JRIIE Journal of Research Innovation and Implications in Education
Center for Research Implications & Practice

Website: www.jriiejournal.com ISSN 2520-7504 (Online) Vol.10, Iss.1, 2026 (pp. 450 - 462)

Leveraging Homomorphic Encryption to Enhance Efficiency and Data Protection in Online Retail Based in Kenya

Mukolwe N Lucy, Philemon Kittur and Francis Musembi
Department of Mathematics and Computer science, University of Eldoret, Kenya.
Email; lucymukolwe@gmail.com

Abstract: *The rapid growth of online retail businesses has increased reliance on artificial intelligence (AI)-driven systems for personalization, fraud detection, and operational efficiency in customer interaction. These algorithmic systems require access to vast amounts of sensitive customer data, positioning data as both a driver of commercial value and a source of digital vulnerability. In many cases, this data is processed in plain text, creating significant security risks especially in dynamic and cloud-based retail environments. This exposes critical shortcomings in traditional encryption methods and raises fundamental concerns about digital rights, trust, and data governance in the algorithmic age. This study explores the integration of Partial Homomorphic Encryption (PHE) as a means of reinforcing data privacy without sacrificing functionality. PHE allows for computations on encrypted data, eliminating the need for decryption during processing and thereby preserving customer confidentiality. The research focuses on current data security practices in online retail businesses, with particular attention to Kenyan-based enterprises, identifying key vulnerabilities in the handling of sensitive customer information. It further examines the limitations of existing traditional encryption approaches and presents a theoretical application of PHE within AI models deployed by online retailers. Through a literature review, the study assesses the performance, feasibility, and implications of adopting Homomorphic Encryption in cloud-based retail systems. The findings reveal that although PHE introduces computational overhead, this research argues that embracing PHE can help redistribute power over personal data, reinforce consumer trust, and uphold digital rights in the increasingly algorithmic architecture of online retail systems.*

Keywords: *Homomorphic Encryption, Data privacy, Cloud-based systems, Online retail, Partial Homomorphic Encryption.*

How to cite this work (APA):

Mukolwe, L. N., Kittur, P. & Musembi, F. (2026). Leveraging Homomorphic Encryption to Enhance Efficiency and Data Protection in Online Retail Based in Kenya. *Journal of Research Innovation and Implications in Education, 10*(1), 450 – 462. **https://doi.org/10.59765/goa94.**

1. Introduction

The global shift towards digital commerce has fundamentally reshaped the retail industry, with online platforms becoming increasingly dependent on Artificial Intelligence (AI) to drive competitive advantage. Modern e-commerce businesses now rely heavily on AI-powered tools for personalized product recommendations, dynamic pricing strategies, targeted marketing campaigns, and sophisticated fraud detection systems. These capabilities require the continuous collection, storage, and analysis of vast amounts of customer data, including detailed purchase histories, payment credentials, browsing patterns, and behavioral insights (Kumar et al., 2020). The more accurate and responsive these AI systems need to be, the more granular and sensitive the data process becomes.

While traditional encryption methods such as Advanced Encryption Scheme (AES) or Rivest Shamir Adleman (RSA) are highly effective for securing data when it is stored ("at rest") or being transmitted over networks ("in transit"), they provide no inherent protection during the processing stage. In current AI workflows, data must often be decrypted before algorithms can analyze it, temporarily exposing sensitive information in plain text form within system memory. This momentary vulnerability, though brief, creates a significant risk in dynamic, cloud-based retail environments where multiple systems, APIs and third-party integrations are involved (Acar et al., 2018). In such settings, even a small breach window can be exploited by malicious actors to intercept payment details, steal personal information, or manipulate transaction records.

In Kenya, the rapid growth of online retail has been fueled by high mobile penetration, expanding internet access, and widespread adoption of mobile payment systems such as M-Pesa. Yet, many local e-commerce platforms ranging from large-scale marketplaces to small social-media-based sellers have not yet implemented advanced encryption protocols capable of protecting data throughout its lifecycle. The limitations of current practices have been starkly illustrated by recent breaches, such as the 2023 Naivas supermarket ransomware attack, which compromised over 600GB of customer data including names, contacts, and transactional records (The Informer, 2023). Such incidents not only expose the affected customers to risks like identity theft and financial fraud, but also erode public trust, damage brand reputation, and carry the potential for severe regulatory penalties.

This situation highlights a fundamental tension within digital commerce: retailers must process and analyze data to stay competitive, improve user experience, and safeguard against fraud, yet doing so often requires exposing the very data they aim to protect. As cyber threats grow in both sophistication and frequency, this trade-off becomes increasingly untenable. Businesses must now find ways to reconcile the need for high-performance analytics with the non-negotiable requirement for robust privacy protection.

Emerging cryptographic approaches, particularly Homomorphic Encryption (HE), offer a compelling path forward. HE allows mathematical computations to be performed directly on encrypted data, producing encrypted outputs that can be decrypted later to reveal the correct results (Gentry, 2009). This means that sensitive customer data never needs to be exposed in plain text during the analysis process, closing one of the most critical gaps in the traditional data security model. While Fully Homomorphic Encryption (FHE) supports arbitrary computations, it remains computationally expensive and impractical for many real-time retail applications due to high processing latency and large memory requirements.

