

Blockchain-Enhanced Chain of Custody with Deep Learning Tampering Detection for Digital Forensics

Hung Truong Xuan^{1,2,*}, Dung Luong The², Tu Tran Anh²

¹Cyber Security and High-Tech Crime Prevention Department, Ministry of Public Security, Ha Noi, Vietnam

²Academy of Cryptography Techniques, Ha Noi, Vietnam

*Corresponding author email: hungtx.ncs@actvn.edu.vn

Abstract

The preservation of digital evidence integrity currently faces critical challenges, notably manual chain-of-custody errors ranging from 15%-20%, the proliferation of sophisticated tampering techniques, and inherent scalability limitations. To address these issues, this paper presents an integrated framework that synergistically combines blockchain-based immutable custody chains, dual-branch convolutional neural networks for tampering detection, and hybrid consensus mechanisms. Through systematic ablation studies, we demonstrate that the proposed smart contract automation effectively mitigates manual custody errors, ensuring a tamper-evident and immutable custody record. Furthermore, our dual-branch architecture, enhanced with adaptive fusion ($\alpha = 0.6$), attains a 98.5% tampering detection accuracy—representing a 5.3% improvement over single-branch baselines. Additionally, the hybrid Proof-of-Authority and Byzantine Fault Tolerant consensus mechanism delivers a throughput of 10,000 transactions per second, marking a 1,429-fold improvement over traditional blockchain implementations. A comprehensive evaluation on the NIST CFReDS, supported by statistical validation ($p < 0.001$), demonstrates the superiority of our approach over six baseline methods. We further provide a detailed failure analysis, a computational cost breakdown, and validation through simulated forensic scenarios, alongside proposed integration pathways for commercial forensic tools such as EnCase and FTK to facilitate practical adoption.

Keywords: Blockchain, chain of custody, digital forensics, smart contracts, tampering detection.

1. Introduction

Digital evidence constitutes a critical component in over 80% of contemporary criminal investigations [1]. However, preserving the integrity of such evidence remains a formidable challenge due to three fundamental vulnerabilities. First, manual chain-of-custody documentation is prone to human error, exhibiting rates as high as 15%-20% in multi-handler scenarios [2]. Second, sophisticated digital tampering techniques have evolved to defeat traditional hash-based verification mechanisms, necessitating advanced detection models [3]. Finally, centralized storage architectures create single points of failure, leaving repositories susceptible to unauthorized modifications and coordinated attacks.

Traditional forensic methodologies predominantly rely on cryptographic hashing (e.g., MD5, SHA-256) for integrity verification, coupled with paper-based custody logs. While these approaches may suffice for rudimentary scenarios, they are increasingly inadequate against advanced threats. Notably, the Message Digest Algorithm 5 (MD5) has demonstrated collision vulnerabilities since 2004, thereby enabling potential evidence forgery [1]. Furthermore, manual documentation introduces errors in approximately 15%-20% multi-jurisdictional

cases [2], while centralized evidence repositories create single points of failure susceptible to insider threats.

Recent studies have explored blockchain technology for immutable evidence logging [4-6] and machine learning for automated tampering detection [3, 7, 8].

In the realm of blockchain-based solutions, M. AlKhanafseh *et al.* [4] achieved 97% integrity through a Blockchain and Long Short-Term Memory (LSTM) integration; however, this method incurred a computational overhead exceeding 200% and necessitated manual verification. Similarly, while Imrankhan *et al.* [6] implemented a Byzantine Fault Tolerant (BFT) consensus mechanism attaining 99.1% reliability, the system was limited to a throughput of only 50 TPS, which is insufficient for enterprise-scale deployment. Furthermore, Lone *et al.* [5] designed frameworks specific to IIoT, which lack the generalization required for heterogeneous evidence types.

Regarding machine learning approaches, Tyagi [3] introduced dual-branch CNNs that achieved 93.2% accuracy in image forgery detection. However, conventional models often rely on fixed fusion weights, limiting their adaptability to complex forgery scenarios. To address this, advanced frameworks such as HiFi-Net

by Guo *et al.* [7] have been proposed, achieving accuracy ranges of 89%-94% on benchmark datasets. Specifically for video forensics, Feng *et al.* [8] leveraged a hybrid Multi-Scale Dilated Convolutional Neural Network and Long Short-Term Memory (MSDCNN-LSTM) architecture to detect frame deletion with a precision of 96.86%.

A significant deficiency in these approaches is the lack of integration with custody chain mechanisms and the absence of failure analysis for real-world deployment. Consequently, a distinct research gap remains, as no existing work provides an integrated blockchain-ML framework that offers a systematic evaluation, ablation studies quantifying component contributions, failure analysis identifying system limitations, scalability validation beyond small-scale prototypes, or practical pathways for tool integration.

This paper addresses the aforementioned challenges by presenting an integrated framework with four primary contributions.

First, we design intelligent Smart Contracts for Chain of Custody to enable autonomous custody validation. Unlike manual systems that suffer from high error rates, our systematic ablation study demonstrates that this automated validation effectively eliminates the 15-20% error rate observed in traditional documentation [2], thereby achieving 100% integrity preservation versus 95% in manual systems.

Second, to enhance tampering detection, we propose an Adaptive Dual-Branch CNN. We extend Tyagi's architecture [3] by incorporating adaptive fusion weights optimized through systematic grid search. Comprehensive ablation studies evaluate the impact of fusion weights ($\alpha \in \{0.4, 0.5, 0.6, 0.7, 0.8\}$), compare single versus dual-branch architectures, and assess backbone networks (ResNet50/101, EfficientNet) alongside robustness to JPEG compression. Our optimized configuration achieves 98.5% accuracy, a statistically significant 5.3% improvement over single-branch baselines ($p < 0.001$, paired t -test).

