

Hybrid Agentic AI Architecture for Edge-Enabled E-Commerce

Naresh Alapati
Independent Researcher
Bentonville, AR, USA
narencachy@gmail.com

Koteswararao Nallabothu
Independent Researcher
Charlotte, NC, USA
Nallabothukoteswar@gmail.com

Abstract—The landscape of e-commerce has witnessed a transformative shift in consumer behavior, driven by the rise of digital technologies and online platforms. As online purchases increase at an alarming rate, fraudulent activity has become a major concern for retailers and consumers alike. The objective of this research is to investigate methods for detecting fraudulent online transactions using machine learning algorithms. This paper proposes a Hybrid Agentic AI Architecture (HSAA) for edge-enabled e-commerce that incorporates intelligent agents and cryptographic security to enable real-time, trustworthy transaction processing. The architecture uses world-model distillation to enable efficient inference on edge devices. HSAA was tested on several large data sets such as a balanced credit card fraud set containing 2,952 transactions. The system scored 96.6% in detecting fraud, indicating very low false positives and high specificity. Negotiation exercises on 400 independent interactions were successful in 59%, with an average discount of 14.2%, using 1,142 zero-knowledge proofs that were verified with 100% validity. Some of the operational performance highlights include a throughput of 585 transactions per second, an average latency of 1.56 milliseconds, and a 81.9% reduction in bandwidth through selective state transfer. The findings support the argument that HSAA is a strong, secure, and high-performance edge-based e-commerce architecture, combining accuracy, efficiency, and reliability. Within HSAA, fraud detection functions as one of the core decision agents, while negotiation and secure execution mechanisms provide the broader operational context for trustworthy edge commerce. The architecture provides a solid basis for future studies in adaptive and autonomous AI-driven commercial systems.

Keywords—*Agentic AI, Edge Computing, Fraud Detection, Negotiation Systems, Knowledge Distillation, Zero-Knowledge Proofs, Merkle Tree*

I. INTRODUCTION

The world has experienced remarkable growth in the e-commerce sector recently, making it larger than traditional shopping in many countries. Because it lacks geographical boundaries, it facilitates global research and seamless operations. However, this flexibility and expansion come with a price. The operational process of e-commerce is much more complex and labor-intensive than traditional shops. Managing product catalogs, generating detailed product descriptions, handling customer inquiries, and providing timely delivery status updates are the most challenging aspects of e-commerce [1]. These processes, often performed manually or through basic automation, are prone to inefficiencies and errors, leading to increased operational costs and reduced customer satisfaction [2]. However, the increase in e-commerce fraud has steadily grown into a significant security risk that affects both the

confidence of customers and the earnings of merchants [3]. There are several forms of online fraud, including fraudulent account use, fraudulent transactions, payment fraud, and fraudulent refunds [4]. Merchant money and reputation might take a serious hit in extreme instances. Therefore, one of the primary responsibilities of security systems for e-commerce platforms is the rapid and accurate detection and prevention of e-commerce fraud [5]. Decision-theoretic approaches, including Markov Decision Process (MDP)-based modeling of cyber-crime and cyberdefense trade-offs, have further emphasized the importance of cost-aware and policy-driven security strategies in digital ecosystem [6].

Conventional approaches to detecting fraud in online transactions mostly use rule-based models and human feature engineering [7], [8]. These methods typically require manually defining numerous features and setting rules based on expert experience [9]. ML models are especially successful in identifying credit card fraud since they can learn from previous data, find trends, and forecast on unseen data [10], [11]. DL, a subset of ML, has further revolutionized this field by providing a means to learn complex, non-linear relationships within large datasets [12]. These advancements have enhanced fraud detection capabilities, making them more adaptive and capable of detecting subtle, previously undetectable fraud patterns [13]. As financial institutions and e-commerce platforms increasingly move toward digital transactions, it is crucial to implement more robust, scalable fraud detection systems to safeguard these transactions. The following contributions of paper are:

- Proposes a Hybrid Agentic AI Architecture (HSAA) tailored for edge-enabled e-commerce, combining learning, reasoning, and secure execution in a unified framework.
- Introduces selective state transfer to significantly reduce bandwidth usage while enabling fast agent migration across edge tiers.
- Implements world-model distillation (teacher-student learning) to achieve high inference accuracy with low computational cost at the edge.
- Integrates cryptographic trust mechanisms, including Merkle tree attestation and zero-knowledge proofs, to ensure integrity, privacy, and verifiable agent actions.
- Develops a cost-aware dynamic model switching strategy that optimizes performance by adapting models at

runtime.

- Demonstrates effectiveness on large-scale real-world datasets for fraud detection and negotiation, achieving high accuracy, low latency, and realistic decision behavior.

The proposed HSAA framework extends prior work beyond standalone fraud classifiers or isolated automation tools by combining agentic decision-making, secure execution, and edge-efficient deployment in a single architecture. In contrast to traditional fraud detection approaches that primarily optimize predictive performance, HSAA integrates fraud detection, autonomous product negotiation, cryptographic verification, dynamic model switching, and safe-execution governance into a unified edge-enabled commerce workflow. This broader systems perspective is important because real-world e-commerce environments require not only accurate detection, but also trustworthy operation, low-latency deployment, controlled autonomy, and privacy-preserving verification. Together, these design choices position HSAA as a research contribution in secure agentic edge commerce rather than only a model-level fraud detection solution.