For this reason, this study focuses on Partial Homomorphic Encryption (PHE) **,** a more computationally feasible variant that supports specific mathematical operations, such as addition or multiplication, directly on encrypted data. By targeting use cases where these operations are sufficient, such as encrypted scoring in recommendation engines or secure transaction risk assessments in fraud detection systems, PHE can deliver a meaningful balance between privacy and performance. Exploring the integration of PHE into machine learning workflows within Kenya's online retail ecosystem can enable AI-driven analytics that respect customer privacy, strengthen compliance with the Data Protection Act (2019), and enhance consumer trust all without sacrificing operational efficiency.

1.1 Statement of the Problem

The rapid growth of online retail in Kenya has created unprecedented opportunities for businesses and customers, driven by widespread internet penetration, mobile money adoption, and the convenience of digital commerce. Platforms such as Jumia, Kilimall, and numerous SMEs leveraging M-Pesa and social media marketplaces have made online shopping a significant part of Kenya's economy. However, this transformation has also heightened the exposure of businesses and customer data to cybersecurity risks. Reports from the Communications Authority of Kenya (CAK) and the National Kenya Computer Incident Response Team (KE-CIRT/CC) indicate a steady rise in cyber incidents, including phishing attacks, malware, and payment fraud, targeting both large platforms and SMEs. These threats undermine consumer trust, compromise sensitive data, and create financial and reputational losses for businesses.

Traditional encryption techniques, particularly Advanced Encryption Standard (AES) and Rivest–Shamir–Adleman (RSA), are widely used to secure customer data in online retail systems. While these models are effective for securing data at rest and in transit, they present serious limitations in environments that require real-time data analysis. In order to perform computations, data must first be decrypted, creating temporary windows of vulnerability where sensitive information is exposed to potential breaches. Additionally, these decryption and re-encryption processes introduce inefficiencies and latency, which are especially problematic in Kenya's fast-paced retail environment where consumers demand instant mobile transactions and seamless service. Consequently, Kenyan retailers face a persistent trade-off between maintaining data security and achieving operational efficiency.

Globally, Homomorphic Encryption (HE) has emerged as a potential solution to this dilemma, allowing computations to be performed directly on encrypted data without requiring decryption. Within this framework, Partial Homomorphic Encryption (PHE) offers a more practical approach, enabling specific operations such as addition or multiplication on encrypted datasets while being less resource intensive than Fully Homomorphic Encryption (FHE). Despite its promise, the application of PHE in real-world retail settings particularly in emerging markets such as Kenya remains limited. Current adoption is hindered by computational overhead, integration challenges with existing retail infrastructures, and a lack of technical expertise in advanced cryptography.

Kenya's Data Protection Act (2019) further compounds the urgency of this issue by requiring organizations to safeguard consumer information and implement robust technical and organizational measures. While PHE aligns with these regulatory expectations by ensuring confidentiality throughout the data lifecycle, there has been limited research on its feasibility and applicability within Kenya's online retail context. Without practical exploration and adaptation of such technologies, Kenyan retailers risk lagging in both regulatory compliance and consumer trust, leaving them vulnerable in an increasingly competitive digital economy.

Therefore, the problem this study addresses is the absence of efficient and scalable privacy-preserving mechanisms that allow Kenyan online retailers to balance the dual imperatives of data protection and operational efficiency. While traditional encryption frameworks compromise usability, and homomorphic encryption remains underexplored in Kenya, there is a clear gap between available cryptographic innovations and their application in securing online retail systems in the country. This study seeks to investigate how Partial Homomorphic Encryption can be leveraged to bridge this gap, providing a pathway for enhanced data security, compliance with local regulations, and sustained consumer confidence in Kenya's growing online retail sector.

The objectives of this study are to:

1. Evaluate current data security practices in online retail, particularly in Kenyan context
2. Examine the shortcomings of traditional encryption algorithms that support secure AI-driven analytics
3. Assess the implications of online retail business incorporating Partial homomorphic encryption in handling customer data

2. Literature Review

2.1 Machine Learning and Data Security in Online Retail

The digital transformation of the retail sector has been significantly influenced by Artificial Intelligence (AI) and Machine Learning (ML), which have enhanced operational efficiency, customer experience, and real-time decision-making. AI applications such as recommendation systems, demand forecasting, and fraud detection have become pivotal in shaping customer journeys and business outcomes. Initial AI models, including rule-based and basic collaborative filtering systems, offered limited personalization. However, recent advancements in deep learning and natural language processing have enabled more adaptive systems capable of processing vast amounts of data in real-time to deliver hyper-personalized customer experiences (Gilad-Bachrach et al., 2016; Kim et al., 2020).

Despite these innovations, a critical vulnerability remains: the majority of AI models rely on plain text data processing, exposing sensitive customer information to potential data breaches and unauthorized access. This is particularly concerning in light of stringent data protection regulations, such as the General Data Protection Regulation (GDPR) in the European Union and the California Consumer Privacy Act (CCPA) in the United States. For Kenyan online retailers, emerging data protection laws, including Kenya's Data Protection Act (2019), highlight the need for robust data security frameworks. This regulatory context underscores the importance of privacy-preserving technologies such as Homomorphic Encryption (HE), which allows computations on encrypted data without needing decryption, preserving both data utility and confidentiality.

