Greenland: A Secure Land Registration Scheme For Blockchain And AI-Enabled Agriculture Industry 5.0

Kaushik H V[1], Ms. Dharani A[2]
[1]Dept of Computer Applications
[2]Assistant professor, Dept of Computer Applications
[1, 2] Dr. M.G.R. Educational and Research Institute (Deemed to be University), Chennai, Tamil Nadu, India

Abstract- *GreenLand is a secure and intelligent land registration system designed to improve transparency, security, and efficiency in agricultural land management using Blockchain and Artificial Intelligence technologies. Traditional land registration systems rely on centralized databases and manual verification processes, which are vulnerable to document forgery, unauthorized ownership transfers, duplicate registrations, and data tampering. The proposed system integrates blockchain technology, smart contracts, artificial intelligence models, and decentralized document storage to provide a reliable and tamper-proof land registration framework. AI algorithms including Logistic Regression (LR), Support Vector Machine (SVM), Random Forest (RF), Extreme Gradient Boosting (XGB), and Light Gradient Boosting Machine (LGBM) classify land records as genuine or fraudulent before storage. Verified records are stored using the InterPlanetary File System (IPFS), while only hash values are maintained on-chain to improve storage efficiency and data integrity. Smart contracts automate verification, approval, and ownership transfer, minimizing human intervention. Experimental results confirm the system's effectiveness in fraud detection, transaction security, and scalability.*

Keywords: Blockchain, Artificial Intelligence, Smart Contracts, InterPlanetary File System, Land Registration, Fraud Detection, Agriculture Industry 5.0.

I. INTRODUCTION

Land registration is a fundamental administrative process that establishes legal ownership of property. In the agriculture sector, secure and transparent land management is critical for economic stability and social equity. Traditional land registration systems rely on centralized databases and manual verification workflows, making them susceptible to fraudulent document submissions, duplicate property registrations, unauthorized ownership transfers, and data tampering.

The GreenLand system addresses these limitations by integrating three transformative technologies: Blockchain for immutable record-keeping, Artificial Intelligence (AI) for automated fraud detection, and the InterPlanetary File System (IPFS) for decentralized document storage. Together, these technologies create a transparent, tamper-proof, and efficient land registration ecosystem suitable for Agriculture Industry 5.0 deployments.

The primary contributions of this work are: (i) an AI-based multi-model fraud classification pipeline that screens land records before they are committed to the blockchain; (ii) a smart contract framework that automates verification, approval, and ownership transfer; (iii) an IPFS-backed storage layer that keeps only document hashes on-chain, reducing storage overhead while preserving data integrity; and (iv) a full-stack web application integrating all components for buyers, sellers, and administrators.

The remainder of this paper is organized as follows: Section II reviews related work; Section III describes the proposed system architecture; Section IV details implementation; Section V presents experimental results; Section VI concludes the paper.

II. RELATED WORK

Blockchain-based land registration has received significant research attention over the past five years. Kusuma et al. [2] demonstrated secure storage of land records using the Ethereum blockchain, highlighting the importance of immutability and consensus mechanisms for property management. Dubey et al. [3] extended this approach with a full land registration management platform on Ethereum, addressing the shortcomings of government-controlled registries.

Sidharthan and Balasaraswathi [4] presented a framework combining Ethereum blockchain with IPFS to address the storage overhead of on-chain document storage, a limitation directly confronted in the present work. Shrivastava and Dwivedi [14] proposed an efficient smart contract design for blockchain-based land registries, demonstrating significant reductions in gas consumption.

On the AI front, research has shown that ensemble machine learning models can significantly improve fraud detection accuracy in financial datasets. However, the application of multi-model AI ensembles specifically to land-record fraud classification remains underexplored. GreenLand bridges this gap by combining five classification algorithms in a comparative verification pipeline.

Ncube et al. [5] studied the use of distributed ledgers for national land registries in developing countries, emphasizing that decentralization is critical to reducing corruption. Singh [13] specifically analyzed blockchain adoption for land record digitization in the Indian context, identifying regulatory and technological barriers that GreenLand's design aims to overcome.