Third, we address the scalability bottleneck by architecting a Hybrid Consensus Mechanism that combines Proof-of-Authority (POA) with Byzantine Fault Tolerance (BFT). This novel consensus model delivers a throughput of 10,000 TPS, which is 1,429 times faster than Bitcoin [9] and 200 times faster than Imrankhan's BFT implementation [6]. Scalability experiments further validate the system's performance across 10 to 100 validator nodes.

Finally, we provide a comprehensive evaluation of the proposed framework. This includes a comparison with six baseline methods on the National Institute of Standards and Technology Computer Forensic Reference Data Sets (NIST CFReDS) [10], statistical validation with confidence intervals and significance

tests, and a detailed failure analysis identifying adversarial scenarios. We also present a computational cost breakdown, a real-world case study with quantified metrics, and integration pathways for commercial forensic tools such as EnCase Forensic and Forensic Toolkit (FTK) to facilitate practical adoption.

2. Related Work

2.1. Traditional Digital Forensics

Foundational principles regarding evidence acquisition, preservation, analysis, and presentation were established by Casey [2]. Conventionally, integrity verification within this domain relies upon cryptographic hashing algorithms, such as Message Digest Algorithm 5 (MD5), Secure Hash Algorithm SHA-1, and Secure Hash Algorithm SHA-256, coupled with manual chain-of-custody logs. However, the National Institute of Standards and Technology (NIST) has identified critical vulnerabilities in these methods [1], specifically noting the susceptibility of MD5 to collision attacks, high documentation error rates of 15-20%, and the inherent risks of centralized storage. Consequently, our proposed smart contract automation is designed to systematically address these deficiencies.

2.2. Blockchain-Based Evidence Management

Recent scholarship has sought to mitigate these risks through distributed ledger technologies. AlKhanafseh and Surakhi [4] pioneered a blockchain-LSTM integration that achieved 97% integrity; however, this approach incurred a substantial computational overhead exceeding 200% and necessitated manual verification. In a specific application domain, Lone *et al.* [5] architected frameworks tailored for the Industrial Internet of Things (IIoT). While achieving 98% integrity, their solution lacks generalization capabilities beyond industrial environments. Addressing this gap, Cuevas *et al.* [11] recently conducted a systematic review of blockchain-based Chain of Custody procedures, confirming their superiority over manual logging. Furthermore, Hung *et al.* [12] proposed a dedicated design for compliant evidence handling, validating the architectural shift towards decentralized verification.

Recent literature from 2025 has further validated the convergence of these technologies. For instance, Igonor *et al.* [13] provided a systematic review highlighting the shift towards automated, smart-contract-based forensic workflows. In parallel, Nargis *et al.* [14] demonstrated a prototype for evidence protection using decentralized ledgers, while El-Kady [15] explored the specific integration of AI models for advanced forensic analysis in blockchain environments.

2.3. Machine Learning Tampering Detection

In the context of automated detection, state-of-the-art methods like the hierarchical approach (HiFi-Net) [7] demonstrate robust performance. Notably, Tyagi [3] introduced dual-branch CNNs that process both spatial and frequency domains, achieving 93.2% accuracy

on CASIA v2.0. Despite these advances, the rigid reliance on a fixed fusion weight ($\alpha = 0.5$) limits the model's adaptability to diverse tampering techniques. Furthermore, Feng *et al.* [8] proposed a hybrid MSDCNN-LSTM framework specifically to capture long-term variation patterns and spatial anomalies in video forensics. In a parallel development, Darwish *et al.* [16] introduced an enhanced approach based on perceptual hash functions to secure video copyright and integrity on blockchain platforms. Collectively, these existing works exhibit notable deficiencies, particularly regarding the absence of systematic ablation studies, failure case analysis, robustness evaluation against compression artifacts, and the lack of integration with custody chain mechanisms.

2.4. Comparative Analysis

Table 1 synthesizes the critical limitations observed in prior literature and highlights how our contributions address these gaps across integrity, accuracy, throughput (TPS), and systematic ablation.

Table 1. Comparison with related work

Work	Integrity	Accuracy	TPS	Ablation
Casey [2]	95%	N/A	N/A	No
AlKhanafseh <i>et al.</i> [4]	97%	N/A	7	No
Tyagi [3]	N/A	93.2%	N/A	Limited
Lone <i>et al.</i> [5]	98%	N/A	50	No
Imrankhan <i>et al.</i> [6]	99.1%	N/A	50	No
Feng <i>et al.</i> [8]	N/A	96.9%	N/A	N/A
HiFi-Net [7]	N/A	94.0%	N/A	Yes
Ours	100%	98.5%	10,000	Yes

3. Proposed Framework

3.1. Overview of the Integrated Framework

Our integrated framework comprises three primary components: (1) blockchain-based chain of custody with smart contract automation, (2) dual-branch CNN for tampering detection with adaptive fusion, and (3) hybrid consensus mechanism for scalability. As illustrated in Fig. 1, the framework processes evidence through parallel pipelines: blockchain for immutable custody logging and deep learning for content tampering detection.

3.2. Blockchain Chain of Custody

Traditional custody documentation is frequently compromised by manual error rates ranging from 15-20%, as noted by Casey [2]. To address this, we implement a permissioned blockchain architecture leveraged by smart contracts to facilitate autonomous validation. The structural integration of the smart contract within the blockchain is depicted in Fig. 2.

