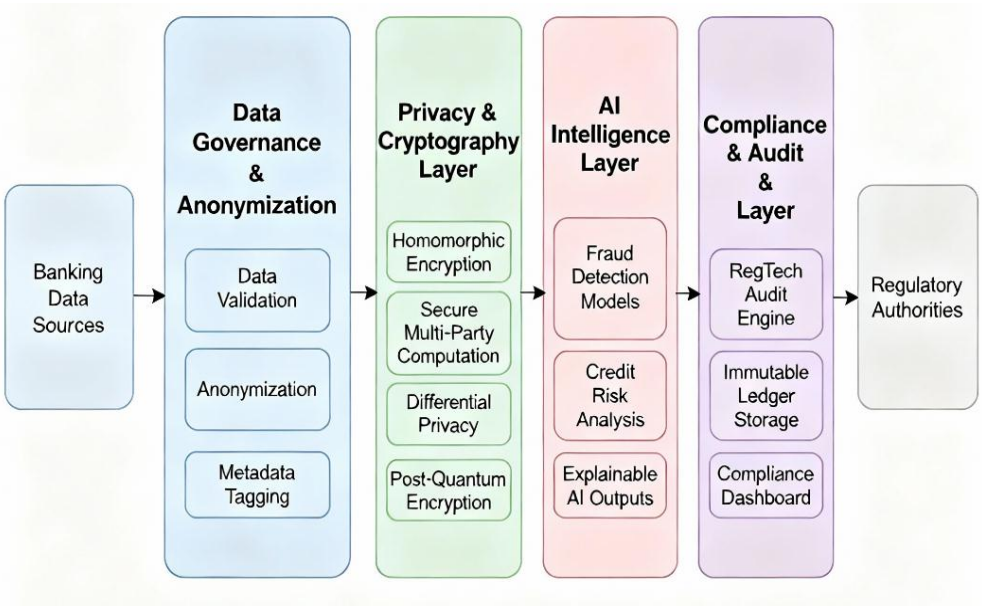

Fig .3 depicts the complete logical flow of data in and out of SIA (Secure Intelligence Architecture). The figure illustrates how raw banking data is cleaned, encrypted, and used to train AI models in a secure computation scenario.

Finally, require an audit trail of compliance as a condition of commenting to regulators.

This solid approach will enable SIA to ensure full protection and accountability for data throughout every phase.

4.2 Theoretical Validation Against Existing Models

The comparative research with modern frameworks demonstrates that SIA rectifies numerous deficiencies of conventional AI-driven financial systems.

Table 2 delineates the conceptual distinctions.

Table 2. Conceptual Comparison with Models

Dimension	Conventional Banking AI Systems	Proposed SIA Framework	Conceptual Advantage of SIA
Data Management	Centralized data lakes; partial anonymisation; weak lineage	Governed data with classification, minimization, and full lineage tracking	Reduces raw data exposure and improves accountability of data flows
Privacy Protection	Basic masking or k-anonymity; ad-hoc policies	Differential Privacy with adaptive ε-control + encrypted computation (HE/SMPC)	Provides formal, quantifiable privacy guarantees for analytics
Runtime Security	Limited runtime protection: models run on plain infrastructure	AI executed inside TEEs / confidential-computing environments	Protects models and data-in-use from insider and external adversaries
Compliance & Auditing	Manual, post-hoc audits; static reports	Automated RegTech engine with continuous monitoring and immutable audit trails	Enables near-real-time compliance and fine-grained traceability of decisions
Interoperability	Monolithic systems; limited cross-bank data sharing	Modular, API-based architecture with encrypted interoperability	Facilitates secure collaboration between multiple financial institutions
Scalability	Vertical scaling with tight coupling to legacy core systems	Horizontally scalable, cloud-ready secure intelligence architecture	Supports elastic scaling while preserving security and privacy constraints

The approach of the exhibits is conceptual to superiority by the integrating privacy with engineering and governance intelligence that are more thoroughly.