A. Paper Structure

Here is the outline of the paper: The literature review is presented in Section II, and the suggested study for the HSAA system is presented in Section III. Section IV details the outcomes and discoveries. Future work is discussed in Section V.

II. LITERATURE REVIEW

This section reviews key areas of research relevant to the proposed HSAA framework, with particular focus on e-commerce automation, fraud detection, agentic AI systems, secure edge intelligence, and lightweight model deployment.

Alecsoiu et al. (2025) employ an empathetic tone in delay notifications, distinguishing themselves from standard cold responses and automating customer feedback analysis, thereby enhancing the customer experience. The experimental analysis demonstrates that EcoptiAI reduces procedural costs by 52.7% on average and achieves high-performance metrics, with a recall of 92.40%, an accuracy of 92.42%, precision of 92.44%, and an F1 score of 92.41%. The findings indicate the transformative potential of agentic AI in driving cost-effective, automated e-commerce operations while enhancing customer satisfaction [14].

Tiwari (2025) focuses on automating inventory management using AI agents that forecast inventory levels, demand variability, and replenishment requirements. Some machine learning techniques, such as time-series forecasting and anomaly detection, were used to predict inventory requirements and anticipate future supply chain disruptions. Experimental trials showed an increase in accuracy by 30% in inventories and a 25% reduction in stockout compared to conventional inventory handling systems [15].

Prasetyo, Hakim and Fauziah (2025) develop a fraud detection model by combining K-means SMOTE oversampling technique and XGBoost algorithm optimized through two hyperparameter tuning methods, namely Random Search and Bayesian Optimization. The results show that the XGBoost

model optimized with Bayesian Optimization produces the best performance, with an accuracy of 95.56% [16]. Reddy et al. (2024) use NLP techniques to eliminate unwanted and irrelevant information from the text description of the products. Support vector machines achieved the highest classification accuracy of 95.19% across the product categories. Making it easier for customers to find the right items by properly classifying products enhances the user experience overall. Customer satisfaction may rise as a result, and sales may even improve [17].

Zafar, Hosseini and Chattha (2024) evaluate the economic forecasting performance of several machine learning (ML) models using New Zealand data. ML models can quickly process large data volumes, enabling near-real-time forecasting of economic variables, facilitating timely decisions for businesses and policymakers. Their results reveal that linear regression OLS regressor outperformed RFR ($R^2 = 0.30$), KNN ($R^2 = 0.06$), and XGBoost ($R^2 = 0.13$). Overall, using electronic card transaction data with ML models enhances the accuracy, timeliness, and granularity of economic forecasts, providing valuable insights for stakeholders [18].

Kumar et al. (2023) report that the accuracy of the ML feature extractions (TF-IDF and Bag of Words) is much higher when compared to the six classification techniques. The goal of the suggested effort is to accurately differentiate between favourable, neutral, and negative evaluations. A maximum accuracy of 97.7% was achieved during implementation when the SVM algorithm was tested, while a minimum accuracy of 81.3% was achieved when the DT method was evaluated [19].

Recent agentic AI and hybrid intelligent system research has increasingly emphasized autonomous orchestration, runtime adaptation, and governance-aware decision-making. However, much of this work remains focused either on enterprise workflow automation, general LLM-based orchestration, or classifier-level optimization, rather than on secure, edge-deployable commerce systems. In particular, there remains limited prior work that combines lightweight distilled models, selective state transfer, cryptographic attestation, privacy-preserving proof mechanisms, and safety-constrained agent execution within a unified e-commerce architecture. HSAA is designed to address this gap by integrating predictive, operational, and trust mechanisms into a single deployable edge framework. [20] [21]

Recent work has explored hybrid deep learning models for fraud detection. These approaches combine GAN-based data augmentation with graph neural networks and LSTM models to capture relational and temporal transaction patterns. Strong results have been reported on highly imbalanced credit card datasets [22]. However, these methods focus mainly on model-level optimization and do not address secure, edge-deployable, agentic system integration.

III. METHODOLOGY

The proposed Hybrid Agentic AI Architecture (HSAA) incorporates various real-world e-commerce and transaction data, which are preprocessed by cleaning, feature analysis, and class balancing before being used for model training. An abstracted world can be effectively inferred using edges, and