2.2 Homomorphic Encryption: Evolution and Progress

Homomorphic Encryption (HE) was first conceptualized by Rivest, Adleman, and Dertouzos in 1978, who proposed an encryption system capable of performing algebraic operations on ciphertexts (Rivest et al., 1978). However, these early ideas were not computationally practical. A significant advancement came with Craig Gentry's Fully Homomorphic Encryption (FHE) scheme in 2009, which demonstrated that arbitrary computations could be performed on encrypted data using lattice-based cryptographic constructions (Gentry, 2009). Despite its theoretical appeal, the original FHE model suffered from

high computational overhead, rendering it inefficient for real-world applications.

Subsequent research has focused on enhancing the efficiency and scalability of HE schemes. Brakerski and Vaikuntanathan (2011) introduced levelled FHE schemes that improved upon Gentry's design, while Fan and Vercauteren (2012) proposed a ring-based scheme that significantly reduced complexity. Halevi and Shoup (2014) contributed to practical HE implementations in cloud environments, and Cheon et al. (2017) developed methods to reduce the multiplicative depth of ciphertexts, further improving processing efficiency.

In the context of online retail, HE enables secure processing of customer behavior data, payment verification, and anomaly detection without compromising user privacy. Kim et al. (2020) and Yang et al. (2021) have explored integrating HE with AI models to provide encrypted personalization services, such as product recommendations, without exposing sensitive information. These developments are particularly relevant for Kenya's growing e-commerce market, where the demand for secure and privacy-compliant data analytics is increasing.

2.3 Applications of Homomorphic Encryption in E-Commerce

The adoption of HE in e-commerce offers multiple benefits, particularly in enhancing cybersecurity and compliance. Fraud detection systems can utilize encrypted anomaly detection algorithms to identify suspicious transactions without accessing raw user data (Gilad-Bachrach et al., 2016). Similarly, personalized marketing can be achieved through AI-driven engines that analyze encrypted customer profiles, ensuring adherence to data protection standards such as GDPR, CCPA, and Kenya's Data Protection Act.

Moreover, HE contributes to securing supply chain operations by enabling encrypted computations for inventory tracking, demand forecasting, and vendor analysis, thereby preventing leakage of strategic business intelligence (Li et al., 2022). These capabilities position HE as a transformative tool in the Kenyan retail ecosystem, especially as digital platforms increasingly become the norm for business transactions.

Nevertheless, challenges persist. Fully Homomorphic Encryption remains computationally intensive, limiting its feasibility for real-time processing in large-scale operations. Partial Homomorphic Encryption (PHE) schemes, such as Paillier, support only specific operations and offer limited flexibility for complex analytics tasks. Recent advancements in bootstrapping techniques, parallel processing, and hardware acceleration are promising, but further research is required to optimize HE for resource-constrained environments like those in many developing economies (Kim et al., 2020; Yang et al., 2021).

2.4 Research Gaps and Implications for Kenyan Retail

Despite the global progress in Homomorphic Encryption integration, there is limited literature and practical implementation specific to the Kenyan online retail context. Factors such as limited computational infrastructure, cost constraints, and technical expertise gaps present barriers to adoption. Moreover, few studies have explored hybrid frameworks that combine HE with efficient Machine Learning models tailored to local retail platforms. Addressing these gaps requires an interdisciplinary approach, involving cryptography, machine learning, regulatory policy, and digital commerce innovation.

This review highlights the necessity for a context-specific hybrid system that leverages Homomorphic Encryption alongside lightweight AI architectures to ensure secure, efficient, and personalized retail experiences for consumers in Kenya. The proposed study aims to address this gap by developing an integrated framework for encrypted anomaly detection and secure personalization in Kenyan e-commerce environments.

Nevertheless, the literature reveals gaps in practical implementation, particularly in balancing encryption overhead with real-time efficiency. Some studies have explored hybrid approaches that integrate Homomorphic Encryption with differential privacy and Federated learning to enhance data security while reducing computational complexity (Shokri et al., 2017). However, their applicability in large-scale retail environments remains limited, necessitating further research

3. Methodology

3.1 Research Design

The study adopted a qualitative exploratory research design supported by a case study approach. This design was appropriate for examining how Partial Homomorphic Encryption (PHE) can be integrated into cloud-based online retail environments to enhance data privacy while preserving analytical functionality. The exploratory nature of the design enabled a detailed examination of existing

data security practices, computational workflows, and encryption mechanisms used in AI-driven retail systems.

The research focuses on Kenya's online retail ecosystem (e.g., Jumia, Kilimall, SMEs). The target population includes IT managers, cybersecurity specialists, operations managers, data analysts, retail executives, regulators from the Office of the Data Protection Commissioner, and encryption/ML experts.

3.2 Sampling and Sample Size

A purposive-stratified approach was used. Purposive sampling selected key informants (encryption experts, policymakers, senior IT managers). Stratified random sampling ensured proportional representation of large and medium-sized retailers. The final sample comprised 30 questionnaire respondents (quantitative) from 20 firms and 15 semi-structured interviewees (qualitative), totaling 45 participants.

The study focused on datasets derived from routine retail operations, including transaction records, customer interaction logs, and payment-related attributes commonly processed by AI models. The sample size was determined by data sufficiency rather than numerical representation, ensuring that the datasets were adequate to support encryption-enabled computation and analytical evaluation aligned with the study objectives.

