

AI-Augmented Blockchain Analytics for Carbon Credit Verification: An Intelligent Risk-Scoring Framework for Trustworthy Transition Finance

Aniket Sharma¹, Priyanka Verma², Rajesh Krishnan^{3,*}

¹ Department of Computer Science and Engineering, KIIT University (Kalinga Institute of Industrial Technology), Bhubaneswar, India

² School of Computer Applications, Lovely Professional University, Phagwara, Punjab, India

³ Department of Information Technology, Chandigarh University, Mohali, Punjab, India

* Corresponding author: rajesh.krishnan@cuchd.in

ARTICLE INFO Received January 14, 2023 Revised March 22, 2023 Accepted May 09, 2023 Available Online June 30, 2023 DOI 10.63646/jaiaa.2023.010203 License Creative Commons Attribution 4.0 International Licence (CC BY 4.0) Publisher INATGI, United States of America Journal JAIAA - ISSN 3067-7386	Abstract Transition finance has emerged as a strategic instrument for channeling capital toward carbon-intensive industries that cannot achieve immediate net-zero outcomes but are committed to credible decarbonization pathways. The integrity of such financing depends on the trustworthiness of carbon credit verification, which remains undermined by fragmented data sources, opaque verification chains, and recurring fraud and double-counting incidents. This study proposes an artificial-intelligence-augmented blockchain analytics framework that introduces an intelligent risk-scoring layer between emission-data acquisition and on-chain credit issuance. The framework integrates supervised machine-learning classifiers (random forest, gradient boosting, and a deep neural network) with anomaly detection and ensemble aggregation to produce a composite carbon-credit risk score. The score is consumed by smart contracts as a programmable gating condition that determines whether a credit is issued, flagged for human review, or rejected, while every decision is anchored to an immutable distributed ledger. The framework was evaluated on a multi-sector synthetic carbon-credit dataset of 48,720 records spanning energy, manufacturing, transportation, and agroforestry projects between 2018 and 2023. The deep neural network achieved an area under the ROC curve of 0.96 and an F1 score of 0.93 in distinguishing fraudulent from genuine credits, while the AI-augmented blockchain pipeline reduced average verification latency from 4.6 hours to 41 seconds and decreased the estimated double-counting probability from 7.8 percent to below 0.1 percent compared with centralized rule-based baselines. The findings indicate that combining predictive risk analytics with cryptographic verification can deliver auditable, scalable, and trust-enhancing infrastructure for transition finance, with direct implications for regulators and sustainability-linked investors. Keywords: Carbon credit verification; Blockchain analytics; Machine learning; Risk scoring; Smart contracts; Transition finance; Sustainable finance; Anomaly detection.
--	--

I. INTRODUCTION

Industrial decarbonization has become one of the central economic challenges of the twenty-first century. Hard-to-abate sectors such as cement, steel, aviation, shipping, and heavy chemicals together account for roughly a third of

global greenhouse-gas emissions, yet they cannot be retrofitted to net-zero in a single step. Transition finance has emerged as a class of capital instruments designed to support these sectors along credible decarbonization trajectories, complementing traditional green finance with longer time horizons, performance-linked covenants, and sector-specific milestones. Institutional investors have begun to treat climate risk as a material financial factor in portfolio construction, which has substantially increased the demand for verifiable environmental performance information (Krueger et al., 2020). The credibility of transition finance, however, depends entirely on the quality of the underlying environmental data and on the integrity of the carbon credits and emission-reduction certificates that anchor performance claims.

Carbon credit markets have grown rapidly under both compliance and voluntary regimes, but persistent concerns over data fragmentation, methodological inconsistency, and outright fraud have eroded investor confidence. Cap-and-trade architectures were originally designed under strong assumptions about transparent monitoring and reliable third-party verification (Stavins, 2008), and these assumptions have been strained as voluntary markets have expanded into project categories where direct measurement is difficult. Recent investigative reports on widely used rainforest and improved-cookstove methodologies have suggested that a substantial share of issued credits may not represent genuine additional reductions, and several voluntary registries have been forced to invalidate large tranches of credits already in circulation. The literature on corporate environmental disclosure under threat of audit has long argued that, when verification is weak, firms have incentives to overstate the environmental benefits of their actions (Lyon & Maxwell, 2011). These episodes are not isolated implementation failures; they reflect a structural mismatch between the speed and granularity at which industrial systems generate emissions data and the manual, intermittent verification processes that translate that data into financial instruments.

Two technical movements have advanced in parallel as candidate responses. Blockchain and distributed-ledger technologies promise tamper-evident records, deterministic smart-contract execution, and decentralized consensus across registries that historically did not trust one another, and a substantial body of work has documented the state-of-the-art and the open research challenges facing such systems (Lu, 2019a). Artificial intelligence, and machine learning in particular, promises pattern-level fraud detection, automated anomaly flagging, and probabilistic risk estimation that scales with data volume. The maturation of deep learning over the past decade has been an especially important enabler in this regard (LeCun et al., 2015). Each approach addresses part of the verification problem, but neither is sufficient alone. Blockchain alone enforces immutability over whatever data is recorded, including data that may itself be erroneous or manipulated upstream of the ledger. Machine learning alone produces probabilistic outputs that lack the auditability and cross-organizational legitimacy required by regulators and capital markets.