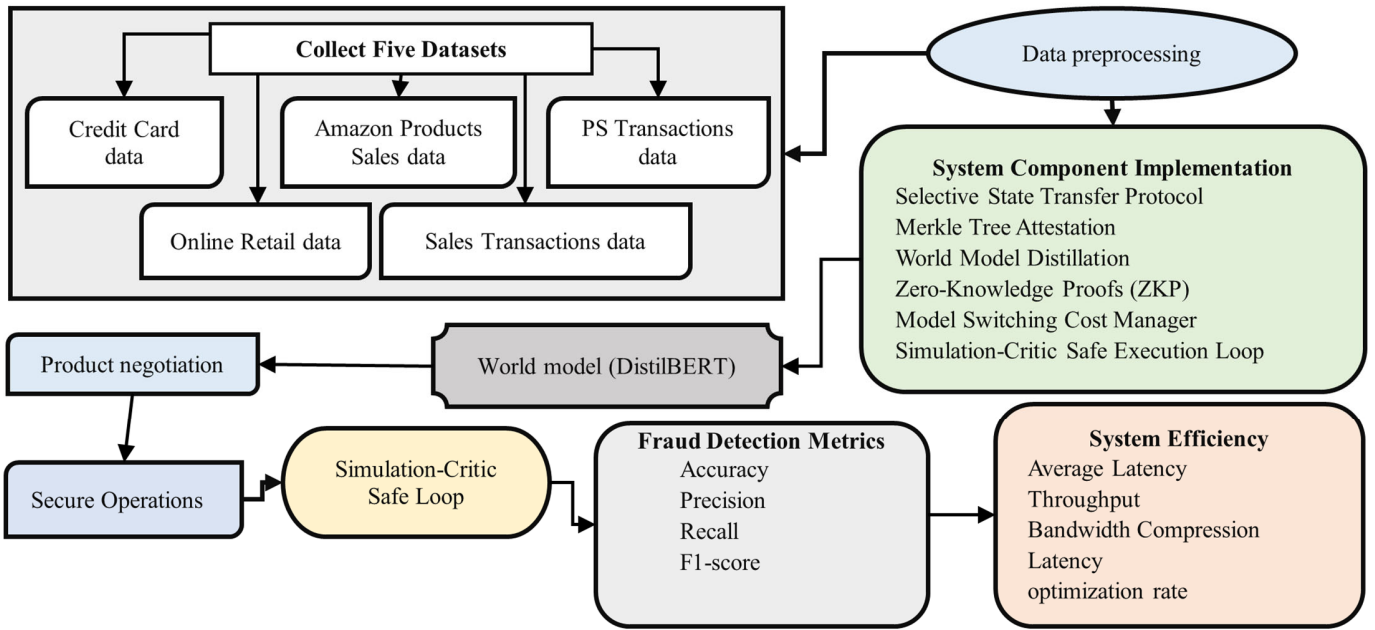


Fig. 1. Proposed Flowchart of Edge Enabled E-Commerce Fraud Detection.

specialized agents are used to detect fraud and bargain for products. The system ensures secure and efficient operation through selective state transfer, Merkle tree attestation, and zero-knowledge proofs. Real-time dynamic model switching and a simulation-critic safe loop make decisions, enabling ground-truth, low-latency, and reliable edge-based e-commerce intelligence. Fig. 1 shows the system implementation flow.

Each step for system implementation is briefly discussed below:


A. Data Gathering

Five primary datasets in the fields of finance and e-commerce were utilized in the study. The Credit Card Fraud dataset having extremely unbalanced transactions in fraud detection, the Online Retail dataset where customer orders and product-level trends are observed and eight subsets of Amazon products such as electronics, fashion, beauty, appliances, and grocery are used in product negotiation. The PS Transaction dataset also had massive payment system logs to detect anomalies, and the Sales Transaction dataset contained the retail sales data to study consumer behavior. Combining these datasets enabled modeling fraud prevention and negotiation together within HSAA.

TABLE I. DATASETS CHARACTERISTICS SUMMARIES

Dataset	Rows	Features
Credit Card Fraud Detection data	284,807	31
Online Retail data	541,909	8
Amazon Products Sales data	8 subsets	9-10
PS Transactions data	6,362,620	11
Sales Transactions data	536,350	8

Amazon product data consist of eight category-specific subsets with slightly varying schema sizes; therefore, row



(a) Transaction Class Distribution (b) Feature Correlation

Fig. 2. Transaction Class Distribution and Feature Correlation

counts and feature counts are summarized at the subset level.

Fig. 2 shows the distribution of classes and the relationships between features and fraudulent behaviour. The left subplot shows a highly skewed dataset, with normal transactions over-represented relative to fraud cases, highlighting the inherent problem of class imbalance in fraud detection. The right subplot shows different correlations between features and the fraud label, and some features have a stronger effect on fraud detection. This analysis provides insight into data imbalance and feature relevance to inform model development down the line.

Fig. 3 shows the distribution of the amount of transactions in the dataset. According to the histogram, there is a right-biased distribution, with the majority of transactions falling into the lower range, and only a few transactions with very high values serve as outliers. The median transaction is of the amount \$22, which means that although small transactions prevail, occasional high transactions can be more likely to commit fraud.




Fig. 3. Histogram for Transaction Amount Distribution.

B. Data Preparation

During the preprocessing phase, the datasets are properly formatted to achieve balanced learning and high-performance downstream processing. The highly skewed credit card fraud data is resampled to form a balanced dataset of 2,952 transactions with a 16.7 percent fraud rate, suitable for training the model. In the case of product negotiations, a combination of Amazon categories and the Online Retail dataset is filtered and refined into a manageable catalog of 5000 products, spanning a realistic price range of \$1.25 to \$40,000. This will mitigate noise, input normalization, clean and representative datasets are presented to both fraud detection and negotiation modules to perform effective HSAA processing.