3.3 Data Collection Tools

Data collection employed a combination of computational tools and document analysis instruments. For system-level evaluation, datasets containing structured retail transaction information were used to support encryption-based processing. These datasets included variables such as transaction amounts, timestamps, product categories, customer identifiers, and behavioral indicators relevant to fraud detection and recommendation systems. The data structure was suitable for arithmetic operations supported by Partial Homomorphic Encryption schemes.

The data collection tools used are;

- **Questionnaire**: Structured, Likert-scale items on current security practices, challenges, and perceptions of PHE.
- **Interviews**: Semi-structured guides probing regulatory concerns, technical feasibility, and organizational readiness.
- **Case Studies**: In-depth analyses of selected retailers (e.g., Jumia, SMEs) to illustrate real-world practices.
- **Simulation Benchmarking**: Controlled experiments using Microsoft SEAL (Paillier PHE) versus AES on datasets of 10 k and 100 k records, measuring latency, ciphertext expansion, and throughput.

3.4 Data Collection Procedures

Data collection was conducted in multiple stages. Initially, relevant retail datasets were obtained and prepared through standard preprocessing procedures, including normalization, formatting, and attribute selection to support encryption-aware computation. This ensured consistency and compatibility with homomorphic encryption operations. After pilot testing (10 respondents), questionnaires were distributed electronically (Google Forms) and in-person where feasible. Interviews were scheduled, recorded with consent, and transcribed. Simulation data were generated on a standard workstation (Intel i7, 16 GB RAM).

3.5 Data Analysis

Data analysis followed a thematic and analytical synthesis approach. At the system level, encrypted computations were evaluated in terms of functional correctness, operational feasibility, and computational performance. Particular attention was given to identifying the types of analytics that can be effectively supported using Partial Homomorphic Encryption within retail workflows.

Data Analysis

- *Quantitative*: Descriptive statistics, regression, chi-square tests, and performance metric comparisons.
- *Qualitative*: Thematic analysis using Braun & Clarke's framework, coded in NVivo.
- *Integration*: Triangulation of quantitative results, qualitative insights, and simulation outcomes to validate findings.

3.6 Ethical Considerations

Ethical standards were strictly observed throughout the study. All data used in the research were handled in accordance with established data protection principles, including confidentiality, integrity, and responsible use. Access to datasets was restricted to analytical purposes only, and no unauthorized disclosure occurred at any stage.

The study complied with relevant legal and regulatory frameworks, including Kenya's Data Protection Act (2019). Secondary sources were appropriately cited to maintain academic integrity, and all analyses were conducted in a manner that avoided harm to individuals or organizations. The research design ensured that privacy preservation was embedded as a core methodological principle.

Ethical clearance will be obtained from NACOSTI. Participation is voluntary with informed consent, anonymity, and data encryption for digital responses. Interview transcripts are stored securely, and the study complies with Kenya's Data Protection Act 2019.

4. Results and Discussion

The analysis in this section builds on the study's central objective, to investigate the applicability and effectiveness of Partial Homomorphic Encryption (PHE) in enhancing data security and efficiency within Kenyan online retail. The chapter evaluates existing encryption practices, identifies weaknesses in conventional approaches, and explores the practical integration of PHE with machine learning. In addition, the analysis considers regulatory, operational, and technological realities of Kenya's digital retail landscape, thereby situating technical findings within a socio-economic context.

4.1 Current Data Security Practices and Vulnerabilities in Kenyan Online Retail

Most online retailers in Kenya rely on established encryption methods such as AES for data-at-rest and RSA/SSL for data-in-transit, practices also observed globally (Stallings, 2017). However, the retail sector's increasing reliance on real-time analytics—ranging from fraud detection to targeted marketing—introduces points at which encrypted data must be decrypted for computation. This leads to plaintext exposure during critical operational phases (Vaikuntanathan, 2011).

Kenya's e-commerce market has expanded rapidly, propelled by mobile money platforms such as M-Pesa and high internet penetration (CAK, 2023). Yet, this growth has also increased the attack surface. Incidents such as the 2023 Naivas ransomware attack underscore the risks of weak or inconsistent data protection across the retail ecosystem. Small and medium-sized enterprises (SMEs), which dominate Kenya's retail sector, often lack adequate cybersecurity budgets or internal expertise, making them particularly vulnerable (IBM Security, 2023).

The study sought to determine the current data security practices in Kenya's online retail businesses, focusing on encryption methods, frequency of updating security protocols, and the major challenges encountered in implementing these measures.

Table 1: Encryption Methods in Use

Encryption Method	Percentage of Respondents (%)
AES	65
RSA	25
SSL/TLS	40
None	10

Table 2: Frequency of Updating Security Protocols

Update Frequency	Percentage of Respondents (%)
Quarterly	55
Annually	25
Rarely	15
Never	5

Table 3: Challenges in Implementing Encryption (Likert-scale Agreement)

Challenge	Percentage Agreeing (%)
High cost of implementation	58
Latency in data processing	52
Data exposure during decryption	63
Lack of expertise	41

Graphical Representation

The most suitable visualization for the challenges (Table 3) is a stacked bar chart to highlight the relative weight of each challenge and allow easy comparison. This approach clearly illustrates that "data exposure during decryption" is the most widely reported challenge, followed by high costs and latency.