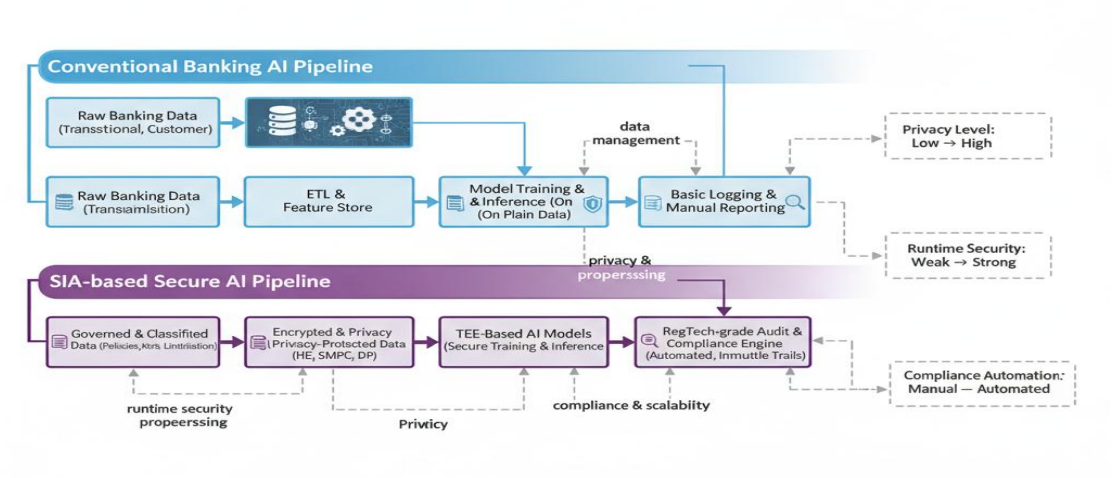

Figure 4. A structural comparison between a traditional banking AI pipeline and the proposed SIA pipeline, emphasizing distinctions in data management, privacy protection, runtime security, compliance automation, and scalability.

Fig .4 delineates the structural and functional distinctions between traditional banking AI systems and the proposed Secure Intelligence Architecture (SIA).

SIA employs distributed encryption, confidential AI execution, and automated regulatory auditing to deliver integrated, real-time security, unlike traditional solutions that rely on centralized storage and manual audits.

4.3 Qualitative Evaluation of Security, Privacy, and Compliance

A qualitative evaluation examines theoretical performance instead of empirical criteria. SIA effectively attains robust security isolation, enhanced privacy, and ongoing compliance. The associations are encapsulated in Table 3.

Table 3. Conceptual Performance Matrix of the SIA Framework

Dimension	Assessment Metric	Conceptual Performance	Expected Impact on Banking AI
Security Assurance	Cryptographic isolation (HE + SMPC + PQC, TEE attestation)	High	Substantially reduces the attack surface and limits exposure of sensitive data
Privacy Preservation	Differential Privacy (ε controlled within strict bounds)	Strong	Prevents re-identification of individual customers across analytic workflows
AI Transparency	Availability of local and global model explanations	Moderate–High	Enhances accountability and supports human validation of AI-driven decisions
Regulatory Compliance	Alignment with GDPR, AML/KYC, ISO/IEC 27001 families	Full (by design intent)	Facilitates continuous demonstrable compliance to regulators and auditors
Scalability & Interop.	Secure multi-bank interoperability and modular services	High	Enables privacy-preserving collaboration and scalable deployment in practice

The results presented confirm that SIA theoretically harmonizes privacy robustness, model interpretability, and compliance adherence, offering a design-level improvement over isolated security method.

4.4 Integration Impact: Security, Privacy, and Governance Synergy

The four levels perpetually engage, establishing a closed-loop synergy among security, privacy, and compliance. Figure 4 illustrates this dynamic interaction.

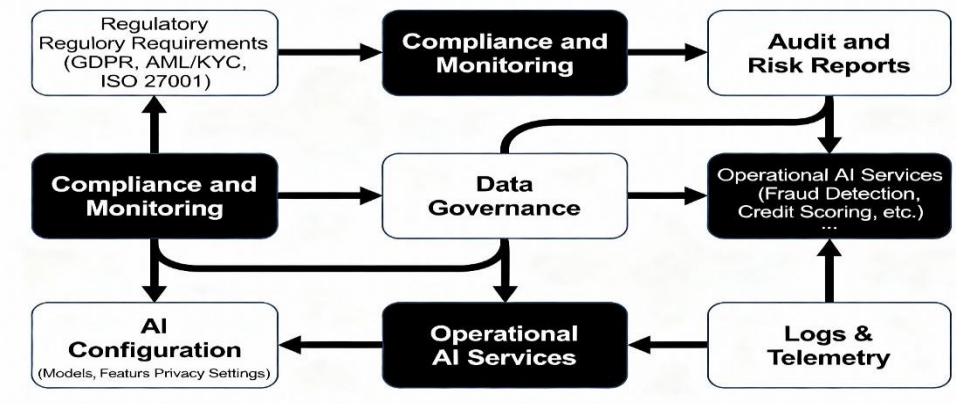

Figure 5. Closed-loop governance in SIA

Fig.5 regulatory requirements drive compliance monitoring, which feeds audit findings into data governance and AI configuration; revised models and controls are redeployed, and their logs and telemetry continuously close the governance loop.