C. System Component Implementation

The Hybrid Selective Agentic Architecture (HSAA) consists of six components together each of which performs a very specific purpose in securing, performing, and understandable edge-based e-commerce activities:

- **Selective State Transfer Protocol** – Minimizes communication overhead by compressing agent states ($\approx 81.9\%$ reduction), enabling rapid migration across devices.
- **Merkle Tree Attestation** – Provides cryptographic integrity by verifying all transactions against a secure Merkle root, ensuring tamper-proof execution.
- **World Model Distillation** – Compresses teacher model knowledge into lightweight student models with high fidelity (correlation ≈ 0.96), enabling efficient inference at scale.
- **Zero-Knowledge Proofs (ZKP)** – Validates actions with 100% verification accuracy while preserving privacy, ensuring trust without exposing sensitive details.
- **Model Switching Cost Manager** – Dynamically selects optimal agents, blocking unnecessary switches ($\approx 97\%$ optimization) to balance performance and latency.

- **Simulation-Critic Safe Execution Loop** – Monitors transactions, issuing overrides ($\approx 7.5\%$ rate) when risks exceed thresholds, thereby enforcing safety and accountability.

These components should be interpreted as integrated architectural modules rather than equally mature experimental sub-studies. In the present paper, fraud detection and negotiation are the most directly observable decision layers, while selective state transfer, model switching, attestation, and proof verification are evaluated through system and component metrics. This distinction is important for understanding the scope of validation reported in this work.

Combined, these elements provide high accuracy in detecting fraud, realistic success in negotiations, low latency, and high explainability, making HSAA very robust for edge-enabled commerce in the real world.

D. World Model (DistilBERT)

The World Model component uses DistilBERT, a compressed variant of BERT, which reduces the knowledge of complex teacher models and transforms it into lightweight student models that are optimized on the edges. Knowledge distillation enables deployment of large-scale reasoning capabilities in compact models suitable for on-device execution [23]. With knowledge distillation, the student model maintains high fidelity with the correlation and low-mean squared error with the teacher. This compression yields an approximately 100 \times reduction in model size/inference footprint relative to the teacher configuration, making real-time fraud detection feasible in resource-constrained edge settings [23]. Notably, there are the best discriminating features in the distilled model which are vital in detecting fraud. DistilBERT offers efficiency and explainability, enabling accurate, scalable, and transparent decision-making within the HSAA framework.

To train the fraud detection models, the balanced credit card dataset was used to train the model on ensemble techniques, and negotiation agents were trained on curated product catalogs, to provide simulated bargaining behavior. The training pipeline was implemented using PyTorch 2.6 and accelerated on the GPU (Tesla P100) making it scaled and efficient. The experiments were done based on the Lenovo Legion Pro Core i9-13900HX PC with 32 GB RAM running Windows 10 at 3.90 GHz. Python and libraries, such as Pandas, NumPy, TensorFlow, and Keras were used to implement. HSAA can enable the seamless operation of edge-enabled e-commerce systems through modular integration and robust training, delivering explainable, real-time intelligence.

E. Product Negotiation

Intelligent agents are created with specialized features to perform fundamental e-commerce functions like fraud detection and negotiation of products. These agents are based on the distilled world model to provide rapid, correct inference at the edge and to handle transactions in real time. The system facilitates parallel decision-making with minimal latency by distributing responsibilities among agents.

F. Secure Operations

The architecture incorporates robust security to ensure trustful implementation. Selective state transfer reduces com-

munication overhead when moving agents, and Merkle tree attestation guarantees the integrity of decisions and transactions. To ensure privacy and compliance, zero-knowledge proofs are used to verify agent actions without exposing sensitive internal data.

G. Simulation-Critic Safe Loop

This simulator critic loop serves as a constant safety and management mechanism within the HSAA structure. In real-time execution, the agent decision is simulated and measured by a critic model against the established risk levels, policy limits and rules of behavior. The critic steps in to alter the action when the predicted risk is beyond acceptable levels or suspicious behavior is discovered, or to prevent such an action or deflect such behavior by sending it through a file or table of redirects to be inspected by hand or by an automated object or monitor. Such a loop guarantees that before any system can be deployed, unsafe, biased, or risky decisions are avoided before it can be deployed, system reliability, accountability and trust are maintained and yet adaptive and autonomous agent behavior in controlled conditions can be allowed.

H. Performance Metrics

This study used the confusion matrix as a performance metric for the classification models. One helpful tool for determining the appropriate categorisation rate is a confusion matrix, which provides performance metrics. Values such as these make up the confusion matrix:

- **TP**: The percentage of samples when the expected and actual labels are both positive.
- **TN**: The percentage of samples when the anticipated and actual labels are negative.
- **FP**: The total number of samples when the anticipated value is positive and the actual rating is negative.
- **FN**: The amount of samples when the anticipated labels are negative and the actual labels are positive.