The operational logic for securing the chain of custody is formalized in Algorithm 1. The smart contract enforces four critical rules: custodian authentication via digital certificates, temporal consistency to ensure monotonic timestamps, geographic plausibility for location-based validation, and role-based access control governed by a permissions matrix.




Fig. 1. Integrated system architecture




Fig. 2. Blockchain chain of custody structure

Algorithm 1 Smart contract chain of custody

```

1: Initialize blockchain with genesis block  $B_0$ 
2: Define smart contract rules  $\mathcal{R} = \{r_1, \dots, r_n\}$ 
3: for each custody transfer event  $e_i$  do
4:   Compute hash:  $h_i = \text{SHA-256}(\text{evidence}_i)$ 
5:   Extract metadata:  $m_i = \{\text{timestamp, custodian, location}\}$ 
6:   Generate signature:  $\sigma_i = \text{Sign}(\text{privkey}_i, h_i || m_i)$ 
7:   Create block:  $B_i = \{h_i, m_i, \sigma_i, h_{i-1}\}$ 
8:   if  $\text{ValidateContract}(B_i, \mathcal{R}) = \text{TRUE}$  then
9:     Append via hybrid consensus (Section 3.4)
10:    Broadcast to network
11:   else
12:     Reject and alert administrator
13:   end if
14: end for
15: return Immutable custody chain
  
```

To validate the efficacy of this approach, we conducted a systematic ablation study by incrementally removing components. As presented in Table 2, the results demonstrate that smart contract automation eliminates manual errors entirely, achieving a 100% integrity rate. This improvement is statistically significant, as confirmed by McNemar’s test ($p < 0.001$).

Table 2. Blockchain component ablation study

Configuration	Integrity	Error Rate
Baseline (Manual)	95.0%	15.2%
+ Blockchain Only	98.2%	8.3%
+ Smart Contracts	99.8%	1.1%
+ All Rules	100.0%	0.0%

3.3. Dual-Branch CNN Tampering Detection

We extend the architecture proposed by Tyagi [3] by incorporating an adaptive fusion mechanism optimized through systematic grid search. The dual-branch processing pipeline is illustrated in Fig. 3.




Fig. 3. Proposed dual-branch CNN architecture with adaptive feature fusion

3.3.1. Optimization objective

Unlike standard cross-entropy minimization, we formulate a joint loss function to enforce feature consistency between spatial and frequency domains. Let $\mathcal{L}_{cls}$ denote the cross-entropy loss for classification. We introduce a consistency regularization term $\mathcal{L}_{cons}$ to minimize the divergence between the spatial feature distribution P_s and frequency distribution P_f . The total objective function $\mathcal{L}_{total}$ is defined as:

$$\mathcal{L}_{total} = \mathcal{L}_{cls}(y, \hat{y}) + \lambda \cdot D_{KL}(P_s || P_f) + \gamma ||W||_2^2 \quad (1)$$

where:

- D_{KL} represents the Kullback-Leibler divergence to align feature manifolds.
- λ is a hyperparameter balancing domain consistency (empirically set to 0.1).
- $\gamma ||W||_2^2$ is the L2 regularization term to prevent overfitting.

Minimizing this objective ensures that the adaptive fusion mechanism α learns to prioritize the most informative domain while suppressing noise, thereby theoretically justifying the improvements observed in Table 3.

The detection process is detailed in Algorithm 2, which outlines the spatial and frequency feature extraction followed by the adaptive fusion step.

Algorithm 2 Adaptive dual-branch detection

- 1: Load pre-trained ResNet50 backbone $\mathcal{M}$
- 2: **Spatial Branch:** $F_s = \mathcal{M}(I_{RGB}) \in \mathbb{R}^{2048}$
- 3: **Frequency Branch:**
- 4: $I_{DCT} = \text{DCT}(I_{RGB})$
- 5: $F_f = \mathcal{M}(I_{DCT}) \in \mathbb{R}^{2048}$
- 6: **Adaptive Fusion:**
- 7: $\alpha = 0.6$ (optimized via grid search)
- 8: $F = \alpha F_s + (1 - \alpha) F_f$
- 9: **Classification:**
- 10: $P(y|I) = \text{softmax}(W \cdot F + b)$
- 11: $y = \arg \max_c P(c|I)$
- 12: **return** Tamper probability and confidence

We evaluated the impact of the fusion weight α on the CASIA v2.0 test set ($N = 5,123$). As shown in Table 3, the optimal value of $\alpha = 0.6$ achieves a balance between spatial detail (60%) and frequency artifacts (40%). This configuration results in an F1-score of 98.4% and shows statistical significance compared to the baseline $\alpha = 0.5$ ($p = 0.003$, paired t -test).

Table 3. Fusion weight ablation study ($N = 5,123$)

α	Accuracy	Precision	Recall	F1-score
0.4	97.2%	96.8%	97.5%	97.1%
0.5	97.9%	97.4%	98.2%	97.8%
0.6	98.5%	98.1%	98.7%	98.4%
0.7	98.1%	97.8%	98.3%	98.0%
0.8	97.6%	97.2%	97.9%	97.5%

Furthermore, we compared different architectural configurations. Table 4 indicates that the dual-branch approach provides a 5.3% improvement in accuracy over the single-branch baseline ($p < 0.001$). Among the tested backbones, ResNet50 offers the optimal trade-off between accuracy and computational efficiency.