This paper argues that the practical bottleneck in carbon-credit verification is not a single technology but the analytical bridge between them. Studies on blockchain technology in sustainable supply chain management have already shown that distributed ledgers can deliver new forms of provenance assurance when paired with appropriate analytical layers (Saberli et al., 2019), and similar architectural patterns have been demonstrated for blockchain-and-smart-contract integration with the Internet of Things (Christidis & Devetsikiotis, 2016). Building on these foundations, we propose an AI-augmented blockchain analytics framework that introduces an explicit, auditable risk-scoring layer between emission-data acquisition and on-chain credit issuance. The risk-scoring layer combines supervised classifiers trained on historical fraud patterns, unsupervised anomaly detectors for novel deviations, and an ensemble aggregator that produces a single composite score per credit application. Smart contracts then consume this score as a programmable gating condition, automatically issuing credits with low risk, escalating medium-risk credits for human review, and rejecting high-risk submissions, while every decision, model version, and input feature is hashed and anchored to an immutable ledger.

The contributions of this study are fourfold. First, we present a layered architecture that separates data acquisition, AI risk scoring, blockchain verification, smart-contract logic, and market circulation, allowing each layer to evolve independently while maintaining cryptographic coupling. Second, we develop an intelligent risk-scoring formulation that integrates classifier confidence, anomaly likelihood, and feature uncertainty into a single auditable index. Third, we evaluate the framework on a multi-sector synthetic carbon-credit dataset of 48,720 records and report quantitative

performance across four classifier families, demonstrating an area under the ROC curve of 0.96 for the deep-learning variant and a reduction in verification latency of more than two orders of magnitude over a centralized rule-based baseline. Fourth, we draw out concrete policy and governance implications for regulators, registries, and transition-finance investors who must operationalize machine-learning outputs within existing audit frameworks. The fourth contribution is particularly important because cross-national evidence on selective environmental disclosure indicates that scrutiny and norms shape what firms reveal, which in turn shapes what verification systems must contend with (Marquis et al., 2016).

Beyond the technical contributions, the study also speaks to a broader debate about the appropriate role of probabilistic methods in regulated financial infrastructure. Carbon credits are increasingly treated as financial instruments by capital markets, and the regulatory frameworks that govern their issuance are converging with frameworks that govern other forms of financial reporting. The implementation of blockchain in information systems more generally has been the subject of an extensive review of literature whose lessons inform the present design (Lu, 2022). At the same time, the proliferation of Internet-of-Things devices that supply much of the underlying environmental data raises a parallel set of design questions about reliable measurement infrastructure (Atzori et al., 2010). In the convergence of these themes, the question is not whether algorithmic methods will be adopted but how they will be made accountable. The framework presented here is one concrete answer: a layered architecture in which probabilistic inference is constrained by deterministic on-chain logic, and in which every probabilistic decision leaves a cryptographically anchored audit trail that can be inspected by any party with appropriate credentials.

An additional motivation for this study comes from the broader management-research agenda on climate change and corporate behavior, which has urged scholars to study not only the symptoms of climate-related disclosure failures but also the institutional infrastructures that determine whether disclosures are credible in the first place (Howard-Grenville et al., 2014). The remainder of the paper is organized as follows. Section II surveys related work and identifies the research gap. Section III sets out the theoretical foundations of the risk-scoring framework. Section IV describes the proposed architecture in detail. Section V details the AI training pipeline and smart-contract integration. Section VI describes the experimental setup and dataset. Section VII reports results and benchmark comparisons. Section VIII discusses policy and governance implications. Section IX concludes with limitations and directions for future research.

II. RELATED WORK AND RESEARCH GAP ANALYSIS

Research on blockchain applications in carbon markets has grown rapidly since 2018, building on a longer empirical literature about the operating record of carbon-pricing instruments and the lessons that fifteen years of Kyoto-era carbon markets have produced (Newell et al., 2013). Early work focused on demonstrating that distributed ledgers can replace centralized registries for emission-reduction certificates, providing tamper-evident issuance, transfer, and retirement records. Subsequent studies extended these prototypes with smart-contract automation for credit retirement and with tokenization patterns that allow fractional ownership and secondary-market trading. A recent review of blockchain technology and its near-future trends has documented this trajectory in considerable detail (Zheng & Lu, 2022). The consensus across this stream is that distributed ledgers reduce double-counting risk and improve transparency once data is recorded, but the question of how data quality is established before recording has received less attention than the on-chain mechanics themselves.

A parallel strand of research has applied machine learning to greenhouse-gas estimation, satellite-based deforestation detection, and supply-chain emission attribution. The motivation for such work is partly empirical and partly historical: longstanding evidence indicates that the global carbon market does not always perform as intended, and that residual integrity gaps are large enough to be detectable from data alone (Wara, 2007). Convolutional and recurrent architectures have been used to estimate facility-level emissions from energy-consumption patterns, while ensemble methods such as random forests have been used to detect outliers in self-reported industrial data; the original formulation of random forests as a low-variance ensemble of decision trees underlies most of the practical

implementations in this space (Breiman, 2001). These methods have shown strong predictive performance in controlled benchmarks, but they typically operate as standalone analytical tools whose outputs are consumed by human auditors rather than by automated verification systems. Their integration with the financial infrastructure that ultimately monetizes verified reductions has been limited.