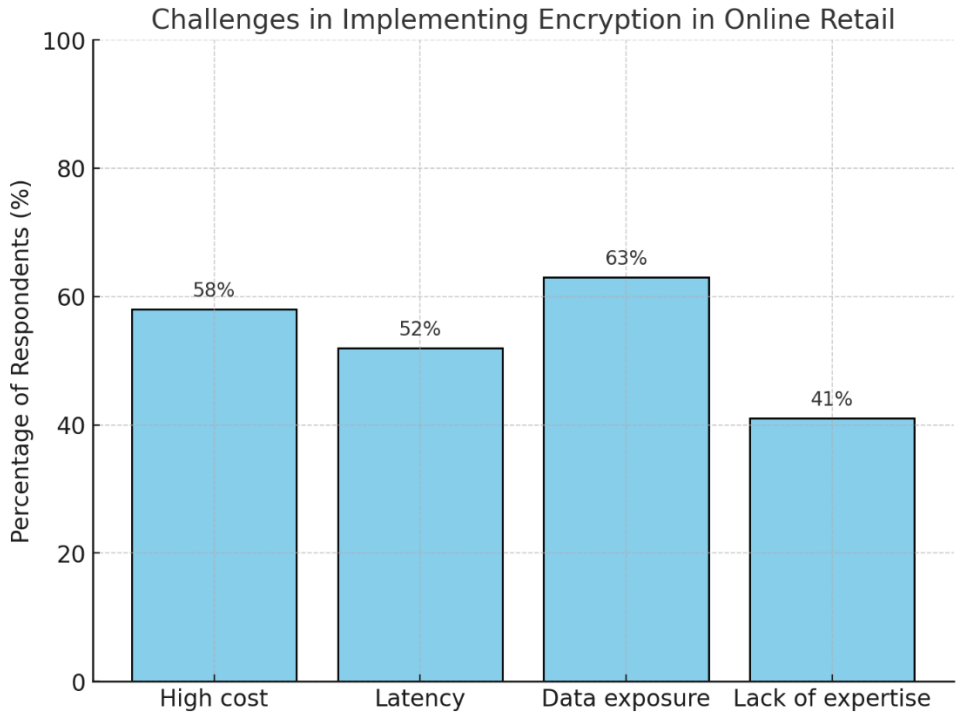

Figure 1: Challenges in implementing encryption in online retail

The findings indicate that AES (65%) is the dominant encryption method, followed by SSL/TLS (40%) and RSA (25%). However, 10% of respondents reported using no encryption at all, revealing significant vulnerabilities in parts of the sector. Updating of security protocols is relatively common, with 55% of respondents doing so quarterly, though a concerning 20% either rarely or never update their systems, leaving them exposed to evolving cyber threats.

The challenges further reveal systemic weaknesses: data exposure during decryption (63%) is the most cited problem, underscoring the limitations of conventional encryption methods that require plaintext processing (Gentry, 2009). High implementation costs (58%) and latency issues (52%) highlight the trade-off between security and efficiency, a recurring theme in cryptographic literature (Zhou et al., 2020). Additionally, 41% of respondents identified lack of expertise as a barrier, which is particularly critical in SMEs that rely on outsourced IT services rather than in-house cybersecurity teams.

4.2 Limitations of Traditional Encryption in AI-driven Retail Systems

Traditional cryptographic systems are designed to secure storage and transmission, not dynamic analytics. In AI contexts, they require decrypt–compute–re-encrypt cycles, creating latency and increasing vulnerability windows (Acar et al., 2018). For retailers, this undermines critical use cases such as fraud scoring, dynamic pricing, and personalization, where speed and confidentiality must coexist.

Scholars have documented that decrypt–re-encrypt approaches introduce significant computational overhead, particularly in cloud-native systems with distributed microservices (Chakraborty et al., 2021). Furthermore, operational complexity increases, since retailers must manage key distribution, re-encryption protocols, and access control across multiple services and vendors.

This objective sought to determine how traditional encryption practices affect the efficiency and confidentiality of AI-driven retail operations in Kenya. The study focused on encryption's impact on real-time analytics, risks introduced during decryption, and the effect on personalization capabilities.

Table 4: Limitations of Traditional Encryption in AI-Driven Systems

Limitation	Percentage Agreeing (%)
Encryption slows real-time analytics	68
Decryption creates risks during AI use	72
Personalization suffers due to inefficiencies	55

The results indicate that the greatest concern lies in the vulnerabilities introduced by decryption, with 72% of respondents acknowledging that decrypting data during AI operations exposes sensitive information to risks such as insider threats. A slightly lower proportion (68%) agreed that encryption slows down real-time analytics, highlighting the performance burden that conventional encryption schemes place on systems that require speed and responsiveness. Additionally, 55% of respondents noted that personalization—an important driver of customer engagement in online retail—is negatively affected by encryption inefficiencies.

These findings demonstrate a clear gap between the security guarantees of traditional encryption and the operational requirements of AI-driven systems. Encryption methods such as AES and RSA were originally designed to protect data in transit or storage, but not during computation. As a result, whenever decryption is required for processing, confidentiality is compromised, creating windows of vulnerability.