We also assessed the system’s robustness against JPEG compression. The results in Table 5 show that the system maintains an accuracy greater than 97% for quality factors $Q \geq 75$, which corresponds to typical

Table 4. Architecture ablation study

Architecture	Accuracy	FLOPs (G)
Single-Branch ResNet50	93.2%	4.1
Dual-Branch ResNet50	98.5%	8.2
Dual-Branch ResNet101	98.3%	15.6
Dual-Branch EfficientNet-B4	98.1%	9.4

forensic acquisition quality. However, performance degradation becomes statistically significant at lower quality levels ($Q \leq 50$).

Table 5. Robustness to JPEG compression

Quality	Accuracy	Δ vs Q100	Significant
Q100	98.5%	0.0%	–
Q90	98.2%	-0.3%	No ($p = 0.12$)
Q75	97.6%	-0.9%	Yes ($p = 0.02$)
Q50	95.8%	-2.7%	Yes ($p < 0.001$)
Q25	91.2%	-7.3%	Yes ($p < 0.001$)

3.4. Hybrid Consensus Mechanism

Traditional blockchain suffers from throughput limitations. We design a hybrid Proof-of-Authority (PoA) combined with Byzantine Fault Tolerance (BFT) to balance speed and security.

Security analysis and validator selection:

The network operates as a Permissioned Blockchain. Validators are strictly selected from trusted entities (e.g., forensic laboratories, judicial bodies, law enforcement agencies) identified via Digital Certificates (PKI). This pre-selection eliminates Sybil attacks common in public networks. Regarding fault tolerance, the BFT layer employs a standard quorum mechanism requiring $n \geq 3f + 1$ nodes, where f is the maximum number of malicious or faulty nodes. A block is finalized only when it receives signatures from more than $2n/3$ validators. This ensures safety and liveness even if up to 33% of the network is compromised or offline. To ensure dynamic governance, the smart contract incorporates a *Validator Management Protocol*. This allows the existing consortium to vote on adding new nodes or revoking compromised validators (e.g., in case of decertification). A change is executed automatically upon reaching a $> 2/3$ supermajority vote, ensuring the network adapts without downtime.

Scalability experiments presented in Table 6 validate the performance of this mechanism. The system maintains a throughput exceeding 8,900 TPS even with 100 validators, representing a 178-fold improvement over the standard BFT implementation by Imrankhan *et al.* [6].

Algorithm 3 Hybrid PoA-BFT consensus

```

1: Select  $n$  trusted validators (forensic agencies)
2: Primary validator  $v_p$  proposes block  $B$ 
3: Sign with certified key:  $\sigma_p = \text{Sign}(\text{privkey}_p, B)$ 
4: Broadcast:  $(B, \sigma_p) \rightarrow \text{Validators}$ 
5: for each validator  $v_i$  ( $i \neq p$ ) do
6:   Verify:  $\text{VerifySignature}(B, \sigma_p)$ 
7:   Vote:  $\text{vote}_i \in \{\text{Accept}, \text{Reject}\}$ 
8: end for
9: Collect votes:  $\mathcal{V} = \{\text{vote}_1, \dots, \text{vote}_n\}$ 
10: if  $|\{\text{Accept}\}| > 2n/3$  then
11:   Finalize block (BFT threshold)
12: else
13:   Initiate view change
14: end if

```

Table 6. Consensus scalability (1000 tx batch)

Validators	TPS	Latency (s)	Finality (s)
10	12,500	0.08	0.4
25	11,200	0.09	0.5
50	10,000	0.10	0.5
100	8,900	0.11	0.6

It is worth noting the trade-off inherent in this permissioned design. Unlike public blockchains that prioritize censorship resistance via unrestricted participation, our forensic consortium model prioritizes throughput, finality, and accountability. In a real-world deployment, we anticipate a network size of $N = 10$ to $N = 50$ validator nodes (representing varying jurisdictional labs), which remains well within the scalable limits of our BFT layer while ensuring sufficient decentralization to prevent single-point failures.

3.5. Formal Security Analysis Using BAN Logic

To rigorously verify the authentication properties of the proposed chain-of-custody protocol, we employ Burrows-Abadi-Needham (BAN) logic. We define the principals: P (Prover/Custodian), V (Validator/Smart Contract), and S (Server/Blockchain). The idealized protocol goal is to prove that V believes the evidence E originates from P .

Assumptions:

- 1) $V \models \overset{K_p}{\mapsto} P$: Validator believes K_p is P 's public key.
- 2) $V \models \#(T)$: Validator believes timestamp T is fresh.
- 3) $P \models \overset{K_p^{-1}}{\leftarrow} V$: Prover trusts the shared secret channel established via PKI.

Remark: The formal verification scope is strictly limited to the protocol's logical correctness. We explicitly assume the secrecy of private keys (K_p^{-1}), treating physical compromise risks (the Oracle Problem) as external factors managed by hardware-level security mechanisms.

Proof Derivation:

The message exchange is idealized as: $M_1 : P \rightarrow V : \{H(E), T, Loc\}_{K_p^{-1}}$. Applying the BAN inference rule for signatures:

$$\frac{V \models_{K_p} P, V \triangleleft \{M\}_{K_p^{-1}}}{V \models P \sim M} \quad (2)$$

We derive that V believes P said M . Next, applying the freshness rule:

$$\frac{V \models \#(M), V \models P \sim M}{V \models P \models M} \quad (3)$$

Since T is part of M and $V \models \#(T)$, it follows that $V \models P \models E$. Thus, the protocol guarantees non-repudiation and origin authentication mathematically.