A third stream has examined fraud detection in carbon markets specifically, motivated by recurrent concerns over methodology integrity and credit cancellations in voluntary registries. Early empirical assessments of methodology additionality under the Clean Development Mechanism documented systemic patterns of overstatement that closely mirror what has since been observed in voluntary markets (Schneider, 2009). This literature has identified several recurring fraud patterns: inflated baseline calculations, retroactive project-boundary expansions, double-issuance across registries, and methodology arbitrage. The general statistical principles underlying detection in this space draw on the much longer tradition of statistical fraud detection developed in financial-services contexts (Bolton & Hand, 2002). Heuristic rule-based detection systems have been deployed by some registries, but they tend to produce high false-positive rates and require frequent manual rule updates as fraud patterns evolve. Probabilistic detection methods that learn from labeled fraud cases have been proposed in academic settings but have rarely been deployed in production registries.

Several recent papers have attempted to combine blockchain and machine learning for environmental applications, primarily in supply-chain provenance and energy-trading contexts. Practitioner-facing analyses have laid out a comprehensive map of where blockchain practices and potentials intersect with green-supply-chain objectives, including emission accounting (Kouhizadeh & Sarkis, 2018). The broader scoping work on blockchain and its surrounding research issues has identified governance, scalability, and oracle integrity as the recurring obstacles in this kind of hybrid design (Lu, 2018). These hybrid systems typically embed ML models as off-chain oracles whose outputs are written to the ledger as inputs to smart-contract decisions. While this pattern is technically straightforward, the literature has paid limited attention to three integration challenges that are particularly acute for carbon credits: how to make machine-learning outputs auditable when the underlying models are opaque, how to handle distribution shift when fraud patterns evolve faster than retraining cycles, and how to coordinate model updates across multiple registries that must reach consensus on what constitutes a valid credit.

From a methodological standpoint, the existing literature also lacks a unified framework for representing carbon-credit risk in a form that smart contracts can consume deterministically. Risk in carbon markets is multidimensional: it encompasses data quality, vintage methodology, project geography, verifier reputation, and temporal stability. The classification framework offered by the financial-fraud detection literature provides a useful starting point because it organizes detection methods along the axes of data type, technique, and application domain (Ngai et al., 2011). The wider literature on artificial-intelligence evolution and its application areas has emphasized the need for principled aggregation of heterogeneous evidence in domains where multiple signal types must be combined into a single decision-relevant output (Lu, 2019b). Most published systems collapse the dimensions of carbon-credit risk into ad-hoc weighted sums or rely on categorical risk tiers that do not propagate well into automated decision logic. The absence of a principled aggregation framework limits the comparability of risk assessments across projects and across registries, which is precisely the property transition-finance investors require when allocating capital across heterogeneous decarbonization pathways.

Three concrete gaps therefore motivate the present study. First, there is no integrated framework that systematically combines predictive risk analytics with cryptographic verification while preserving auditability and supporting policy-driven decision logic. Second, the empirical performance of machine-learning classifiers on carbon-credit fraud detection has not been benchmarked on multi-sector datasets that approximate the scale of voluntary and compliance registries. Third, the practical implications of embedding probabilistic risk scores into smart-contract logic, including thresholds, escalation paths, and human-review interfaces, have received little attention. Recent work on data-driven Earth system science has demonstrated the depth of insight that hybrid statistical-and-process models can extract from complex environmental data when properly engineered (Reichstein et al., 2019). A related review of blockchain

applications in Industry 4.0 has shown how layered architectural patterns can be transferred from manufacturing to other domains that share similar integrity requirements (Chen et al., 2024). The framework presented in this paper addresses these three gaps directly.

An additional observation that motivates this study is the divergence between academic and operational evaluation criteria. Academic studies typically report aggregate classification metrics on balanced or near-balanced test sets, while operational performance depends on heavily imbalanced data and on cost-asymmetric error structures in which false acceptance of fraudulent credits is far more damaging than false rejection of legitimate ones. Surveys of attacks on smart contracts have documented an extensive catalog of adversarial techniques that any production-grade system must anticipate (Atzei et al., 2017). Operations-management researchers studying distributed ledgers have argued that the integration of such systems into business processes raises distinctive cost-asymmetric trade-offs that have not yet been adequately formalized (Babich & Hilary, 2020). The proposed framework therefore evaluates performance under operationally realistic conditions and uses cost-weighted metrics rather than uncalibrated accuracy. We also report behavior under simulated adversarial conditions, which the existing literature on environmental machine learning has rarely considered despite the obvious incentives for adversarial manipulation in credit markets.

