Grenze International Journal of Engineering and Technology, June Issue

Beyond the Hype: A Novel Hybrid Architecture for Layer-2 Scalability Integration & Privacy Preserving Computations

Sumit Chauhan[1], Hirdesh Sharma[2], Dhruv Khandelwal[3], Shivam Khandelwal[4], Saurav Singh Bisht[5] and Dharmendra Kumar Mahato[6]
[1]Chief Executive Officer, Datopic Technologies Pvt. Ltd, Noida, (U.P. – 201307), India
Email: sumit@datopic.com
[2]Associate Professor, Dronacharya Group of Institutions, Greater Noida, UP, India
Email: hirdesharma@gmail.com
[3-6]Dronacharya Group of Institutions, Greater Noida, UP, India
Email: khandelwaldhruv33@gmail.com, Shivamkhandelwal641@gmail.com, Sauravs1969@gmail.com, dharmendrak3327mahato@gmail.com

***Abstract*— Blockchain as a platform to support health care systems in their goal of transparency, security, and the patient as focus is a very active area of study; however, real-world adoption remains limited by several factors, such as: scalability, interoperability and regulatory compliance. The conflict between immutable ledgers and regulation, such as GDPR, is major contributors to this issue. These challenges are a result of the rigid design of first-generation health care blockchain architectures rather than a problem with the underlying blockchain technology itself. This research introduces a multi-layer "thin blockchain" architecture that will address these limitations. This thin blockchain will include a permissioned Layer 1 for identity and trust, a zero-knowledge rollup (ZK-Rollup), Layer 2 for processing a large number of transactions, and an encrypted off-chain data layer that will be compliant with HL7 FHIR. The off-chain data layer will utilize distributed storage systems such as IPFS. Recursive ZKP mechanisms will also be used to improve the off-chain integrity, including the FAITH protocol, which permits the verification of large multimedia datasets efficiently. Additionally, the overall architecture will provide a scalable, interoperable, and Compliant path forward for the implementation of blockchain in health care systems. Privacy-focused applications, such as Blockchain-enabled federated learning (BCFL), and ZKP-verified LLM advisory systems will also be supported within the architecture.**

***Index Terms*— Blockchain, Distributed Ledger Technology (DLT), Self-Sovereign Identity (SSI), Healthcare.**

I. INTRODUCTION

A. The Digital Imperative and the Data Crisis in Healthcare

Data generated by the rapid digitization of healthcare is numerous; they include structured electronic health records (EHRs); high resolution medical images. continuous IoT streams; genomic information [1].

Grenze ID: 01.GIJET.12.2.82

Centralized systems for traditional use in specific institutions cannot provide the support necessary for managing these interconnected data sources. As a result, these systems produce disjointed, insecure data silos that limit clinical decisions, slow research, and create operational inefficiencies [2]. In addition, these centralized databases have been a primary target for hackers. Consequently, there has been an increased number of data breaches resulting from cyber-attacks, compromising the confidentiality and security of millions of patients' records. Therefore, there exists a key issue: healthcare can benefit from sharing data to enhance patient care and accelerate research; however, it must also protect the data to maintain patient privacy.

B. DLT — Opportunity for Healthcare

Distributed Ledger Technology has been proposed as one possible answer to the above Conflict [3]. Its four main characteristics — decentralization, cryptographically secured, immutable, transparent audit trail — make it a potential solution for patient-centric data sharing. Blockchain has been viewed as a viable method to track all accesses to a patient's data; to prove data provenance; and to enable detailed patient-driven control over access permissions within a distributed trust environment.

C. The Gap — Academic Models vs Real World Developments

Despite many studies of the potential for using blockchain in healthcare and notable examples of prototype implementations (e.g., MIT's MedRec), very few Healthcare blockchain systems have been implemented and used in practice. Many Some of the early designs were proof-of-concepts that demonstrated functionality under controlled conditions but did not address the needs of practicing clinicians. The literature identifies three principal reasons for the low level of implementation:

- Scalability limitations, including low throughput and long response time on permissioned networks.
- Regulatory challenges
- Issues of semantic interoperability, in which systems were developed without a standardized data structure, resulting in additional secure data silos separate from those previously existing.

A recent review also highlighted the fact that most Early model designs utilized none of the improved scalability techniques available today (such as Layer-2 Rollups or Sharding) to highlight the developmental shortcomings of the first generation of designs.