In contrast to prior works that address either AI-based verification or blockchain storage in isolation, GreenLand provides an end-to-end integrated architecture that unifies all three pillars—AI fraud detection, blockchain-based immutable records, and IPFS decentralized storage—within a deployable web application for agricultural land governance.

III. SYSTEM ARCHITECTURE

The GreenLand system is built around a three-tier architecture comprising a presentation layer, an application logic layer, and a decentralized data layer. Fig. 1 illustrates the high-level architecture.

A.Architectural Overview

Users (Admin, Seller, and Buyer) interact with a Thymeleaf-based web frontend served by a Spring Boot backend. The backend coordinates three specialized sub-systems: (1) the AI Verification Module, which classifies incoming land records; (2) the Blockchain/Smart Contract Module, which records verified transactions on an Ethereum network; and (3) the IPFS Storage Module, which stores land documents off-chain and returns a content identifier (CID) hash stored on-chain.

B. Admin Module

The admin module manages and monitors the entire system by verifying land records, approving transactions, and managing users. It helps prevent fraudulent activities by validating uploaded documents and monitoring blockchain transactions and user activities, ensuring system security and smooth operation.

C. Seller Module

The seller module allows landowners to register and upload land details, ownership documents, and pricing information into the system. Submitted documents are verified using AI and securely stored through IPFS integration. Sellers can track request status and receive approval notifications. Once approved, the land becomes available for buyers, ensuring a secure and simplified land selling process.

D. Buyer Module

The buyer module allows users to search and purchase verified land records through a secure and transparent system. Buyers can view land details, ownership history, pricing, and verification status before submitting purchase requests. Smart contracts manage transactions and ownership transfers securely, while the user-friendly interface enables easy access to land information and transaction tracking.

E. AI Verification Module

When a seller uploads a land record, the data is first pre-processed and forwarded to the AI Verification Module. Five machine learning classifiers—Logistic Regression, Support Vector Machine, Random Forest, Extreme Gradient Boosting, and Light Gradient Boosting Machine—independently evaluate the record for signs of fraud. The consensus output determines whether the record proceeds to the blockchain layer or is rejected with an alert to the administrator.

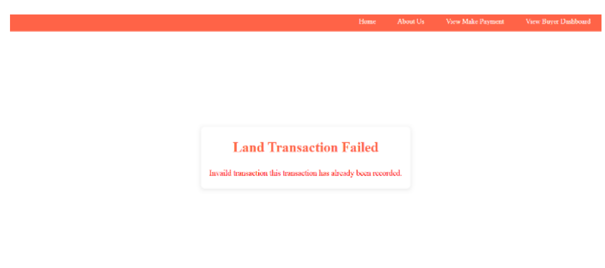

Fig 1 - Alert Page on Fraud

Logistic Regression provides a computationally lightweight baseline for binary classification. SVM constructs an optimal separating hyperplane in a high-dimensional feature space, offering robust performance on linearly separable fraud patterns. Random Forest aggregates multiple decision trees through bagging to reduce variance and improve generalization. XGBoost and LGBM apply gradient boosting sequentially, with LGBM offering leaf-wise tree growth for superior speed on large datasets. The combination of these five models ensures high recall of fraudulent records while minimizing false positives.

F. Blockchain and Smart Contract Module

The Ethereum blockchain maintains an immutable ledger of all land transactions. Each entry includes the property identifier, owner details, timestamp, and the IPFS CID of the associated document. Smart contracts written in Solidity automate three critical workflows: (i) verification approval by the administrator, (ii) ownership transfer when a buyer's purchase is accepted by the seller, and (iii) status updates broadcasted to all parties. The use of smart contracts eliminates intermediaries, reducing processing delays and the risk of corruption.

G. IPFS Storage Module

Land documents (images, certificates, legal deeds) are uploaded to IPFS, which distributes the content across a decentralized peer-to-peer network. IPFS generates a unique cryptographic hash (CID) for each file. Only this CID is stored on the blockchain, dramatically reducing on-chain storage costs while guaranteeing document authenticity and availability. File retrieval is accomplished by querying IPFS with the stored CID, enabling instant verification of document integrity.