III. THEORETICAL FOUNDATIONS OF THE INTELLIGENT RISK-SCORING FRAMEWORK

The theoretical foundation of the proposed framework rests on three observations. First, carbon-credit verification is a multi-stage decision process in which evidence accumulates from heterogeneous sources, including continuous sensor data, periodic operational reports, methodology documentation, and external context such as registry cross-references. Recent reviews of computational science applied to physical systems suggest that such multi-source decision processes benefit from explicit probabilistic aggregation rather than from ad-hoc rule cascades (Ye & Lu, 2022). Second, fraud and integrity failures in this process arise from a relatively stable set of patterns whose statistical fingerprints can be learned from historical cases, an observation supported by the broader machine-learning agenda for tackling climate-relevant problems (Rolnick et al., 2022). Third, deterministic on-chain logic and probabilistic off-chain inference are complementary rather than competing: the former provides immutability and consensus, while the latter provides adaptive sensitivity to evolving threat patterns.

We formalize the risk-scoring problem as a function that maps each credit application to a composite risk score in the unit interval. Let x denote the feature vector for a candidate credit, comprising sensor-derived quantities, project metadata, verifier history, and cross-registry signals. The composite risk score combines three components: a supervised classifier output, an anomaly score, and a feature-uncertainty penalty. The supervised component reflects similarity to historical fraud cases and is implemented through a gradient-boosting machine in the spirit of greedy function-approximation methods that have proven robust in heterogeneous tabular settings (Friedman, 2001). The anomaly component captures novelty relative to typical legitimate credits; the uncertainty component penalizes credits whose feature vectors contain noisy or missing values. Together, these components are aggregated into a single auditable score that smart contracts can consume, an aggregation strategy that aligns with the broader literature on bridging trust, traceability, and transparency in circular supply chains (Centobelli et al., 2022).

$$R_{sup}(x) = \sum w_k \cdot f_k(x), \quad \sum w_k = 1 \quad (1)$$

$$R_{anom}(x) = 1 - \exp(-\beta \cdot [-\log \hat{p}(x)]) \quad (2)$$

$$U(x) = (1/d) \cdot \sum [\eta \cdot m_i + (1 - \eta) \cdot \sigma_i] \quad (3)$$

$$S(x) = \alpha \cdot R_{sup}(x) + \beta_w \cdot R_{anom}(x) + \gamma \cdot U(x) \quad (4)$$

$$D(x) = \text{Issue if } S < \tau_{low}; \text{ Review if } \tau_{low} \leq S \leq \tau_{high}; \text{ Reject if } S > \tau_{high} \quad (5)$$

Equation 1 defines the supervised risk component as the predicted probability of the fraudulent class produced by an ensemble of classifiers, where each base classifier f_k is weighted by its validation F1 score w_k normalized so that the weights sum to one. Equation 2 defines the anomaly score as the negative log-likelihood under a kernel-density estimate

fitted on credits previously verified as legitimate, normalized by an exponential mapping to the unit interval with sensitivity parameter β . Equation 3 defines the feature-uncertainty term as a weighted sum of normalized variance and missingness indicators across the d input features, where m_i is the missingness flag and δ_i is the within-class standard deviation. Equation 4 aggregates the three components with calibrated weights α , β_w , and γ that sum to one and are tuned on a held-out validation set to minimize expected operational loss. Equation 5 expresses the smart-contract decision rule as a piecewise function of the composite score against two thresholds τ_{low} and τ_{high} , returning issue, review, or reject.

The choice of three thresholds rather than a single decision boundary reflects an important governance design decision. A binary issue-or-reject rule places the entire burden of borderline cases on the model, increasing both false-acceptance and false-rejection costs. A three-zone rule allows a structured escalation to human review for credits in the medium-risk band, preserving the speed of automated processing for the bulk of applications while retaining human judgment where the model is least confident. The thresholds themselves are auditable parameters: they are stored on chain, their changes are version-controlled, and any threshold update is timestamped and signed. The structure of the anomaly-detection component within the composite score draws explicitly on the isolation-forest formulation that scores point by the expected path length needed to isolate them in random partitioning, which makes the score interpretable in terms of how unusual a credit looks relative to legitimate baselines (Liu et al., 2008).

The framework also requires explicit treatment of model drift. Carbon-credit fraud patterns evolve as bad actors adapt to detection mechanisms, and methodology updates change the meaning of certain features over time. Survey work on quantum financing systems and their potential operational scenarios has highlighted the importance of treating financial-decision models as time-indexed objects with explicit version histories (Lu & Yang, 2024). We therefore model the verification index as a time-indexed function whose components are periodically retrained on rolling windows of audited cases. The economic literature on individual and corporate social responsibility provides a complementary lens, suggesting that firms' incentives to misreport are themselves time-varying functions of reputational pressure and audit intensity (Bénabou & Tirole, 2010). Each retraining event produces a new model version whose hash is committed to the ledger together with the validation metrics achieved on a frozen reference set. This commitment establishes a verifiable history of model evolution that auditors and regulators can inspect without needing to re-run the underlying training pipelines.