D. Dissertation: Developing a Hybrid, Multi-Layered Architecture

The shortcomings of earlier designs demonstrate architectural design flaws rather than inherent problems with blockchain technology. It is impractical technologically, economically, and legally to store large amounts of sensitive medical data directly on the blockchain. This paper will propose a multi-layered A hybrid architecture that integrates solutions based on blockchain engineering, Zero-Knowledge cryptography and health informatics. The architecture is composed of a "thin" blockchain comprising three layers:

Layer 1, Trust Layer: A permissioned blockchain for identities (DIDs), access rights, and settlements [6].

Layer 2, Scaling Layer: A ZK-Rollup system that performs high-frequency transactions off-chain while maintaining cryptographic guarantees.

Off-Chain Data Layer: Patient data compliant to HL7 FHIR standards are encrypted and stored in decentralized systems such as IPFS, integrity of the data is ensured via robust ZKP mechanisms [7]. This framework provides a foundation for developing an efficient, interoperable, regulation-compliant healthcare blockchain ecosystem.

II. The Evolving Landscape of Blockchain in Healthcare

To build a solution, one must first understand the problem area and the progress made in past efforts. This section reviews the key application areas for blockchain in healthcare and examines the important architectural change that has shaped the field.

A. Core Application Domains

Patient-Centricity: Secure EHRs and Self-Sovereign Identity (SSI)

The EHR application remains the leading application; however, traditional systems continue to separate patient information across different provider-controlled areas, limiting both patient ownership and continuity of care. The early blockchain models, such as MIT's MedRec model, proposed a patient-centered approach. The MedRec model utilized cryptographic keys that allowed individuals to manage access to their records rather than institutions. The concept of SSI (Self Sovereign Identity) or DID (Decentralized Identity) has evolved from the earlier blockchain models. In this model, blockchain utilizes a globally unique and portable identifier to provide long-term management of data access and patient data.

TABLE I. CORE APPLICATION DOMAINS

Use Case	Core Problem Solved	Key Blockchain Features Utilized	Foundational Models / Examples
Secure Electronic Health Records (EHRs) & Self-Sovereign Identity (SSI)	Fragmented data silos; lack of patient data control; data breaches.	Decentralization, Cryptographic Access Control, and Patient-Controlled Consent [1, 1]	MedRec (MIT); Patient-Centric Architecture; Decentralized Identity (DID).
Pharmaceutical Supply Chain Integrity	Counterfeit pharmaceuticals, Lack of pedigree, and temperature chain violations.	Immutable, Transparent, Auditable Transaction Log.	Pharmatrust 1
Clinical Trial & Research Data Management	data integrity & falsification, inefficient management of consent, data-sharing silos.	immutability (Data Integrity), Smart Contracts (Automating Consent), Tokenization (Data Sharing Incentives).	Avaneer (IBM) 11

Provenance: Pharmaceutical Supply Chain Integrity

The next application area is Provenance: Pharmaceutical Supply Chain Integrity. The immutable nature of blockchain has shown significant promise for providing supply chain protection. Counterfeit medications, diverted stock, and issues with cold chain handling pose significant risks to patient safety. A blockchain-based provenance can track every transfer of product from manufacturer to distributor to pharmacy using an immutable, shared audit trail. Examples of real-world applications of this include both PharmaTrust and MediLedger. These companies have provided evidence of the practical applicability of tracking the authenticity of products and identifying issues in distribution.

Trust: Clinical Trial and Research Data Management

Lastly, there is the application area of Trust: Clinical Trial and Research Data Management. Clinical trials depend upon the accuracy and integrity of the data collected. The use of blockchain for clinical trials will provide an auditable history and prevent data tampering while also ensuring reproducibility. Smart contracts can simplify patient consent processes, and token-based systems can incentivize data sharing for privacy-focused research. An example of a company utilizing blockchain to support the integrity of clinical trials includes IBM's Avaneer network [9].

B. Early Systems and Architecture

In the early days of blockchain in healthcare, many of the first blockchain systems were either on-chain or quasi-on-chain storage. They maintained a majority of the patients' data and metadata in the ledger. However, prototypes such as MedRec attempted to establish a decentralized method of controlling access to the data; however, they had no clear definition as to what portion of the data would be stored on the ledger. As a result of this decision, there was one major issue that arose: the ability to store large medical datasets (such as MRI images, genomic sequences, etc.) on a blockchain was neither feasible nor cost-effective. The larger amounts of data resulted in slower block propagation times, decreased throughput rates, and created unsustainable demand for storage space for network participants. As a result, the industry shifted away from the "fat" blockchain model and moved towards a thin blockchain or hybrid model.