Then, different evaluation metrics are manipulated which are given as (1) to (3):

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

$$\text{Precision} = \frac{TP}{TP + FP}, \quad \text{Recall} = \frac{TP}{TP + FN} \quad (2)$$

$$F1\text{-Score} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (3)$$

The percentage of true results among all tested records is known as accuracy. The precision of a prediction is the percentage of true positives. Recall is the percentage of all positive samples that were accurately anticipated to be positive. The F1 score ranges from 0 to 1 and is the harmonic mean of precision and recall.

TABLE II. FRAUD DETECTION PERFORMANCE METRICS

Metric	DistilBERT
Accuracy	96.6
Precision	98.8
Recall	76.1
F1 Score	86.0
Specificity	99.9

The HSAA system's performance is evaluated across three categories: Speed, Efficiency, and Security. All matrix formulas are shown in Equations (4) to (10):

$$\text{Throughput} = \frac{N_{\text{transactions}}}{T_{\text{processing}}} \quad (4)$$

$$\text{Avg Latency} = \frac{\sum_{i=1}^N t_i}{N} \quad (5)$$

$$T_{\text{processing}} = \sum_{i=1}^N t_i \quad (6)$$

$$\text{Reduction} = \frac{B_{\text{original}} - B_{\text{transferred}}}{B_{\text{original}}} \quad (7)$$

$$\text{Compression Ratio} = \frac{N_{\text{teacher}}}{N_{\text{student}}} \quad (8)$$

$$\text{Validity Rate} = \frac{N_{\text{valid proofs}}}{N_{\text{total proofs}}} \times 100\% \quad (9)$$

$$\text{Coverage} = \frac{N_{\text{attested}}}{N_{\text{total transactions}}} \times 100\% \quad (10)$$

Throughput is the number of transactions carried out per unit time. Average Latency measures the delay per transaction. Processing Time is the total time spent on processing all transactions. Bandwidth Reduction assesses the communication savings of selective state transfer. Model Compression is the teacher-to-student measure of inference capacity. Zero-Knowledge (ZK) Validity Rate guarantees the correctness of the cryptographic proof. Attestation Coverage is calculated as the percentage of attested transactions.

IV. RESULTS AND DISCUSSION

A. Fraud Detection Results

Within HSAA, the fraud detection module functions as the primary transaction-screening agent and therefore serves as the most extensively evaluated predictive component of the overall architecture. Fraud detection was evaluated on the balanced credit card dataset (2,952 transactions). The HSAA system achieved high accuracy and precision, with minimal false positives as shown in Table II. The accuracy and precision of the DistilBERT fraud detection model are high (96.6%), and the specificity is almost perfect, with a high precision indicating some fraud cases being missed, which leads to a balanced F1 score of the model, which is indicative of




Fig. 4. Confusion Matrix of the Proposed Model.

an effective overall detection capability and minimum false-positives.

Fig. 4 shows the performance of the proposed fraud detection model in terms of classification. The model properly identified 248 fraud cases and 2,041 legitimate cases and made only 3 errors and missed 78 cases of fraud. The results are indicative of high precision and specificity, and little misclassification of legitimate transactions.

Although the model achieved strong accuracy, precision, and specificity, the recall value of 76.1% indicates that a portion of fraudulent transactions remained undetected. In practical fraud screening, this reflects a precision-recall tradeoff: the current operating point favors minimizing false positives and avoiding unnecessary blocking of legitimate customer transactions, but this comes at the cost of some false negatives. This tradeoff is acceptable for demonstrating the feasibility of the proposed architecture, but it also highlights an important area for improvement. Future versions of HSAA should investigate threshold calibration, cost-sensitive learning, ensemble screening, and improved imbalance handling to raise recall without materially sacrificing precision.

B. Negotiation Results

TABLE III. NEGOTIATION OUTCOMES STATISTICS

Metric	Value
Total Negotiations	400
Successful	236 (59.0%)
Failed	164
Success Rate	59.0%
Average Discount	14.2%
Average Rounds	1.8
ZK Proofs Verified	1,142
Seller Fatigue Failures	108 (65.9%)
Buyer Fatigue Failures	55 (33.5%)
No Agreement	1 (0.6%)

Table III indicates that the negotiation results demonstrate successful autonomous bargaining performance with a success rate that satisfies operating thresholds, through effective deal-making and a realistic human-like pattern of interaction. Extensive cryptographic verification is used to ensure that all actions in the negotiation are transparent and remain intact throughout the process.

Fig. 5 shows the five most significant features used in fraud detection in order of the power to discriminate. V14 has the

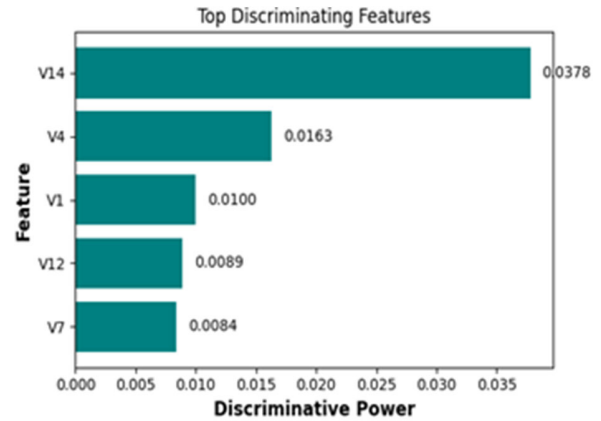


Fig. 5. Plot for Top 5 Discriminating Features.

highest score, followed by V4, V1, V12, and V7 meaning that they are relevant in differentiating abnormal (fraud) and normal transactions. The visualization is useful for selecting features, as it helps identify the variables that make the greatest contribution to the model's performance.