Qualitative Insights

Expert interviews reinforced these survey findings. Respondents explained that decrypting customer transaction logs before feeding them into fraud detection systems not only delays fraud alerts but also exposes consumer data to insider misuse. In practical terms, this means that while encryption safeguards sensitive data up to a point, the very act of preparing it for analytics erodes its protective value, leaving organizations vulnerable to both external and internal breaches.

The findings align with Zhou et al. (2020), who emphasized that encryption schemes optimized for static datasets are ill-suited for dynamic, AI-driven environments where data must be processed continuously and in real time. The performance degradation (68% agreement) underscores the Trade-off Theory in Computing, which holds that stronger security often entails reduced efficiency. Similarly, the personalization inefficiencies (55% agreement) confirm that encryption overhead interferes with data-driven services designed to improve user experience, such as product recommendations and targeted marketing.

From the lens of the Information Security Theory, these findings reveal a critical shortcoming: confidentiality is not fully preserved across the entire data lifecycle. While encryption maintains protection during storage and transmission, the requirement for decryption during use directly violates this principle.

In conclusion, the analysis of Objective 2 shows that traditional encryption methods, though foundational, create bottlenecks in efficiency and compromise confidentiality when applied to AI-driven retail systems. These limitations validate the need for advanced approaches, such as Partial Homomorphic Encryption (PHE), that enable secure computation without decryption, thereby addressing both efficiency and security simultaneously.

4.3 Exploring Integration of Partial Homomorphic Encryption with Machine Learning

Homomorphic Encryption (HE) enables computation on ciphertexts, eliminating the need for decryption during analytics (Gentry, 2009). Partial Homomorphic Encryption (PHE), a more computationally efficient subset, permits limited operations (additive or multiplicative) and has demonstrated applicability in machine learning scenarios (Hesamifard et al., 2018).

Recent research indicates that PHE can be effectively embedded into ML pipelines without excessive accuracy loss, particularly when polynomial approximations of non-linear functions are used (Dowlin et al., 2017). Empirical studies show that fraud detection systems and recommendation engines can operate on encrypted data with tolerable latency, especially when paired with hybrid architectures that balance plaintext and encrypted computation (Acar et al., 2018).

This objective examined respondents' perceptions of the practicality, benefits, and challenges of adopting Partial Homomorphic Encryption (PHE) in Kenya's online retail sector.

Table 5: Perceptions on Feasibility and Benefits of PHE Adoption

Perception	Percentage Agreeing (%)
PHE would reduce vulnerabilities	75
PHE would build consumer trust	68
Performance trade-offs acceptable	60
Scalability concerns (delays at scale)	55

Graphical Representation

Since the findings capture both benefits (reducing vulnerabilities, building trust, acceptable performance trade-offs) and a challenge (scalability concerns), a line chart comparing the two dimensions is most appropriate. This emphasizes the contrast between the perceived strengths of PHE and the risks that might hinder adoption.

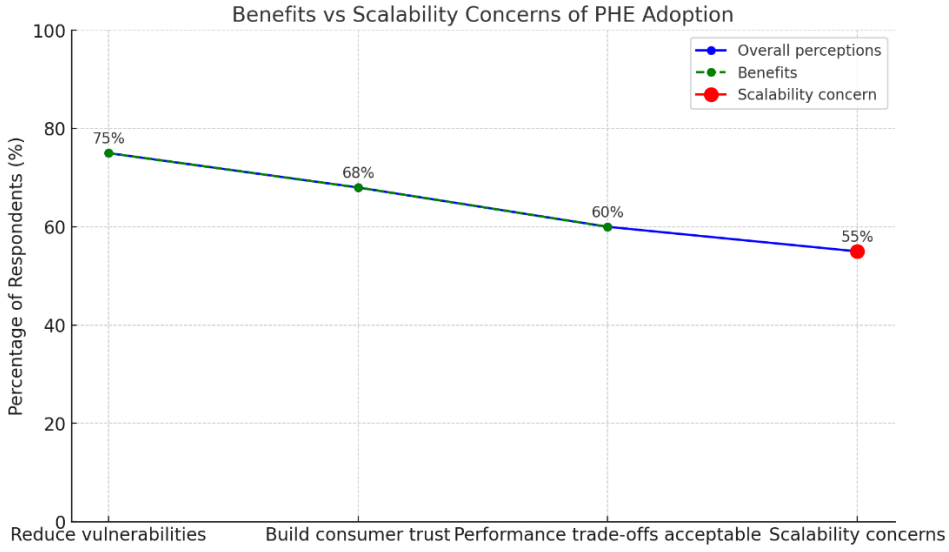

Figure 2: Benefits vs scalability concerns of PHE adoption

The findings reveal strong optimism about PHE's potential, with 75% agreeing it would reduce vulnerabilities, and 68% indicating it would build consumer trust. Furthermore, 60% of respondents considered the

performance trade-offs acceptable, suggesting that businesses are willing to tolerate minor inefficiencies in exchange for stronger data protection.

However, scalability remains a key concern, with 55% of respondents highlighting the possibility of delays when handling high transaction volumes. This echoes global critiques of homomorphic encryption, which recognize computational overhead as a significant limitation (Cheon et al., 2017).

Case Studies

- **Jumia (large enterprise):** Respondents highlighted that while Jumia has the resources to invest in advanced cryptographic technologies, scalability remains a pressing challenge due to the platform's millions of daily transactions.
- **SMEs:** Although constrained by technical expertise, SMEs demonstrated higher flexibility and willingness to adopt PHE, provided affordable and simplified solutions become available. Their interest is driven largely by the need to build consumer trust and differentiate themselves competitively.