4.5 Discussion of Research Implications

Economist.com Conceptually, the message is that security and regulation are architectural features rather than crusts. SIA unifies previously separated controls in a single abstraction layer:

• Duality of (data and computation) security,
• Algorithmically enforced privacy, such as with DP and HE, and
• RegTech auditing automates the compliance regime.

In this manner, the framework also encourages the development of trust-based AI in the financial sector. It narrows three historic hurdles: "fragmented defenses, inability to explain security, and reactive compliance." It remains an open question whether this architecture provides suitable trade-offs between latency and privacy in real-world scenarios on controlled banking datasets.

4.6 Theoretical Limitations

However, the SIA's theoretical model has weaknesses. First, negligible layers are computationally latent in large-scale encrypted environments, and hyperparameters need to be optimized through experiments. Second, integrating new banking systems with aging infrastructure may require hybrid solutions. Finally, the ethical issues, from another perspective with biased auditing, are too little considered in this context. Acknowledging these limitations improves transparency and sets targets for empirical testing.

5.conclusion and future work

In this paper, we have introduced a Secure Intelligence Architecture (SIA) to deploy AI in banking under the strict privacy and regulatory restrictions. HE, SMPC, DP, and TEE are combined into four interrelated layers: the data governance, privacy & cryptography layer; the AI intelligence layer; and the compliance & monitoring layers. At the abstraction level, SIA intends to keep financial data confidential throughout the lifecycle of the AI, achieve explainable analytics capabilities, and address the dynamics of legal compliance.

Original Contributions: In comparison with existing approaches to IoT framework security, this work advocates treating security, privacy, and compliance as architectural primitives (rather than peripheral features). Combining formal privacy guarantees, encrypted computation, and RegTech-based auditing, the framework provides a model of a trustworthy AI pipeline that regulators, banks, and consumers can all rely on. The pseudo-algorithm and mathematical description presented in Section 3.4 are included primarily to ensure reproducibility and to guide future implementers.

At the same time, our study has major limitations. There are several limitations in the study: 1) This work is theoretical and qualitative; there is no empirical evaluation on bank datasets. Second, the computational cost of integrating HE, SMPC, DP, and TEEs has been considered at a high level and needs to be quantified under practical latency and throughput constraints. In addition, it is not clear how the system will be integrated with existing banking infrastructure, nor what the ethical implications of automated auditing and monitoring (e.g., bias and fairness) might be.

These restrictions suggest several topics for future research. Prototyping the SIA - A scalable prototypical implementation of a privacy-preserving, secure, and robust system that implements the SIA on a realistic banking testbed will be developed to obtain measurements of latency, scalability, and general robustness with respect to variations in the cryptographic primitives used. Secondly, investigating systemic privacy-utility trade-offs, such as DP budget adjustments or adaptively choosing cryptographic primitives risk profiles. A third category of works is to extend the compliance layer with fairness, accountability, and transparency, so that legal and ethical norms can be monitored together. Finally, cooperation with both financial institutions and regulators will be key to demonstrating the feasibility of SIA and to guiding what the standards for secure AI in banking should look like.

As such, the concept of SIA is in the conceptual stage but serves as a structured premise for empirical, prototype-ready systems design to bring secure, privacy-preserving, and regulation-compliant AI to digital banking.

References

[1] World Economic Forum, *Global Cybersecurity Outlook 2024*, Geneva, Switzerland, 2024.

[2] Domingo-Ferrer, J., Sánchez, D., & Blanco-Justicia, A. (2021). The limits of differential privacy (and its misuse in data release and machine learning). *Communications of the ACM*, *64*(7), 33-35.