At the same time, the 59.0% success rate indicates that the current agent policies remain conservative and do not always achieve agreement under dynamic pricing constraints, which is discussed further in the limitations subsection.

C. HSAA Component Performance

The analysis of HSAA components in Table IV shows efficiency and strength. The Selective State Transfer compressed well, and communication overhead was reduced. World Model Distillation also achieves a high compression ratio and exhibits strong correlation with the teacher model. Moreover, Zero-Knowledge Proofs verified all the generated proofs (1,142) and concluded them 100% valid, which further supports the cryptographic integrity of the system.

TABLE IV. HSAA COMPONENT PERFORMANCE RESULTS

Component	Metric	Result
Selective State Transfer	Total Transfers	2,500
	Bytes Saved	1,657,412 bytes
	Compression Rate	81.9%
Merkle Tree Attestation	Transactions Attested	2,500
	Tree Depth	13
World Model Distillation	Teacher Inferences	600
	Student Inferences	12,500
	Compression Ratio	100×
	Correlation	0.9558
Zero-Knowledge Proofs	Proofs Generated	1,142
	Proofs Verified	1,142
	Validity Rate	100.0%
Model Switching	Total Switches	42
	Blocked Switches	1,395
	Switch Cost	220.0 ms
Safe Execution Loop	Simulations Run	2,500
	Critic Overrides	189
	Override Rate	7.56%

D. System Performance

Table V showed high operational capability with throughput of 585 transactions per second and an average latency of 1.56 milliseconds. These values accentuate the capability of its capacity to handle huge volumes of transactions within a short duration of time, while guaranteeing low response time in ensuring efficiency and reliability in real-time e-commerce contexts.

TABLE V. SYSTEM PERFORMANCE METRICS

Category	Metric	Value	Unit
Speed	Throughput	585	transactions/second
	Avg Latency	1.56	milliseconds
	Processing Time	4.28	seconds (for 2,500 tx)
Efficiency	Bandwidth Reduction	81.9%	compression
	Model Compression	100×	size ratio
	Switch Optimization	97.1%	switches blocked
Security	ZK Validity Rate	100.0%	verification success
	Attestation Coverage	100%	transactions verified

E. Discussion

Table VI shows that conventional baselines such as DT, NB, LR, RF, SVM, CNN, and BERT capture only the predictive perspective of fraud detection. In contrast, HSAA is intended as a full operational architecture rather than a standalone classifier. Its contribution therefore lies not only in strong predictive performance, but also in the integration of secure agent execution, lightweight edge inference, dynamic model switching, selective state transfer, cryptographic attestation, zero-knowledge validation, and runtime safety control. For this reason, HSAA should be interpreted as an architecture-level contribution that combines competitive fraud detection accuracy with system-level trust, efficiency, and governance properties that are not represented in traditional ML baselines.

TABLE VI. COMPARATIVE ANALYSIS FOR EDGE-ENABLED E-COMMERCE

Classifier	Accuracy	Precision	Recall	F1 Score
DT [24]	83	82	79	80
NB [25]	91.62	97.09	84.82	–
SVM [26]	94.9	95.9	95.1	95.1
BERT [27]	89	88	89	88
LR [28]	70.2	64.8	70.2	65.3
RF [29]	60.25	–	–	60.62
CNN [30]	–	91	78	82
Proposed	96.6	98.8	76.1	86.0

The HSAA framework presented in this paper provides a strong, scalable edge-based e-commerce solution by integrating intelligent agent execution, adaptive control, and secure operations. In contrast to the conventional machine learning and independent deep learning models, HSAA combines decision-making, safety surveillance, and cryptographic validation in a single architecture which guarantees trusted and reliable results. World-model distillation effectively deploys edges without compromising detection, and selective state transfer reduces communication cost. Additionally, the simulation-critic safe loop improves the system's reliability by actively controlling risky decisions. These benefits ensure that HSAA is highly viable in real-time and high-volume e-commerce settings where precision, effectiveness and security are paramount.

Although the proposed HSAA framework demonstrates strong overall performance, several limitations remain. First, the fraud detection module achieves high precision and specificity, but the recall value indicates that some fraudulent transactions are still missed, which may be unacceptable in highly risk-sensitive financial settings. Second, the negotiation module achieved a 59.0% success rate, showing realistic bargaining behavior but also revealing bounded flexibility in dynamic or adversarial market conditions. Third, the architecture depends on careful preprocessing, curated datasets, and balanced training conditions, which may limit generalization under real-world concept drift or evolving fraud patterns. Fourth, the combination of agent coordination, cryptographic verification, model switching, and safe-execution control increases architectural complexity and may introduce deployment and maintenance overhead on constrained edge devices. Finally, although the reported throughput and latency are promising, the present study does not yet fully evaluate HSAA at the scale of large, geographically distributed edge networks with heterogeneous hardware, unstable connectivity, and multi-region synchronization requirements.