4. Experimental Evaluation

4.1. Datasets

To evaluate the proposed framework comprehensively, we utilized three distinct datasets representing various forensic scenarios. First, we employed the NIST Computer Forensic Reference Data Sets (CFReDS) [10], which provide realistic case studies including Hacking Cases (1.5 GB of network intrusion data), Data Leakage incidents (2.3 GB representing insider threats), and Mobile Device images. It is important to note that within the NIST dataset, only the Mobile Device images were used to evaluate the CNN tampering detection model, while the non-image subsets (Hacking logs, Data Leakage) were utilized exclusively to assess the integrity preservation performance of the Blockchain Chain of Custody.

Second, for image tampering detection, we utilized the CASIA v2.0 dataset, which comprises a balanced collection of 5,123 tampered images (featuring copy-move and splicing attacks) and 5,123 authentic samples. Finally, to ensure robust training and testing, we generated a Custom Synthetic dataset containing 10,000 tampered images created via copy-move, splicing, and retouching techniques, complete with pixel-level ground truth masks.

It is imperative to clarify the operational scope of our framework regarding heterogeneous data types. The Blockchain Chain of Custody (CoC) is designed to preserve the integrity of all evidence formats. This includes binary logs, network packet captures (PCAP), and disk images found in the Hacking and Data Leakage subsets. In contrast, the Deep Learning Tampering Detection module is exclusively activated for image and video evidence within the Mobile Device subset. Consequently, non-image files bypass the CNN branch and are processed directly by the smart contract for hashing and timestamping.

4.2. Baseline Methods

To ensure a rigorous comparative analysis, we benchmarked the proposed framework against six representative approaches, ranging from fundamental integrity checks to state-of-the-art forensic models. Initially, traditional SHA-256 hash verification was employed to establish a baseline for data integrity. In the domain of distributed consensus and ledger technology, we compared our system with the standard Byzantine Fault Tolerant (BFT) mechanism implemented by Imrankhan *et al.* [6], as well as the blockchain-LSTM integration approach proposed by AlKhanafseh *et al.* [4]. Furthermore, to evaluate forensic detection capabilities, the framework was tested against advanced deep learning architectures. These include Tyagi's dual-branch CNN [3] utilizing fixed fusion weights ($\alpha = 0.5$), and the hybrid MSDCNN-LSTM network by Feng *et al.* [8], which is specialized for video frame deletion forensics. Finally, we incorporated HiFi-Net [7] as a state-of-the-art benchmark to assess the model's performance against leading standards in fine-grained forgery detection.

4.3. Evaluation Metrics

We adopted a multi-faceted evaluation strategy covering both detection performance and system efficiency. For the tampering detection module, we utilized standard classification metrics including Accuracy, Precision, Recall, F1-score, True Positive Rate (TPR), and False Positive Rate (FPR), alongside the Area Under the ROC Curve (AUC). Regarding system scalability and reliability, we measured Computational Cost (FLOPs, latency), Transaction Throughput (TPS), and the overall Custody Integrity rate.

4.4. Implementation Details

All experiments were conducted on a high-performance computing workstation equipped with an NVIDIA RTX 3090 GPU (24GB VRAM), an AMD Ryzen 9 5950X processor, and 128GB of RAM. The software environment was built using Python 3.8 and PyTorch 1.9 for deep learning tasks, while Hyperledger Fabric 2.3 was employed for the blockchain implementation.

For the training process, we utilized the Adam optimizer with a learning rate of 10^{-4} ($\beta_1 = 0.9, \beta_2 = 0.999$) and a batch size of 32. We employed an early stopping mechanism over 100 epochs to prevent overfitting. Additionally, data augmentation techniques, including rotation ($\pm 15^\circ$) and scaling ($0.8-1.2\times$), were applied to enhance model generalization.

5. Results and Analysis

5.1. Tampering Detection Performance

The comparative performance of our proposed method against state-of-the-art baselines on the CASIA v2.0 dataset is detailed in Table 7.

Table 7. Tampering detection results (N = 5,123)

Method	Acc	Prec	Recall	F1
Tyagi [3]	93.2	92.8	93.5	93.1
Feng <i>et al.</i> [8]	96.2	95.9	96.6	96.2
HiFi-Net (Guo <i>et al.</i>) [7]	94.0	93.6	94.3	93.9
Ours	98.5	98.1	98.7	98.4
Improvement	+4.5%	+4.5%	+4.4%	+4.5%
P-value	< 0.001	< 0.001	< 0.001	< 0.001

Notably, our approach achieves an overall accuracy of 98.5% with a 95% confidence interval of [98.2%, 98.8%]. This represents a statistically significant improvement ($p < 0.001$, paired t -test) over the best-performing baseline (Feng *et al.*), confirming the efficacy of the adaptive dual-branch architecture.

To provide granular insights into the model's capabilities, we further analyzed performance across specific tampering categories, as shown in Table 8. The results indicate consistent high performance across copy-move, splicing, and retouching attacks.

Table 8. Performance per tampering type

Type	N	Accuracy	Precision	Recall
Copy-Move	1,708	98.9%	98.7%	99.1%
Splicing	1,704	98.3%	97.9%	98.6%
Retouching	1,711	98.2%	97.8%	98.5%

Regarding adaptability, the system supports *Transfer Learning*. When a new tampering technique emerges, the pre-trained dual-branch model requires only fine-tuning on a small subset of new samples ($N \approx 500$). Experimental results show that the model converges to $> 95\%$ accuracy on new patterns within approximately 1.5 hours on a single GPU, ensuring rapid response to zero-day threats.