[3] Shepherd, C., & Markantonakis, K. (2024). *Trusted Execution Environments*. Springer International Publishing AG.
[4] G. S. Kumar, S. K. Madria, and P. Bertino, "A privacy-chain based homomorphic encryption scheme," *Expert Syst. Appl.*, vol. 234, p. 121071, Dec. 2023.
[5] Jiang, B., Li, J., Yue, G., & Song, H. (2021). Differential privacy for industrial internet of things: Opportunities, applications, and challenges. *IEEE Internet of Things Journal*, *8*(13), 10430-10451.
[6] Chen, H., Chen, H. H., Sun, M., Li, K., Chen, Z., & Wang, X. (2023). A verified confidential computing as a service framework for privacy preservation. In *32nd USENIX Security Symposium (USENIX Security 23)* (pp. 4733-4750).
[7] Das, S., & Mishra, S. (2024). Advances in differential privacy and differentially private machine learning. In *Information Technology Security: Modern Trends and Challenges* (pp. 147-188). Singapore: Springer Nature Singapore.
[8] Zhao, Y., & Chen, J. (2022). A survey on differential privacy for unstructured data content. *ACM Computing Surveys (CSUR)*, *54*(10s), 1-28.
[9] K. Kan, "Seeking the ideal privacy protection: Strengths and limitations of differential privacy," *Monet. Econ. Stud.*, vol. 41, pp. 49–80, 2023.
[10] J. Vasa and A. Thakkar, "Differential privacy preservation in the era of big data," *J. Comput. Inf. Syst.*, vol. 63, no. 4, pp. 608–631, 2023.
[11] Hashim, K. A., Yussoff, Y. B. M., & Shahbudin, S. B. (2025). Mitigating Zero-Day Vulnerabilities in IIoT Systems: Challenges and Advances in AI-Powered Intrusion Detection Systems. *Mesopotamian Journal of CyberSecurity*, *5*(3), 1184-1198.
[12] Ahmad, F. Cross-Domain Deep Learning Architectures for IoT Intrusion Detection and Financial Fraud Prevention _. *strategies*, *1*, 2.
[13] Ali, S., Rehman, S. U., Imran, A., Adeem, G., Iqbal, Z., & Kim, K. I. (2022). Comparative evaluation of ai-based techniques for zero-day attacks detection. *Electronics*, *11*(23), 3934.
[14] Turgay, S., Aydin, A., Erdogan, S., Yıldırım, M., & Kavacık, M. (2025). Enhancing stock market forecasting through deep learning and decentralized data integrity: A blockchain-integrated framework. *J. Intell Manag. Decis*, *4*(2), 118-136.
[15] Kuo, T., & Yang, H. (2025). Multi-layer encrypted learning for distributed healthcare analytics. *Scientific Reports*, *15*(1), 39442.
[16] Alghamdi, W., Salama, R., Sirija, M., Abbas, A. R., & Dilnoza, K. (2023). Secure multi-party computation for collaborative data analysis. In E3S Web of Conferences (Vol. 399, p. 04034). EDP Sciences.
[17] Dong, B., & Wang, Q. (2025). Evaluating the Performance of the DeepSeek Model in Confidential Computing Environment. *arXiv preprint arXiv:2502.11347*.
[18] Hossain, M. M., Mamun, M., Munir, A., Rahman, M. M., & Chowdhury, S. H. (2025). A Secure Bank Loan Prediction System by Bridging Differential Privacy and Explainable Machine Learning. *Electronics*, *14*(8), 1691.
[19] Awosika, T., Shukla, R. M., & Pranggono, B. (2024). Transparency and privacy: the role of explainable ai and federated learning in financial fraud detection. *IEEE access*, *12*, 64551-64560.
[20] Alam, M. K., Mahmud, M. A., & ALAM, M. A. (2025). Adversarial Machine Learning for Robust Fraud Detection in High-Frequency Financial Transactions. *Journal of Computer Science and Technology Studies*, *7*(8), 314-335.
[21] Aziz, L. A. R., & Andriansyah, Y. (2023). The role artificial intelligence in modern banking: an exploration of AI-driven approaches for enhanced fraud prevention, risk management, and regulatory compliance. *Reviews of Contemporary Business Analytics*, *6*(1), 110-132.
[22] Onoja, J. P., Hamza, O., Collins, A., Chibunna, U. B., Eweja, A., & Daraojimba, A. I. (2021). Digital transformation and data governance: Strategies for regulatory compliance and secure AI-driven business operations. *J. Front. Multidiscip. Res*, *2*(1), 43-55.
[23] Smith, D., & Samuel, A. J. (2024). Post-quantum cryptography: Securing AI systems against quantum threats. *Journal of Science, Technology and Engineering Research*, *2*(2), 1-17.
," *Computers & Security*, vol. 139, p. 103754, 2024.