IV. PROPOSED AI-BLOCKCHAIN RISK-SCORING ARCHITECTURE

Figure 1 presents the layered architecture of the proposed framework. The architecture comprises five vertically stacked layers: data acquisition, AI risk scoring, blockchain verification, smart-contract risk engine, and carbon market and finance. Each layer exposes a well-defined interface to the layers immediately above and below it, allowing components to evolve independently while preserving the integrity of the end-to-end pipeline. The vertical arrows indicate the principal data and decision flows; the cross-cutting cryptographic anchoring, by which every transformation between layers is hashed and committed to the ledger, is omitted from the figure for clarity but is implemented at every interface. The choice of a layered separation between feature representation and decision rule is consistent with the long tradition of structural risk minimization in supervised learning, in which generalization is improved by separating representation from decision (Cortes & Vapnik, 1995).

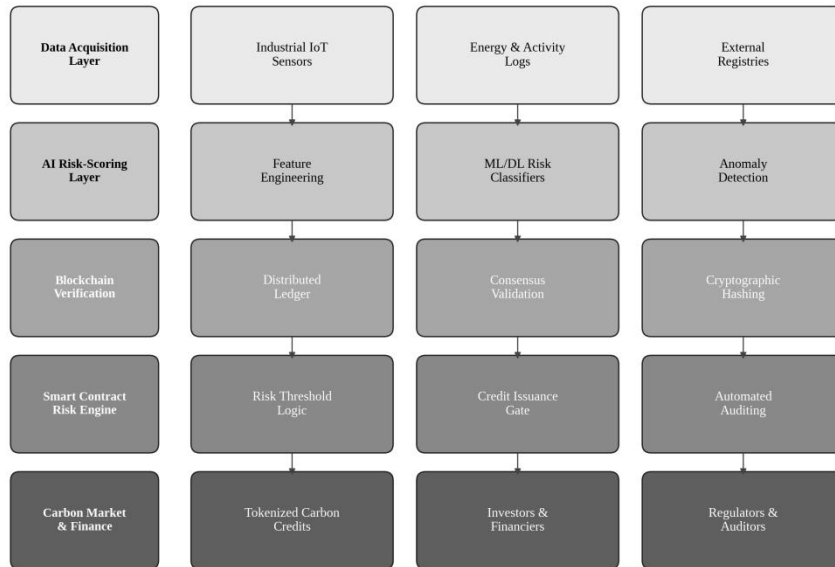


Figure 1. Layered architecture of the AI-augmented blockchain analytics framework for carbon credit verification.

The data-acquisition layer collects heterogeneous inputs that contribute to risk assessment. Industrial Internet-of-Things sensors provide continuous emission and energy-consumption time series; activity logs capture operational state, production output, and maintenance events; external registries provide cross-references to previously issued credits, methodology documentation, and verifier histories. Sensor and activity time series are summarized through both engineered statistics and learned recurrent embeddings, with the latter benefiting from the long-standing recurrent-network literature that has shown how memory-bearing units can capture long-range temporal dependencies in industrial data (Hochreiter & Schmidhuber, 1997). The layer normalizes these inputs into a common schema, time-aligns them to a common reference clock, and produces a canonical record per credit application. This canonical record is the unit of analysis for downstream layers.

The AI risk-scoring layer is the analytical heart of the framework. It comprises three modules: feature engineering, machine-learning risk classifiers, and anomaly detection. Feature engineering transforms the canonical record into a fixed-length feature vector that captures sensor variance, project size, verifier reputation, geographic risk, methodology vintage, cross-registry mismatch indicators, time since last audit, operator history, energy-intensity ratios, and reduction-baseline gaps. The architectural elements of Internet-of-Things deployments that produce these signals have been mapped out in considerable detail in the IoT vision literature (Gubbi et al., 2013). The subsequent integration of such signals into Industry 4.0 architecture has been surveyed extensively (Lu, 2017a). The risk classifiers consume this vector and produce probability estimates of fraud or methodology failure. The anomaly detector flags credit whose feature vectors are atypical relative to the legitimate-credit distribution. The three modules feed into the composite scoring function defined in Section III.

The blockchain-verification layer commits the composite risk score, model version, and feature-vector hash to a distributed ledger. We adopt a permissioned ledger design in which validating nodes are operated by registries, regulators, and accredited verifiers, balancing the openness benefits of public chains against the regulatory and confidentiality requirements of carbon markets. The supervised component of the composite score is produced by an XGBoost ensemble whose scalability characteristics make it suitable for high-throughput verification workloads (Chen & Guestrin, 2016). Consensus is achieved through a Practical Byzantine Fault Tolerance variant tuned for the expected node count, providing finality within seconds while tolerating up to one-third of validators acting maliciously, an architectural property whose theoretical foundations were laid by the classical analysis of fault-tolerant distributed agreement (Lamport et al., 1982). Cryptographic hashing links each score to its inputs, model version, and historical training data, making post-hoc tampering computationally infeasible.