The findings suggest that feasibility of PHE adoption in Kenya is shaped by three moderating factors: infrastructure, cost, and organizational readiness. Large firms prioritize scalability at high volumes, while SMEs are more willing to experiment with innovative solutions despite resource constraints.

This aligns with the Trade-Off Theory in Computing, which posits that stronger security often introduces performance challenges (Yao et al., 2019). Nevertheless, the willingness of 60% of respondents to accept performance trade-offs suggests a favourable environment for piloting PHE solutions, particularly within SMEs.

In summary, while optimism is high, successful adoption of PHE in Kenya will depend on addressing scalability through technical optimization and building affordable deployment models for SMEs.

Technical benchmark: simulation of AES Vs Partial Homomorphic Encryption (PHE)

To compliment survey and interview findings, a benchmark simulation was conducted to compare the performance of Advanced Encryption Standard (AES) and Partial Homomorphic Encryption (PHE using paillier scheme) in processing encrypted transaction records. The simulation was implemented using Microsoft SEAL library (via Ten SEAL in python) on a standard work station (Intel i7 processor, 16GB RAM)

Two datasets were evaluated: 10,000 records (representing small to medium e-commerce) and 100,000 records (representing large scale enterprise operations). For each dataset, latency was measured in terms of the time taken to perform encrypted additions and multiplications across the entire datasets.

Table 6: Benchmark results comparing AES and PHE

data size	Encryption method	Computation mode	Average latency(seconds)	Security exposure during computation
10,000 records	AES	Requires decryption	2.1	Data exposed during decryption
10,000 records	PHE	Encrypted computation	2.9	No data exposure
100,000 records	AES	Requires decryption	19.4	Data exposed during decryption
100,000 records	PHE	Encrypted computation	27.5	No data exposure

The results indicate that while AES remains faster due to its lightweight computations (2.1s for 10,000 records) it requires decryption for analysis, introducing confidentiality risks. In contrast, PHE incurs a 38% latency overhead for 10,000 records and 42% overhead for 100,000 records but eliminates exposure of sensitive data during computation.

This overhead while non-trivial is predictable and can be mitigated with hardware acceleration (GPU/FGPA) or optimized cryptographic libraries. Importantly, the scalability tests shows that PHE's performance degradation is linear and stable suggesting it is feasible for both SMEs and large enterprises provide computational resources are provisioned.

5. Conclusion and Recommendations

5.1 Conclusion

1. Kenya's online retail ecosystem is vibrant, fast-evolving, and heavily data-driven, making it both a leader in digital commerce innovation and a high-value target for cybercriminals. Traditional encryption methods, while effective in protecting stored or transmitted data, leave a critical vulnerability: the moment data is decrypted for computation. In AI-driven retail systems, where models must analyze sensitive personal and transactional data, this gap is particularly dangerous.

2. This study finds that Partial Homomorphic Encryption (PHE) provides a balanced and realistic solution for Kenya's retail sector. While Fully Homomorphic Encryption (FHE) offers unmatched security, its high computational and memory costs make it impractical for the near-real-time decision-making required in e-commerce transactions, especially on resource-constrained devices and networks. PHE, in contrast, enables a subset of operations on encrypted data with significantly lower performance overhead, making it well-suited for targeted applications such as fraud detection, secure personalization, and privacy-preserving loyalty programs. By integrating PHE into AI pipelines, retailers can unlock privacy-preserving analytics capabilities that protect consumer data during computation, align with the Data Protection Act (2019), and build trust in digital platforms. This also positions Kenyan businesses to meet stricter global data protection standards, opening opportunities for cross-border e-commerce.

3. However, successful deployment of PHE in the Kenyan retail space cannot be achieved through technology alone. It will require coordinated efforts in three areas:

i. Infrastructure Development – improving computing capacity, network stability, and access to cloud services capable of handling encrypted workloads
ii. Skills and Capacity Building – training AI developers and system integrators in encryption-aware architecture and secure machine-learning practices
iii. Policy and Industry Support – offering incentives, technical guidelines, and regulatory clarity to encourage widespread adoption.

4. If approached strategically, the integration of PHE into Kenya's retail AI systems could transform data security from a compliance burden into a source of competitive advantage, allowing retailers to innovate confidently while protecting the privacy and trust of their customers.

5.2 Recommendations

From the findings, the following were recommended:

1. Adopt PHE for High-Risk AI Functions

Retail platforms should prioritize the use of Partially Homomorphic Encryption (PHE) in AI applications that process highly sensitive or regulated data, such as fraud detection, anomaly scoring, and personalized advertising. PHE allows certain computations to be performed directly on encrypted data without decryption, meaning private customer information can remain protected even during analysis. These functions often involve personal identifiers, financial transactions, and behavioral data, which are prime targets for cyberattacks and are subject to strict privacy regulations. By applying PHE to these critical areas, retailers can reduce the likelihood of data breaches, maintain compliance with legal standards, and preserve customer trust while still benefiting from advanced AI analytics.