The thin blockchain or hybrid model separates functions. On-chain is limited to a very small amount of information that is needed to establish trust, identity (DIDs), access control, and metadata references (hashes). All other information, such as the encrypted patient data, is stored off-chain in high-performance external systems such as HIPAA-compliant cloud services or distributed networks such as IPFS.

This separation represents an acknowledgment that the true value of a blockchain lies not in the mass storage of data. Rather, it provides a stable trust anchor for identity, authentication, and provenance. This understanding forms the basis for the architecture described in this paper.

III. OBSTACLES TO DEPLOYMENT: AN ANALYSIS OF THE PROBLEM SPACE

The hybrid "thin blockchain" was able to address the too simplistic assumption that data would be stored on-chain; however, there were much deeper structural problems to be addressed.

There continue to be substantial technical, legal, and social barriers to the use of blockchain technology in the health care industry. There needs to be a systematic and integrated solution to address each of the three.

A. The Bottleneck of Scalability: Throughput, Latency, and Storage Problem

Even though the majority of the data has been removed from on-chain storage, healthcare ecosystems generate extremely large volumes of transactions. Millions of patients, thousands of healthcare providers, and billions of

IoT devices continuously perform access events, consent updates, and clinical logs, all of which stay on-chain and overwhelm standard L1 networks. As such, both public and permissioned blockchains have inherent limits on throughput, long confirmation periods, and expensive transaction processing. These limitations do not allow for real-time clinical workflows.

Research has shown that a significant proportion of the existing healthcare blockchain prototypes have not used modern scalability techniques (such as Layer-2 rollups, sharding, or batching), and thus the scalability problem remains largely unaddressed. Without a separate scalability layer, healthcare blockchain systems continue to exist at proof-of-concept scales.

B. The Myth of Interoperability: Why Blockchain Is Not a Solution to Data Silos

Many believe that blockchain provides for interoperability. However, blockchain does not understand the data it stores. If different organizations place their proprietary or incompatible data formats onto a blockchain, they will establish cryptographically secure data silos rather than interoperable systems.

In order to provide true semantic interoperability, there needs to be a standardization of data structures. In healthcare, the standard is HL7 FHIR. HL7 FHIR defines modular, API accessible resources, including Patient, Observation, and Medication Request. Without requiring compliance with FHIR before encryption and storage, it is impossible to share data between organizations using blockchain technology. Thus, any scalable architecture must require FHIR compliance as a prerequisite.

C. The Regulatory Paradox: Immutability vs. Compliance

The largest barrier is regulatory. Blockchain's immutability is at odds with many countries' privacy laws, particularly the GDPR "right to be forgotten," which requires that all personal data be erased. Previous versions of blockchain models were unable to solve these conflicting requirements.

HIPAA Security Requirements

HIPAA requires clear access controls, confidentiality, and data integrity. A well-designed hybrid architecture satisfies all of these requirements. Confidentiality: All ePHI continues to be encrypted and stored off-chain. Auditability: All grants of access, revocations of access, and access requests are permanently recorded on-chain. As such, blockchain technology offers advantages over traditional systems in terms of traceability [9].

GDPR: Implementing "Functional Deletion"

To balance GDPR with the immutable nature of blockchain, hybrid systems employ a two-step functional deletion process. First, the patient data is deleted from off-chain storage systems, such as IPFS. Second, the corresponding pointer or hash is deleted on-chain. The on-chain hash is then left as an orphaned, non-identifiable artifact, unlinked from all personal data. This approach complies with GDPR requirements while maintaining the integrity of blockchain.

D. Sociotechnical Obstacles: Adoptions, Trust, and Organization-Fit

Even if a technically perfect and legally compliant system is designed and built, it will likely fail to gain acceptance from stakeholders due to other sociotechnical obstacles.

Healthcare institutions are generally slow to adopt new technologies because they disrupt workflows, incur costs, and increase the risk of being first in line. Research utilizing the Technology-Organization-Environment (TOE) and Fit Viability Model (FVM) has shown:

Drivers of Adoption

- High levels of perceived security and transparency, High levels of organizational readiness, Leadership support and Strong workflow alignment ("fit") barriers to Adoption
- Lack of participant engagement (network effects), Poor vendor support, Perceived system complexity, and Negative reaction to centralized government-mandated implementation

These research findings indicate that successful implementation of blockchain technology for health care must occur through a collaborative governance model, based upon a consortium model, and cannot be accomplished through unilateral governmental mandate.