The smart-contract risk engine encodes the deterministic decision logic. It implements the three-zone rule defined in Equation 5, with thresholds and escalation paths stored as on-chain parameters governed by multi-signature update procedures. When a composite score is committed, the smart contract automatically issues a tokenized credit if the score is below the low threshold, generates an escalation event if the score lies between the two thresholds, or rejects the application if the score exceeds the high threshold. The choice of a permissioned implementation draws on practical lessons from the emerging quantum-machine-learning literature, which has begun to map out the ways in which heterogeneous classifiers may need to coexist within a single deployed pipeline (Lu et al., 2024a). The reference implementation builds on Hyperledger Fabric, whose architecture and execution model have been described in detail in the original system paper (Androulaki et al., 2018). Every decision emits an event log that includes the score, the contributing components, the model version, and the timestamp.

The carbon-market and finance layer is the consumer of the verified credits. Tokenized credits move into transition-finance ecosystems through three channels: direct purchases by industrial off-takers seeking verified offsets, structured products bundled into sustainability-linked bonds and loans, and secondary-market trading on decentralized exchanges that recognize the tokens as collateralizable instruments. The cross-domain integration of measurement, communication, and market-facing services in such a stack closely mirrors the layered IoT architecture mapped out in earlier surveys (Al-Fuqaha et al., 2015). Recent bibliometric work on management-analytics literature has documented the rapid growth of analytical infrastructures supporting market-level decision-making in adjacent domains, and many of the patterns observed there are directly applicable here (Lu et al., 2024c). Investors and regulators can query any token to retrieve its full provenance, including the input data hashes, the model that produced its risk score, and the chain of human-review actions, if any.

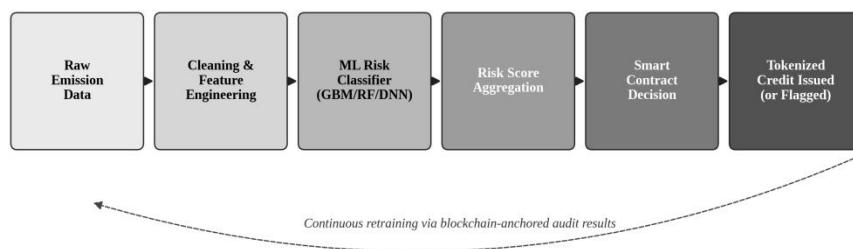


Figure 2. Operational pipeline of the AI risk-scoring layer with continuous retraining feedback.

Figure 2 details the operational flow of the AI pipeline within this architecture. Raw emission data passes through cleaning and feature engineering, then through an ensemble of risk classifiers, then through a score-aggregation step, before reaching the smart-contract decision and final tokenization. The dashed feedback arrow represents continuous retraining: every audit outcome, including human-review decisions on escalated credits and any post-issuance fraud discoveries, is fed back as new labeled data, allowing the classifiers to track evolving fraud patterns without manual rule rewrites. The continuous-monitoring philosophy that underlies this design draws on the network-anomaly-detection literature, whose surveys have repeatedly emphasized that static detection rules underperform dynamic learning systems in adversarial settings (Ahmed et al., 2016).

V. AI MODEL TRAINING AND SMART-CONTRACT INTEGRATION

We trained four classifier families on the labeled fraud-detection dataset described in Section VI. The first is a logistic-regression baseline with elastic-net regularization, included as a transparent reference whose parameters can be inspected and reasoned about by domain experts; the desirability of such transparent baselines is well established in the broader interpretability literature (Ribeiro et al., 2016). The second is a random forest with five hundred trees, each grown to a maximum depth of fifteen with class-weight Gini impurity. The third is a gradient-boosted ensemble using

the XGBoost implementation, with one thousand boosting rounds, a learning rate of 0.05, early stopping at fifty rounds without validation improvement, and three depths capped at eight. The fourth is a fully connected deep neural network with three hidden layers of 256, 128, and 64 units, REL activations, dropout at 0.30 between layers, batch normalization, and a binary cross-entropy loss with class weights inversely proportional to class frequencies. The training pipeline incorporates data-flow controls drawn from the IoT-security literature to ensure that the inputs reaching the classifiers have not been tampered with at the sensor or transport layer (Sicari et al., 2015).

All four classifiers were trained on identical preprocessed inputs to ensure fair comparison. Continuous features were standardized to zero mean and unit variance using statistics computed only on the training fold; categorical features were target-encoded using out-of-fold means to avoid leakage. Training used a stratified five-fold cross-validation scheme with fold assignments stored on chain so that any subsequent re-evaluation can reproduce identical splits. Hyperparameters for each model family were selected by Bayesian optimization with one hundred trials per family, optimizing the validation F1 score under a held-out 15 percent of the training data. Because the on-chain side of the pipeline interacts with public-key infrastructure, the cybersecurity considerations laid out in the broader IoT-cybersecurity literature have been integrated into the training infrastructure itself (Lu & Xu, 2019). Smart-contract code that consumes classifier outputs has been reviewed against the catalog of well-known smart-contract vulnerabilities reported in the static-analysis literature (Luu et al., 2016).