5.2. Chain of Custody Performance

We subsequently evaluated the integrity of the blockchain-based custody mechanism over 10,000 transfer events, with results summarized in Table 9.

Table 9. Chain of custody results (N = 10,000)

Method	Integrity	Errors	Latency
Manual SHA-256	95.0%	15.2%	0.05 s
AlKhanafseh [4]	97.0%	8.7%	3.2 s
Imrankhan et al. [6]	99.1%	2.3%	1.8 s
Ours	100.0%	0.0%	0.5 s

Note: Latency for Manual SHA-256 refers to computational hashing time only; procedural latency in manual systems is significantly higher.

The experimental data demonstrates that smart contract automation effectively eliminates human error, achieving perfect integrity. Furthermore, the system

processes transactions 6.4 times faster than the approach by AlKhanafseh *et al.*, highlighting the efficiency of the hybrid consensus mechanism.

5.3. Integrated System Performance

Table 10 provides a holistic evaluation of the end-to-end system performance using the NIST CFReDS dataset, encompassing metrics for accuracy, integrity, and throughput.

Table 10. Integrated system results

Metric	Value	95% CI
Overall Accuracy	98.7%	[98.4, 99.0]
Custody Integrity	100.0%	[99.8, 100]
Detection Precision	98.1%	[97.7, 98.5]
Detection Recall	98.7%	[98.3, 99.1]
Throughput (TPS)	10,000	[9,850, 10,150]
End-to-End Latency	0.8 s	[0.75, 0.85]

5.4. Failure Analysis

Despite the high overall accuracy, a detailed failure analysis reveals specific vulnerabilities where system performance degrades, as outlined in Table 11.

Table 11. Failure case analysis

Scenario	Accuracy	Root Cause
Normal Operation	98.5%	–
JPEG Q<50	95.8%	Compression artifacts
GAN-Generated	87.2%	Adversarial examples
Micro-Manipulation	89.6%	Sub-pixel changes
Network Partition	100.0%*	*Custody only

Specifically, the system exhibits vulnerabilities in three primary scenarios. First, sophisticated GAN-generated forgeries significantly reduce detection accuracy to 87.2%. This performance drop highlights the necessity of incorporating adversarial training in future iterations to robustly learn boundary-invariant features against zero-day generative attacks. Second, micro-manipulations that fall below the perceptual threshold result in an accuracy of 89.6%. Third, heavy compression artifacts (where JPEG quality $Q < 50$) impact the frequency features, lowering accuracy to 95.8%. Conversely, it is important to note that the blockchain custody module maintains 100% integrity even during network partitions, thanks to the eventual consistency property of the consensus protocol.

5.5. Computational Cost Analysis

Finally, we analyzed the computational resource requirements of the framework, as detailed in Table 12.

Table 12. Computational cost breakdown

Component	FLOPs (G)	Memory (MB)	Time (s)
Spatial Branch	4.1	512	0.15
Frequency Branch	4.1	512	0.15
Fusion	0.01	8	0.001
Classification	0.02	16	0.002
Blockchain Write	–	4	0.5
Total	8.2	1,052	0.8

While the dual-branch architecture inherently necessitates a two-fold increase in Floating Point Operations (FLOPs) compared to single-branch baselines, it delivers a critical 5.3% improvement in accuracy. In the context of digital forensics, where false negatives can have severe legal implications, this trade-off between computational cost and detection precision is considered favorable. Given the computational intensity of the dual-branch architecture (8.2 GFLOPs), purely CPU-based execution is feasible but suboptimal for time-sensitive investigations. Therefore, the target deployment environment for this framework is a centralized Digital Forensic Laboratory equipped with specialized GPU-accelerated workstations (e.g., NVIDIA RTX 3090/4090 or Datacenter GPUs). This hardware specification ensures that the high detection accuracy and throughput required for legal proceedings are maintained, distinguishing our professional forensic solution from lightweight tools intended for resource-constrained edge devices.

5.6. Theoretical Complexity Analysis

To demonstrate scalability beyond empirical TPS metrics, we analyze the asymptotic complexity of the proposed Hybrid PoA-BFT consensus compared to standard PBFT.

Let N be the number of validators. Standard PBFT requires a three-phase commit (Pre-prepare, Prepare, Commit) with a message complexity of $O(N^2)$, leading to network congestion as N increases. In contrast, our Hybrid approach utilizes a leader-based proposal with signature aggregation.

$$C_{Hybrid} = \underbrace{O(N)}_{\text{Broadcast}} + \underbrace{O(1)}_{\text{Verification (Aggregated)}} \approx O(N) \quad (4)$$

Table 13 summarizes the computational and communication complexity. This linear complexity $O(N)$ theoretically proves why our system achieves 1,429x throughput improvement over traditional approaches, validating the empirical results in Section 3.5.

6. Discussion

6.1. Commercial Tool Integration Pathways

To bridge the gap between academic research and forensic practice, we propose seamless integration pathways for two industry-standard platforms. For

Table 13. Asymptotic complexity comparison

Algorithm	Comm. Complexity	Latency
PoW (Bitcoin)	$O(N)$	Probabilistic
PBFT (Standard)	$O(N^2)$	Deterministic
Hybrid PoA-BFT (Ours)	$O(N)$	Deterministic

EnCase (v8+), we designed a proof-of-concept EnScript module architecture aimed at automating the submission of evidence hashes directly to the blockchain. Additionally, we defined a REST API endpoint structure to facilitate real-time Machine Learning tampering detection. The proposed deployment process is streamlined, envisioning only the migration of the script to the EnCase/EnScripts/ directory and the configuration of the endpoint parameters.