IV. A Novel Multi-Layered Hybrid Framework

For Scalable and Secure Health Data Exchange Section 3 has identified the challenges in health care, which cannot be solved by a single Layer-1 blockchain. The requirements of scalability, interoperability, and regulatory compliance cannot be satisfied by Layer-1 blockchain alone. A viable architecture will therefore need to share responsibilities among different layers. Our suggestion is a multi-layered hybrid architecture. This architecture is

shown in Figure 1 and contains a Layer 1 that is based on the minimal trust principle and a Layer 2 that uses a high-throughput scaling mechanism; additionally, it includes a Layer 3 that represents the off-chain data layer, which is based on open standards.

A. Layer 1: The Trust Layer (Main Chain)

Layer 1 (L1) provides the basic level of trust for this system. It works as a permissioned blockchain that uses an efficient consensus method, such as PBFT or PoS, and is operated by a group of trusted healthcare providers. In this architecture, L1 is intended to have a very limited scope:

- DID registry: This ensures a decentralized way of managing unique identifiers for patients, providers, and devices, which are essential for securely interacting with each other.
- Access control through smart contracts: This allows the enforcement of specific and revocable permissions, determining whether and under what conditions. Under certain conditions, certain FHIR resources can be accessed.
- Settlement and metadata validation: The L1 chain does not contain any medical data. It only anchors the identities, permissions, and integrity.

Layer 2: The Scaling Layer (ZK-Rollups)

This layer addresses the scalability issue. High-frequency, low stakes transaction move off the main chain and onto a Layer 2 (L2) network. For this purpose, we propose to use the Zero-Knowledge Rollups (ZK-Rollups). We chose ZK-Rollups because they represent the best of both worlds. They collect tens to hundreds or thousands of transactions — e.g., when a patient grants consent, a prescription is fulfilled, or an IoT device logs a measurement — from the main chain. Then, they calculate a single compact Zero-Knowledge Proof (ZKP), which verifies the validity of all the bundled transactions, and post this proof to the L1 chain. This enables the system to process a large number of transactions — often thousands per second — while at the same time, taking advantage of the full security and decentralization of the underlying L1 chain.

Off-Chain Layer

The Data Layer (Encrypted FHIR on IPFS) All medical data — including EHRs, images, and laboratory values — is stored in encrypted format off-chain. For this reason, HL7 FHIR is used as the semantic standard for representing all clinical data as interchangeable resources. Each encrypted FHIR resource is stored in distributed systems like IPFS. This generates a content-addressable ID (CID) for each file. The CID is recorded on the L1 so that integrity can be verified without having to store sensitive data on-chain. Additionally, deleting files in accordance with GDPR is possible by simply removing the off-chain file and cancelling the corresponding on-chain reference.

B. Detailed Description: Scalable Layer 2 solution (ZK-Rollups) Healthcare data access needs have a very direct, simple cause-and-effect nature:

High-volume events exceed Layer 1 throughput, leading to an immediate need for off-chain execution while maintaining trust. ZK-Rollups provide proofs of correctness. Each ZK proof can replace hundreds or thousands of Layer 1 transactions, thus enabling real-time use cases such as IoT telemetry ingestion, dynamic consent management, and multi-party care coordination. ZK-Rollups also support public auditability, providing transparency to regulators, providers, and auditors while protecting clinical information.

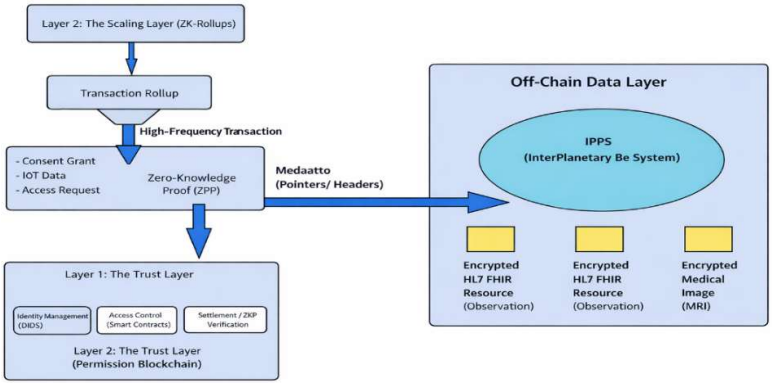

Figure.1 Proposed Multi-Layered Hybrid Architecture

C. Detailed Description

The Decentralized Data and Privacy Layer Offloading large amounts of data off-chain has improved scalability; however, a significant new problem has been created:
How can a user trust that a retrieved off-chain file is the true version referenced on chain without computing the heavy hash again?