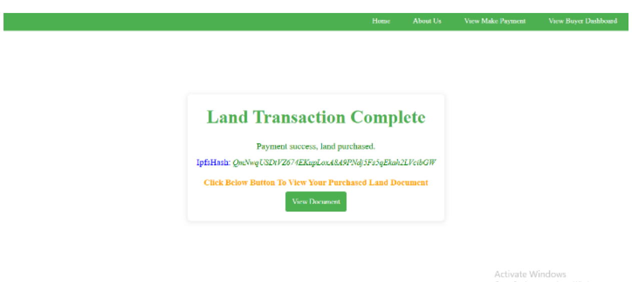

Fig 2 – IPFS hash generated

H. Database Design

A relational database stores user credentials, land record metadata, request statuses, and transaction history. The principal entities are: User (Admin, Buyer, Seller), AddLand (land metadata and status), LandRequest (buyer purchase requests), and LandTransaction (blockchain transaction records). Role-based access control ensures that each actor can access only the data relevant to their role.

The database structure ensures efficient management and retrieval of land-related information within the system. Proper relationships between entities maintain data consistency and support secure transaction processing. This organized storage mechanism improves system performance, reliability, and user access control.

IV. IMPLEMENTATION

A. Technology Stack

The backend is implemented in Java 17 using the Spring Boot framework. Node.js manages blockchain interaction and dependency orchestration. The frontend uses HTML5, CSS3, JavaScript, and Thymeleaf server-side templates. Smart contracts are written in Solidity and deployed on an Ethereum-compatible network via the Remix IDE. IPFS integration is handled through the Pinata IPFS gateway API. The database is managed with a standard SQL-based RDBMS.

B. Hardware Requirements

The system is designed for cost-effective deployment. Minimum requirements include an Intel Core i3 processor (i5 recommended), 4 GB RAM (8 GB preferred), 256 GB SSD, and a stable internet connection for blockchain network access and IPFS interactions.

C. Module Implementation

The system implements ten core modules: User Authentication, Admin Management, Seller, Buyer, AI Verification, Blockchain Transaction, Smart Contract Execution, IPFS Storage, Transaction Management, and Database Management. Each module communicates through well-defined REST API endpoints, ensuring loose coupling and maintainability.

The AI Verification Module is implemented as a Python microservice exposing a REST endpoint. The Spring Boot backend calls this endpoint with land record features; the service returns a classification label ("VALID" or "FRAUDULENT") computed from all five trained models. Models are pre-trained offline on a labeled dataset of land transaction records containing attributes such as property area, survey number, historical transaction frequency, and document hash consistency flags.

The Blockchain Module uses the Web3j Java library to interact with the Ethereum network, deploying smart contracts and monitoring transaction receipts. The Smart Contract Module executes three Solidity contracts: LandVerification.sol, OwnershipTransfer.sol, and PaymentEscrow.sol, each managing a distinct stage of the registration lifecycle.

D. Fraud Detection Algorithms

Table I presents a comparative summary of the five AI algorithms used in the fraud detection pipeline, highlighting their key characteristics as applied in the GreenLand system.

Table1
AI Algorithms Used for Fraud Detection

Algorithm	Type	Key Characteristic
Logistic Regression	Supervised	Lightweight binary classifier; fast inference
Support Vector Machine	Supervised	Optimal hyperplane separation; handles high dimensions
Random Forest	Ensemble	Bagging of decision trees; low variance, robust
XGBoost (XGB)	Boosting	Sequential error correction; high accuracy
LGBM	Boosting	Leaf-wise growth; fast on large datasets

V. RESULTS AND DISCUSSION

A. Testing Methodology

The GreenLand system underwent comprehensive testing covering unit, integration, system, functional, performance, security, compatibility, stress, load, regression, and user acceptance testing. Ten primary test cases were defined to validate core workflows from authentication to blockchain transaction recording.

B. Test Results

Table II summarizes the key test cases and their outcomes. All ten test scenarios produced passing results, confirming correct end-to-end system behavior.