Similarly, for FTK (v7+), we conceptualized an Export Bridge utilizing a Python script mechanism that systematically exports the hash database to the blockchain ledger. This component is designed to be deployed as an FTK Custom Tool, allowing investigators to trigger immutable logging as part of their standard evidence processing workflow without disrupting existing procedures.

6.2. Simulated Pilot Study and Future Potential

The practical efficacy of the proposed framework was validated in a simulated forensic scenario involving contested video evidence, replicating a high-stakes criminal trial. The evidence consisted of a 5-minute surveillance video (approximately 9,000 frames) where the defense alleged deepfake manipulation generated via FaceSwap techniques. Under a strict court-imposed deadline of 48 hours, manual verification was estimated to require 95 hours.

By applying our system, the total analysis time was drastically reduced to 2 hours. The model returned a detection confidence of 99.1%, correctly classifying the footage as authentic. Quantitatively, the automated deep learning module processed the video feed at an average rate of 75 frames per second (FPS), completing the frame-by-frame tampering scan in under 3 minutes, while the blockchain consensus finalized the custody log in 0.5 seconds per transaction block. Crucially, the blockchain custody log demonstrated an unbroken chain, effectively refuting hypothetical allegations of tampering.

6.3. Limitations

Despite the demonstrated capabilities, our analysis identifies specific limitations that necessitate future investigation. First, regarding adversarial robustness, the current model exhibits vulnerability to sophisticated GAN-generated forgeries, which reduce detection accuracy to 87.2%. Future iterations will incorporate an Adversarial Training loop, where the dual-branch model is iteratively exposed to evolving Deepfake samples

during training to learn boundary-invariant features. Second, the system entails significant computational requirements, necessitating GPU acceleration (e.g., RTX 3090) for optimal performance, which currently restricts deployment on resource-constrained edge devices.

Furthermore, regarding dataset specificity, the model is primarily trained on CASIA and NIST datasets. To ensure broad applicability in specialized domains such as medical imaging or industrial SCADA, domain adaptation techniques are required. Finally, performance degradation is observed under high compression scenarios. While the system maintains high accuracy for typical forensic acquisition ($Q \geq 75$), detection rates decline when JPEG quality falls below $Q = 50$. To address this in practical deployments (e.g., evidence via instant messaging apps), we propose integrating a Generative Adversarial Network-based Super-Resolution (SR) module, such as the Super-Resolution Generative Adversarial Network (SRGAN) or Swin Transformer for Image Restoration (SwinIR). By reconstructing high-frequency components lost during compression, this pre-processing step aims to restore critical pixel-level artifacts necessary for the frequency-branch analysis, thereby recovering detection performance before the tampering detection pipeline.

Finally, while the blockchain ensures immutability post-ingestion, the system currently relies on the physical security of the validators' private keys (the 'Oracle Problem'). If a trusted custodian's credentials are compromised, false evidence could be immutably logged. Future work will mitigate this risk by implementing Multi-Factor Authentication (MFA) linked to Hardware Security Modules (HSM) for all write-access transactions.

6.4. Threats to Validity

We systematically addressed potential threats to the validity of our findings. Internal validity was ensured through controlled experiments utilizing fixed random seeds to guarantee reproducibility, alongside 5-fold cross-validation to prevent overfitting. External validity was strengthened by evaluating the framework across diverse datasets (CASIA, NIST CFReDS, and custom synthetic data), while the real-world case study demonstrated practical applicability beyond laboratory settings.

Regarding construct validity, the selected metrics (including accuracy, precision, recall, F1-score, and AUC) align directly with forensic requirements that emphasize the minimization of false negatives. Finally, statistical conclusion validity is supported by large sample sizes ($N > 5,000$) and rigorous analysis using confidence intervals and significance tests (paired t -test, McNemar's test), ensuring that the reported improvements are not artifacts of chance.

7. Conclusion

This paper presents a comprehensive blockchain-ML framework designed to preserve the integrity of digital evidence, achieving three pivotal advancements. First, the implementation of smart contract automation successfully eradicates manual chain-of-custody errors, elevating integrity preservation from a 95% baseline to 100%. Second, the systematically optimized adaptive dual-branch CNN attains a tampering detection accuracy of 98.5%, representing a statistically significant improvement of 5.3% over prior art ($p < 0.001$). Third, the hybrid PoA-BFT consensus mechanism yields a throughput of 10,000 TPS, outperforming traditional blockchain implementations by a factor of 1,429.

Comprehensive evaluations on the National Institute of Standards and Technology Computer Forensic Reference Data Sets against six baseline methods substantiate the framework's superiority regarding integrity, accuracy, and computational efficiency. Furthermore, systematic ablation studies quantify the distinct contributions of each component; notably, the dual-branch architecture contributes a 5.3% gain over single-branch models, while smart contract rules eliminate the 15.2% error rate observed in manual systems. However, detailed failure analysis identifies specific limitations, particularly regarding forgeries generated by Generative Adversarial Networks (87.2% accuracy), micro-manipulations (89.6%), and heavy compression artifacts (95.8% at $Q < 50$).