A Framework for Off-Chain Data Integrity (the "FAITH" Model)

In order to address this new problem, the architecture will incorporate concepts from the FAITH framework. The FAITH framework uses recursive Zero-Knowledge Proofs to validate the integrity of large files.
Process:

- The client verifies this small proof in milliseconds.
- No full-file hashing is required at the user's end. The results of FAITH demonstrated a 98% faster validation of multi-GB medical images than previous methods. This is a critical aspect of clinical workflows.

Case Study: HL7 FHIR Resource Integration

FHIR is designed modularly, allowing each patient's clinical record to be composed of encrypted individual resources (e.g., Allergy Intolerance, Observation, Medication Request. Layer 1 will store pointers (CIDs) to each resource. This structure allows specific consent options to exist, such as:

- Grant READ access to Medication Request/id-123 for 24 hours

This model of patient-controlled access is the ultimate goal of the early systems, such as MedRec. It can only be realized by utilizing L1 trust, L2 scalability, and FHIR-compliant off-chain storage.

V. Future Directions for using the Architecture

Applications Beyond Secure Data Management Secure data management is just one aspect of what the architecture can do. Its design provides an extensible environment to use the architecture as a base for the future generation of privacy-preserving healthcare applications. The on-chain trust and scalability features of the architecture, combined with ZKPs, enable new forms of AI, identity, and analytics applications in healthcare that have never been possible before.

A. Privacy Preserving Analysis: BCFL

High-quality training data is required by most modern AI applications. Due to legal and privacy issues, however, no two health care organizations will share their patients' data with other hospitals. To solve this problem, Federated Learning (FL) was introduced. FL allows multiple hospitals to develop AI models that train on their local hospital data; they then share the trained AI model parameters with all other participating hospitals. FL relies on a central entity However, to aggregate the model parameter updates from each hospital. This centralized aggregator represents a potential single point of failure, has no inherent auditing capability, and has no method for providing incentives for proper behaviour among hospitals. Our Architecture's Enhancement: BCFL as a Decentralized Aggregator By using the Layer 1 trust layer as the centralized aggregator, FL transforms into BCFL:

- Parameter updates for the AI models are transmitted via Layer 2 and settled on Layer 1.
- Layer 1 maintains an auditable log of all transactions.
- Proof-based validation ensures malicious updates or "free riding" cannot occur.

Two current state-of-the-art architectures that are built upon the same principles as our architecture is:

- MCGP (2024): A blockchain-driven system for multi-continental glucose prediction that uses on-chain incentives and slashing mechanisms.
- BFEL (2025): An enhanced federated edge learning system that performs verifiable aggregation of AI model updates on-chain.

These two examples show that the architecture supports privacy-preserving, cross-institutional AI development.

B. Verifiable, Personalized Medicine: ZKPs + LLMs

Personalized medicine is increasingly supported through LLMs that help clinicians make medical decisions. However, patients' confidential data must be revealed for the LLM to generate personalized recommendations; this would represent a serious breach of their right to privacy.
Solution: ZKP Verified LLMs ZKP-LLM (2025) is an example of a framework that can allow:

- Patient's client-side application to perform computations using the client-side encrypted FHIR data.
- Zero-Knowledge Proof to be produced (example: "Patient is allergic to penicillin")

- LLM will verify the zero-knowledge proof and generate safe, personalized recommendations [8]

This mechanism removes the traditional privacy–utility trade-off for providing personalized care without transmitting the patient's underlying data.

C. Patient Identity

DID-Based Identity Management- Patient identity is the foundation of all access and authorization in the health system. Traditionally, patient identity management has been fragmented, resulting in multiple issues, including misidentification, duplicate records, and errors with respect to accessing patient data.