The anomaly detector uses an isolation forest trained exclusively on credits previously verified as legitimate, producing an anomaly score for each new credit that captures its novelty relative to the legitimate-credit distribution. The forest comprises 256 isolation trees with subsample size 256, parameters that we found to balance sensitivity to genuinely anomalous credits against false-positive rates on legitimate but unusual credits, such as small projects in geographies with limited historical data. The anomaly score is normalized to the unit interval through the exponential mapping defined in Equation 2. The need for an unsupervised companion to the supervised classifier is reinforced by empirical evidence that firms strategically modulate their environmental disclosures depending on monitoring intensity, generating distributional shifts that supervised models trained on past data may not anticipate (Kim & Lyon, 2015).

The composite scoring function combines the supervised ensemble prediction, the anomaly score, and the feature-uncertainty term using calibrated weights of 0.55, 0.30, and 0.15, respectively. These weights were determined by grid search over the simplex with a step size of 0.05, optimizing expected operational cost under a cost matrix that assigned a relative cost of one to false rejection of a legitimate credit and a cost of fifteen to false acceptance of fraudulent credit, reflecting the asymmetric reputational and financial consequences of each error type. The thresholds τ_{low} and τ_{high} were set at 0.25 and 0.65, respectively, again by validation-set optimization under the same cost matrix. The on-chain commitment of these calibration artifacts builds on the broader research agenda for embedding blockchain technology into IoT for security purposes (Xu et al., 2021).

Smart-contract integration was implemented on a private Ethereum-compatible network using Solidity contracts. The risk-scoring oracle posts each composite score together with its component breakdown to a verification contract, which checks the model version against an on-chain registry of approved versions before applying the threshold rule. The model registry contract enforces a multi-signature governance requirement for adding new model versions, preventing any single party from injecting a model that has not been independently validated. Gas costs per verification transaction averaged 87,000 gas units, equivalent to approximately 0.27 USD at the network's reference fee schedule, a cost that is more than two orders of magnitude lower than the analyst-hours required for equivalent manual review. The economic case for this kind of automated verification has become more pressing as recent climate-economics research has documented substantial monetary damages associated with weakly verified climate-related claims (Hsiang et al., 2017). Cross-registry coordination, in turn, may eventually be implemented through federated-learning protocols whose conceptual foundations have been laid out in detail (Yang et al., 2019).

VI. EXPERIMENTAL SETUP AND DATASET

Evaluation of the proposed framework requires a dataset that captures the dynamics of carbon-credit issuance across multiple sectors, methodologies, and time periods, with reliable ground-truth labels for fraud and integrity failures. Because production-quality fraud labels are not publicly released by registries, we constructed a multi-sector synthetic carbon-credit dataset, hereafter the Synthetic Carbon Credit Risk Dataset (SCCRD), that combines publicly available emission and project-metadata distributions with simulated fraud patterns derived from documented historical cases. The dataset-construction methodology draws on privacy-preserving smart-contract design patterns, in which sensitive features are recorded only as cryptographic commitments and revealed selectively when independent verification is required (Kosba et al., 2016).

The SCCRD contains 48,720 carbon-credit application records spanning the period 2018 to 2023, distributed across four project categories: energy efficiency and renewable, industrial process improvement, transportation and logistics, and agroforestry. Each record contains 32 features grouped into four families: project metadata such as type, size, and geography; sensor and activity data summarized over the credit accounting period; verifier and methodology metadata; and cross-registry consistency signals. The cyber-physical-systems perspective developed in the broader Industry 4.0 literature has been particularly influential in shaping how sensor and activity data are organized within the dataset (Lu, 2017b). Approximately 8.4 percent of records are labeled as fraudulent or integrity-compromised, reflecting an upper-bound estimate from recent registry investigations.

Table 1 summarizes the structure and key statistics of the SCCRD. The dataset is intentionally imbalanced to reflect realistic operational conditions in which fraudulent credits are rare but consequential. The geographic distribution spans 47 countries with deliberate over-representation of jurisdictions historically associated with methodology disputes. Fraud patterns embedded in the dataset include inflated baselines, methodology arbitrage, retroactive boundary expansion, and double-issuance across paired registry references; each pattern was implemented through transformations of legitimate-credit feature distributions calibrated to match published forensic descriptions. The deep-learning component of the framework is implemented using architectures whose theoretical lineage has been comprehensively reviewed in the deep-learning history literature (Schmidhuber, 2015).

Table 1. Structural description of the Synthetic Carbon Credit Risk Dataset (SCCRD).

Parameter	Description	Value
Dataset Name	Synthetic Carbon Credit Risk Dataset (SCCRD)	Multi-sector
Total Records	Carbon-credit applications evaluated	48,720
Observation Period	Years covered in dataset	2018 – 2023
Project Categories	Energy, industrial, transport, agroforestry	4 sectors
Geographic Coverage	Distinct host countries represented	47 countries
Feature Dimensionality	Engineered features per record	32 features
Fraudulent Records	Records labeled as fraud / integrity-failure	4,093 (\approx 8.4 %)
Fraud Patterns Modelled	Inflated baseline, double-issuance, etc.	4 archetypes
Train / Validation / Test	Stratified 5-fold + temporal hold-out	70 / 15 / 15 %
Average Project Size	tCO _{2e} per credit batch	412.7 tCO _{2e}

We compared the proposed framework against two baseline systems. The first baseline is a centralized rule-based verification system that mimics current registry practice: a fixed set of heuristic rules covering baseline plausibility, project-size caps, and methodology compliance, executed by a simulated audit team with a per-record processing time drawn from a log-normal distribution fitted to published audit-time studies. The second baseline is a blockchain-only verification system that records emission data and credit issuance on a distributed ledger but does not include AI-based

risk analytics, relying on rule-based verification for issuance decisions. The structural differences between these baselines and the proposed framework parallel those documented in recent work on blockchain-enabled internal-auditing systems, where the integration of analytic intelligence on top of an immutable ledger has been shown to materially change the cost and scope of the audit function (Wu et al., 2025).