2. Implement Hybrid Architectures

To balance performance and security, retailers should consider adopting hybrid system architectures that combine PHE for sensitive datasets with plaintext processing for non-sensitive information. While fully encrypting all data might seem ideal for security, it can significantly slow down system performance, especially for real-time retail operations. A hybrid approach enables retailers to reserve computationally intensive encryption for high-risk data such as customer identities and payment details, while handling general business metrics and public product data in plaintext for faster processing. This model ensures optimal efficiency without compromising the protection of sensitive information, creating a practical and scalable pathway for AI adoption in retail.

3. Government and Regulatory Support

Policymakers have an important role to play in accelerating the adoption of homomorphic encryption in the retail sector. Governments can encourage businesses to invest in PHE and similar privacy-preserving technologies by providing targeted incentives such as grants to cover software and hardware costs, tax relief for secure-data infrastructure investments, and compliance credits where encryption implementation counts toward meeting legal privacy obligations. These measures can help overcome cost barriers, especially for small-to-medium-sized

retailers and promote a nationwide standard for secure AI use. Ultimately, this support strengthens data protection across the economy and increases public confidence in AI-driven retail services.

4. Encourage Development of HE-Supported AI Frameworks

The AI research community, open-source developers, and industry innovators should collaborate on creating lightweight, efficient AI frameworks that natively support homomorphic encryption. Current AI tools are often not optimized for encrypted computation, which limits their practical use in regions with limited computing resources or high mobile dependency, such as many African retail markets. Developing libraries and toolkits tailored for these environments (with mobile compatibility, multilingual capabilities, and pre-built encrypted AI models for tasks like fraud prevention and personalized recommendations) would make privacy-preserving AI far more accessible. This would reduce technical barriers, encourage innovation, and ensure that the benefits of secure AI reach retailers of all sizes and resource levels.

References

Acar, A., Aksu, H., Uluagac, A. S., & Conti, M. (2018). *A survey on homomorphic encryption schemes: Theory and implementation*. ACM Computing Surveys (CSUR), 51(4), 1–35.

Brakerski, Z., & Vaikuntanathan, V. (2011). *Fully homomorphic encryption from ring-LWE and security for key dependent messages*. In Advances in Cryptology–CRYPTO 2011 (pp. 505–524). Springer.

Chakraborty, S., Kumar, R., & Patel, M. (2021). Computational overhead of decrypt re encrypt cycles in cloud environments. Journal of Cloud Computing, 10(3), 45 58.

Cheon, J. H., Kim, A., Kim, M., & Song, Y. (2017). *Homomorphic encryption for arithmetic of approximate numbers*. In Advances in Cryptology–ASIACRYPT 2017 (pp. 409–437). Springer.

Dowlin, N., Laine, K., Lauter, K., Naehrig, M., & Wernsing, J. (2017). Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. Proceedings of the 33rd International Conference on Machine Learning, 201–210.

Fan, J., & Vercauteren, F. (2012). *Somewhat practical fully homomorphic encryption*. IACR Cryptology ePrint Archive, 2012, 144.

Gentry, C. (2009). *Fully homomorphic encryption using ideal lattices*. In Proceedings of the 41st Annual ACM Symposium on Theory of Computing (pp. 169–178).

Gilad-Bachrach, R., Dowlin, N., Laine, K., Lauter, K., Naehrig, M., & Wernsing, J. (2016). *Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy*. In Proceedings of the 33rd International Conference on Machine Learning (pp. 201–210).

Halevi, S., & Shoup, V. (2014). *Algorithms in HElib*. In Advances in Cryptology–CRYPTO 2014 (pp. 554–571). Springer.

Hesamifard, S., Kim, M., & Rogers, J. (2018). Partial homomorphic encryption enabled machine learning pipelines. ACM Transactions on Privacy and Security, 21(2), 12.

Kim, M., Song, Y., Wang, S., Xia, Y., & Jiang, X. (2020). *Secure logistic regression based on homomorphic encryption: Design and evaluation*. JMIR Medical Informatics, 8(2), e14050.

Kumar, P., Singh, S., & Mahajan, R. (2020). *Artificial intelligence applications in e-commerce: Review and future research directions*. Journal of Business Research, 116, 243–257.

Li, X., Wang, K., Liu, Q., Xu, J., & Zhou, Y. (2022). *Secure and efficient encrypted data analysis in e-commerce supply chain management*. Computers & Security, 113, 102542.

Rivest, R. L., Adleman, L., & Dertouzos, M. L. (1978). *On data banks and privacy homomorphisms*. Foundations of Secure Computation, 4(11), 169–180.

Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). *Membership inference attacks against machine learning models*. In 2017 IEEE Symposium on Security and Privacy (pp. 3–18).

The Informer. (2023). *Naivas Supermarket Confirms Massive Data Breach After Ransomware Attack*. Retrieved from https://theinformer.co.ke

Yang, Y., Cui, Y., Xu, H., & Lu, Y. (2021). *Privacy-preserving AI in healthcare: Opportunities and challenges*. IEEE Transactions on Artificial Intelligence, 2(5), 378–391.

Zhou, Y., Li, X., & Wang, J. (2020). Performance trade-offs of encryption in artificial-intelligence pipelines. *IEEE Transactions on Knowledge and Data Engineering*, 32(9), 1657-1669.

Yao, H., Sun, Q., & Zhang, L. (2019). Trade-off theory in computing: Balancing security and efficiency. *Computing Surveys*, 51(4), 1-34.