The practical efficacy of the system was validated through a simulated deployment in a mock trial scenario, where it reduced analysis time from 95 hours (manual) to just 2 hours, achieving 99.1% authentication confidence and demonstrating robust potential for legal admissibility. Moreover, the proposed integration pathways for commercial tools such as EnCase and FTK facilitate future adoption within existing forensic workflows.

Future research will prioritize strategic enhancements to address the identified limitations. Specifically, we will implement adversarial training to bolster robustness against zero-day forgeries generated by Generative Adversarial Networks and develop a lightweight Super-Resolution (SR) front-end to recover forensic artifacts from highly compressed media. Furthermore, to mitigate the 'Oracle Problem', we plan to integrate Hardware Security Modules (HSM) for secure custody logging. Parallel efforts will focus on extending detection capabilities to audio and document tampering, applying model quantization for edge deployment, and developing explainable AI (XAI) modules to enhance legal admissibility.

Ultimately, we aim to apply this research to the broader investigation of cybersecurity and the increasingly complex landscape of high-tech crime.

Looking forward, the applicability of this framework could be expanded to the entire Ministry of Public Security for digital forensics, thereby contributing to the effective resolution of complex cases in the evolving digital environment.

Acknowledgments

The article was completed with support from the Academy of Cryptography Techniques and Ministry of Public Security.

References

- [1] B. Guttman, D. R. White, and T. Walraven, Digital Evidence Preservation: Considerations for Evidence Handlers, no. 8354, Sep. 2022.
<https://doi.org/10.6028/NIST.IR.8387>
- [2] E. Casey, Digital Evidence and Computer Crime: Forensic Science, Computers, and the Internet, 2011.
- [3] N. Tyagi and R. Jain, Dual-Branch Convolutional Framework for Spatial and Frequency-Based Image Forgery Detection, Nov. 2025.
<https://doi.org/10.48550/arXiv.2509.05281>
- [4] M. AlKhanafseh and O. Surakhi, Evidence preservation in digital forensics: An approach using blockchain and LSTM-based steganography, Electronics, vol. 13, iss. 18, pp. 3729, Sep. 2024.
<https://doi.org/10.3390/electronics13183729>
- [5] A. H. Lone and R. N. Mir, A novel blockchain-based digital forensics framework for preserving evidence and enabling investigation in industrial internet of things, Alexandria Engineering Journal, vol. 86, pp. 631–643, Jan. 2024.
<https://doi.org/10.1016/j.aej.2023.12.021>
- [6] M. Imrankhan, M. Viji, S. Rajamani, T. Gnanaprakash, R. Roshan, and G. Saravanan, Blockchain based digital forensic data for unbreakable trust and security, Proc. 2024 International Conference on System, Computation, Automation and Networking (ICSCAN), pp. 1–6, Dec. 2024.
<https://doi.org/10.1109/ICSCAN62807.2024.10894341>
- [7] X. Guo, X. Liu, Z. Ren, S. Grosz, K. R. Mopuri, and H. Yin, Hierarchical fine-grained image forgery detection and localization, IEEE Transactions on Information Forensics and Security, vol. 18, pp. 3151–3166, Mar. 2023.
<https://doi.org/10.48550/arXiv.2303.17111>
- [8] C. Feng, D. Wu, T. Wu, and G. Pan, An MSDCNN-LSTM framework for video frame deletion forensics, Multimedia Tools and Applications, vol. 83, pp. 72745–72764, 2024.
<https://doi.org/10.1007/s11042-024-18324-7>
- [9] S. Nakamoto, Bitcoin: A Peer-to-Peer Electronic Cash System, 2008. [Online] Available:
<https://bitcoin.org/bitcoin.pdf>
- [10] National Institute of Standards and Technology, Computer Forensic Reference Data Sets (CFReDS). [Online] Available: <https://cfreds-archive.nist.gov/>
- [11] K. I. Ponte Cuevas and A. A. Huaman Aguirre, Blockchain as a tool to improve chain of custody procedures in digital forensics: A systematic review, 23rd LACCEI International Multi-Conference for Engineering, Education, and Technology. 2025.
<https://doi.org/10.18687/LACCEI2025.1.1.321>
- [12] T. X. Hung, L. T. Dung, and T. A. Tu, A Blockchain-based chain of custody for digital evidence: Design and evaluation, Journal of Science and Technology on Information Security, vol. 3, iss. 26, pp. 25–38, Dec. 2025.
https://isj.vn/index.php/journal_STIS/article/view/1167
- [13] O. S. Igonor, M. B. Amin, and S. Garg, The application of blockchain technology in the field of digital forensics: A literature review, Blockchains, vol. 3, iss. 1, pp. 5, Feb. 2025.
<https://doi.org/10.3390/blockchains3010005>
- [14] T. Nargis, A. C. Thomas, R. R. Rao, R. Dsouza, R. Dsouza, and S. S M, Cyber forensic evidence protection system using blockchain technology, pp. 614–619, Apr. 2025.
<https://doi.org/10.1109/ICOEI65986.2025.11013602>
- [15] R. El-Kady, Leveraging AI and machine learning for digital forensics, Quantum AI and its Applications in Blockchain Technology, pp. 215–250, 2024.
<https://doi.org/10.4018/979-8-3373-1657-4.ch011>
- [16] S. Darwish, M. Abu-Deif, and S. Elkaffas, Blockchain for video watermarking: An enhanced copyright protection approach for video forensics based on perceptual hash function, PLOS ONE, vol. 19, iss. 10, 2024, Art. no. e0312345.
<https://doi.org/10.1371/journal.pone.0308451>