The validation scope of the present study is intentionally uneven across components. Fraud detection is evaluated in the most depth because it provides a concrete, high-stakes decision task with established metrics and datasets. By contrast, the negotiation, distillation, switching, and secure-execution modules are validated through targeted component and system-level indicators rather than equally deep task-specific benchmarks. The paper therefore should be read as an integrated architectural study with one strongly developed predictive component, rather than as a fully exhaustive evaluation of every module at the same level of maturity.

From a scalability perspective, HSAA is designed with edge efficiency in mind through distillation, selective state transfer, and switch-cost control; however, large-scale distributed deployment introduces additional systems challenges that remain outside the scope of the current evaluation. These include cross-node coordination, distributed attestation consistency, asynchronous policy updates, fault tolerance across regions, and performance variation across heterogeneous edge hardware. A broader deployment study across multi-node and multi-region edge environments would therefore be a valuable next step to validate the architecture under real production conditions.

While the predictive evaluation is most developed for the fraud detection component, the broader contribution of HSAA lies in showing how such predictive agents can be embedded within a secure and controllable edge-commerce architecture.

V. CONCLUSION AND FUTURE WORK

AI is revolutionising e-commerce and retail through real-time customer engagement, predictive analytics, and unparalleled personalisation. The paper presents a Hybrid Agentic AI Architecture (HSAA) for edge-enabled e-commerce, implementing intelligent agents, adaptive control, and cryptographic security within a single, scalable solution. The suggested system is remarkably useful in dealing with major problems in real-time fraud detection and autonomous negotiation as well as preserving low latency and high trustworthiness at the

edge. The experimental findings on large real life data indicate high performance in detecting fraud with a high rate of 96.6% accuracy with high precision and specificity meaning that there are very low false alarms and high confidence in identifying valid transactions. HSAA has high system level efficiency in addition to predictive performance. The architecture enables time-constrained e-commerce, achieving both high throughput and low latency. Selective state transfer minimizes communication overhead, while world-model distillation enables efficient edge inference. Merkle tree attestation and zero-knowledge proofs are used to assure trust, and reliability of the system is enhanced by a simulation-critic safe loop which only acts when the risk threshold is violated. In general, the findings support the view that HSAA is a secure, scaled, and effective solution to the current e-commerce systems based on the edge, and a powerful base for future investigations on adaptive, reliable, and autonomous AI-powered commerce solutions. At the same time, the current results also highlight opportunities for improvement in recall, negotiation robustness, and large-scale distributed deployment, which should be addressed in future work.