Layer 1 DID

Proposed Layer 1 Trust Layer will function as a secure DID Registry: Patient controls private key(s) tied to decentralized identifier(s), Each provider, payer, device, etc., will have a unique registration and All access control decisions will reference the DID Graph.

VI. CONCLUSION

Blockchain's utility in healthcare for more than a decade has been limited by architectural issues — not technological ones. The early systems had problems with scalability, semantic interoperability, and regulatory compliance. This resulted in many proofs-of-concept that were unable to transition from conceptual models to practical applications. This paper addresses these architectural challenges with a hybrid multi-layered architecture, which includes a permissioned trust layer; a ZK- Rollup layer for scalability, and an off-chain standards-compliant data layer. This paper describes what we refer to as a "thin blockchain" design, Trust and Identity — a Minimal Layer 1 DIDs, Access Control & Settlement, Scalability — Layer 2 ZK-Rollup System with High Throughput, Low Latency Processing, Interoperability — MANDATORY HL7 FHIR Data Structures to Avoid Cryptographic Silos, Regulatory Alignment — Off-Chain Encrypted Storage Enables GDPR Compliant Deletion & Strong HIPAA Auditing and, Privacy Preserving Computation — FAITH-style ZKPs in DAPPS for Efficient Integrity Verification & Support for Advanced Applications like BCFL, ZKP Verified LLMs.. Instead of asking if blockchain can enable secure electronic data exchange in the healthcare industry, we are now discussing how it can be effectively and responsibly implemented. It provides a solid foundation for the development of the next generation of patient-centric digital ecosystems.

REFERENCES

[1] Kasyapa, M. S. B., & Vanmathi, C. (2024). Blockchain integration in healthcare: a comprehensive investigation of use cases, performance issues, and mitigation strategies. Frontiers in Digital Health, 6:1359858.

[2] Michael, J., Cohn, A. L. A. N., & Butcher, J. R. (2018). Blockchain technology. The Journal, 1(7), 1-11.

[3] Cryptography and Network Security" by William Stallings, Fourth Edition.

[4] Abdalbasit Mohammed Qadir; NurhayatVarol"A Review Paper on Cryptography" IEEE, June 2019.

[5] Sharma, H., Pandey, B., Kumar, R., Mishra, S., Bhardwaj, M. K., & Pant, S. (2025). A Lag for Optimized Columnar Repositioning Methodology for Data Security. (GIJET), 11.

[6] Anjula Gupta, Navpreet Kaur Walia "Cryptography Algorithms: A Review" IJEDR, Volume 2 Issue 2, 2014.

[7] Khairun Nahar, Partha Chakraborty (2020) "Improved Approach of Rail Fence for Enhancing Security" International Journal of Innovative Technology and Exploring Engineering 9(9):583-585.

[8] Watanabe, H., & Uchikoshi, M. (2025, April). Generating Privacy-Preserving Personalized Advice with Zero-Knowledge Proofs and LLMs. Companion Proceedings of the ACM Web Conference 2025. arXiv:2502.06425.

[9] Rihartanto Rihartanto; Supriadi Supriadi; Didi Susilo Budi Utomo "Image Tiling Using Columnar Transposition"IEEEXplore,11 April 2019.

[10] Ashwin Ramesh "Enhancing the Security of Hill Cipher using Columnar Transposition" IJERT Volume 4 Issue 7 July 2015.

[11] T. Rajani Devi "Importance of Cryptography in Network Security" IEEEXplore, 10 June 2013.

[12] Sharma, H., Aggarwal, G., & Kumar, S. (2022). Literature Survey for Design and Analysis of a Multi-Level Algorithmic Technique for Encryption and Decryption. Grenze International Journal of Engineering & Technology (GIJET), 8(2).

[13] Mishra, A., Chaturvedi, R. P., Sharma, H., Sharma, R., & Asthana, S. (2023, November). Multi-Scale Optimized Feature Network for Polyp Segmentation. In 2023 International Conference on Computing, Communication, and Intelligent Systems (ICCCIS) (pp. 444-448). IEEE.

[14] Yaga, D., Mell, P., Roby, N., & Scarfone, K. (2019). Blockchain technology overview. arXiv preprint arXiv:1906.11078.

[15] Kumar, R., Pandey, B., Sharma, H., Mishra, S., Chaudhary, N., Yadav, S., Gambhir, I.: Script Recognition of Printed Manuscript Images Using DCT and PCA. AIP Conf. Proc. 3297, 080001 (2025).