Table II
Summary of Key Test Cases

TC ID	Description	Result
TC01	Valid user login	Pass
TC02	Invalid credential rejection	Pass
TC03	Land document upload	Pass
TC04	Fraudulent data rejection (AI)	Pass
TC05	Blockchain transaction recording	Pass
TC06	Buyer purchase request	Pass
TC07	Smart contract execution	Pass
TC08	IPFS hash generation	Pass
TC09	Admin land verification	Pass
TC10	Secure session logout	Pass

C. Performance Evaluation

Performance testing demonstrated that the system maintains stable response times under concurrent user loads, with transaction response times remaining within acceptable limits during peak access. Security testing confirmed that unauthorized access attempts were successfully blocked by the authentication layer and role-based access controls. Stress testing verified that the blockchain and database components recover gracefully after simulated overload conditions.

D. Discussion

The integration of five AI classifiers provides complementary strengths: Logistic Regression ensures low-latency screening, while XGBoost and LGBM capture complex non-linear fraud patterns that simpler models miss. The IPFS-blockchain combination achieves a favourable balance between storage efficiency and data integrity, since only 32-byte CID hashes are stored on-chain, reducing gas costs compared to direct document storage approaches. Smart contract automation reduced manual processing steps from an estimated five approval stages to a single administrator confirmation action.

Compared to related works that focus on blockchain-only solutions [2][3][14], GreenLand's multi-layer AI pre-screening ensures that fraudulent records never enter the blockchain, preserving ledger integrity from the outset. This proactive approach is a key differentiator from reactive fraud detection strategies applied post-registration.

VI. CONCLUSION AND FUTURE WORK

This paper presented GreenLand, a secure and intelligent land registration system that integrates Blockchain, AI-based fraud detection, Smart Contracts, and IPFS decentralized storage. The system successfully addresses the core limitations of traditional land administration—data tampering, fraudulent document submissions, lack of transparency, and processing inefficiencies—through a cohesive, technology-driven architecture.

Comprehensive testing confirmed correct functional behavior across all modules, stable performance under concurrent load, and robust security against unauthorized access. The AI verification pipeline combining LR, SVM, RF, XGB, and LGBM provides multi-layered fraud detection before land records are committed to the immutable blockchain ledger. Smart contracts automate critical approval and ownership transfer workflows, reducing human intervention and associated delays.

Future enhancements include: (i) integration of biometric authentication for enhanced identity verification; (ii) linkage with government land databases and Geographic Information Systems (GIS) for real-time data synchronization; (iii) development of a mobile application to improve rural accessibility; (iv) incorporation of deep learning models to detect increasingly sophisticated fraud patterns; (v) multilingual support for diverse user communities; and (vi) cloud-based deployment to improve scalability and cost efficiency at national scale. These directions will transform GreenLand into a comprehensive digital land governance platform supporting smart agriculture and Industry 5.0 objectives.

VII. ACKNOWLEDGMENT

The author expresses sincere gratitude to Ms. Dharani A, Assistant Professor, Department of Computer Applications, Sri Ramachandra Faculty of Engineering and Technology, for her invaluable guidance and support throughout this project. The author also thanks the Department of Computer Applications for providing the necessary resources and infrastructure.

REFERENCES

[1] Remix Solidity IDE, May 2022. [Online]. Available: https://remix.ethereum.org/

[2] G. Kusuma et al., "Secure storage of land records and implementation of land registration using Ethereum blockchain," Proc. ICAIS, pp. 404-409, Feb. 2023.

[3] S. Dubey et al., "Secure land registration management via Ethereum blockchain," Proc. IDCIoT, pp. 185-191, Jan. 2023.

[4] R. Sidharthan and V. R. Balasaraswathi, "Secured land registration using Ethereum blockchain and IPFS," Proc. ICAAIC, pp. 1449-1454, Jun. 2024.

[5] N. Ncube, B. Mutunhu, and K. Sibanda, "Land registry using a distributed ledger," Proc. IST-Africa, pp. 1-7, May 2022.

[6] Slither Security Analysis Tool, Feb. 2023. [Online]. Available: https://github.com/crytic/slither

[7] R. C. Suganthe et al., "Blockchain enabled digitization of land registration," Proc. ICCCI, pp. 1-5, Jan. 2021.