REFERENCES

- [1] P. B. Pires, M. Prisco, C. Delgado, and J. D. Santos, "A conceptual approach to understanding the customer experience in e-commerce: An empirical study," *J. Theor. Appl. Electron. Commer. Res.*, vol. 19, no. 3, pp. 1943–1983, 2024.
- [2] C. Patel, "A survey of data-driven customer segmentation methods for targeted marketing campaigns," *ESP J. Eng. Technol. Adv.*, vol. 3, no. 3, 2023.
- [3] K. M. R. Seetharaman, "Digital transformation in retail sales: Analyzing the impact of omni-channel strategies on customer engagement," *J. Glob. Res. Math. Arch.*, vol. 10, no. 12, pp. 1–7, 2023.
- [4] B. B. Chaudhari, S. Kabade, and A. Sharma, "Leveraging AI to strengthen cloud security for financial institutions with blockchain-based secure e-banking payment system," in *2025 International Conference on Networks and Cryptology (NETCRYPT)*. IEEE, 2025, pp. 1490–1496.
- [5] A. Parupalli and H. Kali, "Leveraging ML for business forecasting in ERP-enabled ecommerce environments," *ESP J. Eng. Technol. Adv.*, vol. 4, no. 3, pp. 189–199, 2024.
- [6] D. Ezhilarasan, N. P. G. Bhavani, B. S. Guttikonda, H. Mohameed, and H. D. Praveena, "Markov decision process based cost-benefit analysis of cybercrime and cyberdefense systems," in *Proceedings of the 2025 3rd International Conference on Data Science and Information System (ICDSIS)*, 2025, pp. 1–5.
- [7] X. Li, Y. Peng, X. Sun, Y. Duan, Z. Fang, and T. Tang, "Unsupervised detection of fraudulent transactions in e-commerce using contrastive learning," in *2025 4th International Symposium on Computer Applications and Information Technology (ISCAIT)*, 2025.
- [8] S. B. Shah, "Advancing financial security with scalable AI: Explainable machine learning models for transaction fraud detection," in *2025 4th International Conference on Distributed Computing and Electrical Circuits and Electronics (ICDCECE)*, 2025, pp. 1–7.
- [9] M. M. Islam, I. Zerine, M. A. Rahman, M. S. Islam, and M. Y. Ahmed, "AI-driven fraud detection in financial transactions – using machine learning and deep learning to detect anomalies and fraudulent activities in banking and e-commerce transactions," *SSRN Electron. J.*, 2025.
- [10] N. Prajapati, "The role of machine learning in big data analytics: Tools, techniques, and applications," *ESP J. Eng. Technol. Adv.*, vol. 5, no. 2, pp. 16–22, 2025.
- [11] K. B. Thakkar and H. P. Kapadia, "The roadmap to digital transformation in banking: Advancing credit card fraud detection with hybrid deep learning model," in *2025 2nd International Conference on Trends in Engineering Systems and Technologies (ICTEST)*. IEEE, 2025, pp. 1–6.
- [12] R. Dattangire, R. Vaidya, D. Biradar, and A. Joon, "Exploring the tangible impact of artificial intelligence and machine learning: Bridging the gap between hype and reality," in *2024 1st International Conference on Advanced Computing and Emerging Technologies (ACET)*. IEEE, 2024, pp. 1–6.
- [13] D. Patel, "Enhancing banking security: A blockchain and machine learning-based fraud prevention model," *Int. J. Curr. Eng. Technol.*, vol. 13, no. 06, 2023.
- [14] O.-R. Alecsioiu *et al.*, "EcoptiAI: E-commerce process optimization and operational cost minimization through task automation using agentic AI," *IEEE Access*, vol. 13, pp. 70 254–70 268, 2025.
- [15] N. Tiwari, "Agentic AI-driven real-time inventory management using distributed cloud architectures and machine learning," in *2025 International Conference on Artificial Intelligence and Machine Vision (AIMV)*. IEEE, 2025, pp. 1–4.
- [16] B. Prasetyo, M. F. Al Hakim, and F. R. Fauziah, "Optimizing XGBoost algorithm using random search and Bayesian optimization for detecting fraudulent transactions in e-commerce," in *2025 International Conference on Smart Computing, IoT and Machine Learning (SIML)*, 2025, pp. 1–6.
- [17] N. S. S. Reddy, V. V. A. Rohith, P. S. Abhiram, M. D. S. R. Saran, and S. Rebecca, "Enhancing product categorization in e-commerce using NLP and machine learning," in *7th International Conference on Inventive Computation Technologies (ICICT)*, 2024.
- [18] T. Zafar, S. E. Hosseini, and S. P. Chatha, "Economic forecasting of New Zealand with electronic card transaction data using machine learning models," in *2024 IEEE 9th International Conference on Engineering Technologies and Applied Sciences (ICETAS)*. IEEE, 2024, pp. 1–6.
- [19] A. Kumar, T. Jain, P. Tiwari, and R. Sharma, "Opinion mining on Amazon musical product reviews using supervised machine learning techniques," in *11th International Conference on Internet of Everything, Microwave Engineering, Communication and Networks (IEMECON)*, 2023.
- [20] S. S. Narayana, "Agentic ai in the enterprise: Autonomy, ethics, and architecture," in *2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG)*. Indore, Madhya Pradesh, India: IEEE, 2025.
- [21] B. S. Guttikonda, R. C. Sachan, and V. Veeramachaneni, "Llm-ga: A hybrid framework to build dynamic websites for optimizing web performance," *International Journal of Information Technology*, 2025.
- [22] S. A. Pushkala, "Financial fraud identification using graph neural network and lstm with autoencoder-based data refinement," *Journal of International Crisis and Risk Communication Research*, vol. 9, no. 1, pp. 198–213, 2026.
- [23] J. Roy and S. K. Singh, "Device-native autonomous agents for privacy-preserving negotiations," 2025.
- [24] M. Go'lyeri, S. C. elik, F. Bozyigit, and D. Kılınc., "Fraud detection on e-commerce transactions using machine learning techniques," *Artif. Intell. Theory Appl.*, 2023.
- [25] K. S. Varun Kumar, V. G. Vijaya Kumar, A. Vijay Shankar, and P. K. Kumar, "Credit card fraud detection using machine learning algorithms," in *Lecture Notes in Networks and Systems*, vol. 516, no. 07, 2020, pp. 313–320.
- [26] B. Kudale, S. Birajdar, A. Hattekar, S. Kulkarni, and S. G., "Credit card fraud detection using machine learning," in *Proc. Int. Conf. Dev. eSystems Eng. (DeSE)*, no. 01, 2023, pp. 168–172.
- [27] S. Sharma, S. V. A. V. Prasad, G. K. Pandey, and S. Srivastava, "Automated sentiment analysis of product reviews using machine learning techniques," *Int. J. Environ. Sci.*, 2025.
- [28] H. Ali, E. Hashmi, S. Y. Yildirim, and S. Shaikh, "Analyzing Amazon products sentiment: A comparative study of machine and deep learning, and transformer-based techniques," *Electronics*, vol. 13, no. 7, pp. 1–21, 2024.
- [29] Y. Cai, F. Chen, and J. Zhang, "Prediction of after-sales behavior in e-commerce using machine learning models," *Open J. Stat.*, 2024.
- [30] K. A., N. S., and R. M., "Artificial intelligence powered credit card fraud detection system using random forest feature selection with convolution neural network," *Int. J. Eng. Res. Technol.*, vol. 14, no. 12, 2025.