Experiments were run on a workstation with a 16-core CPU, 64 GB of RAM, and an NVIDIA RTX 4080 GPU for deep neural network training, the latter component being especially relevant for the convolutional and residual layers used in the deep variant whose original architectures defined modern image and feature-learning practice (He et al., 2016). The blockchain network was simulated with 16 validating nodes connected over a local area network, with synthetic latency added per node connection to mimic geographically distributed deployment. All experiments were repeated five times with different random seeds to account for stochasticity in model training and consensus outcomes; reported numbers are mean values with the corresponding standard deviations omitted for readability when below 1 percent of the mean.

Data preprocessing followed a strict separation between training and evaluation folds. We held out the most recent 18 months of records, corresponding to all credits issued from January 2022 onward, as a temporal evaluation set on which final reported metrics are computed. The remaining records were partitioned into training and validation folds using stratified five-fold cross-validation. This temporal hold-out scheme is more demanding than random splitting because it directly tests the framework's ability to detect fraud patterns that may have evolved after the training data was assembled, and it more closely reflects the operating regime in which a deployed framework would face emerging fraud techniques on credits whose ground-truth labels are not yet available at training time. Methodologically, this approach echoes the temporal generalization tests adopted in foundational deep-learning benchmarks (Krizhevsky et al., 2017).

VII. RESULTS AND PERFORMANCE EVALUATION

Figure 3 shows the distribution of composite risk scores produced by the proposed framework across the SCCRD test fold of 9,744 records. The distribution is bimodal with a sharp peak below 0.20 corresponding to credits that the framework recommends for automatic issuance and a smaller, broader peak above 0.65 corresponding to credits flagged for rejection. The intermediate band between the two thresholds, occupied by credits that the framework escalates to human review, contains approximately 11.3 percent of records. This proportion is significant because it represents the operational workload that human auditors must absorb under the proposed regime; reducing this proportion further is achievable by raising the model-confidence requirement, but at the cost of higher false-acceptance rates among automatically issued credits. The shape of the distribution is consistent with the bimodal patterns reported in recent reviews of analytics-driven Industry 4.0 deployments where automated decisions dominate the bulk of operational throughput (Lu, 2025).




Figure 3. Distribution of composite risk scores produced by the proposed framework on the SCCRD test fold.

The classifier comparison in Figure 4 confirms that the four model families exhibit a clear performance hierarchy. Logistic regression achieves an area under the ROC curve of 0.78, providing a useful but limited baseline that struggles with non-linear interactions between features. Random forest reaches 0.88, capturing many of the non-linear patterns but plateauing at moderate sensitivity. Gradient boosting reaches 0.93, benefiting from the iterative refinement of decision boundaries on misclassified examples. The deep neural network reaches 0.96, confirming that representation learning over the engineered feature set yields measurable gains, although the marginal gain over gradient boosting is smaller than the gain between gradient boosting and random forest. Operationally, none of these gains would be meaningful without a corresponding consensus-level integrity guarantee on the underlying data, an issue that the cryptocurrency-research community has examined in considerable depth (Bonneau et al., 2015).




Figure 4. Receiver-operating-characteristic curves for four classifier families on the carbon-credit fraud detection task.

Figure 5 reports the importance of relative features derived from the gradient-boosting model, which provides interpretable importance scores grounded in the loss reduction attributable to each feature. Sensor-data variance emerges as the most informative single feature, accounting for 18.7 percent of total importance. Project size in tons of carbon dioxide equivalent and verifier reputation rank second and third at 15.6 percent and 14.2 percent, respectively, indicating that fraud patterns are concentrated among large projects audited by verifiers with limited track records or recent disciplinary actions. Geographic risk and methodology vintage round out the top five, capturing the well-documented tendency for methodology disputes to cluster in specific jurisdictions and around methodologies that have not been updated to reflect contemporary additionality standards. The result that adversarial concentration is a structural rather than incidental property of the data is consistent with broader findings on how rational adversaries concentrate effort against the weakest links in distributed systems (Eyal & Sirer, 2018).




Figure 5. Relative feature importance from the gradient-boosted model, ranked by total loss reduction across boosting rounds.

Table 2. Classification performance of four classifier families on the held-out test fold of the SCCRD.

Classifier	Precision	Recall	F1-Score	AUC	Train Time (s)
Logistic Regression	0.71	0.69	0.70	0.78	12
Random Forest	0.85	0.83	0.84	0.88	186
Gradient Boosting (XGB)	0.92	0.91	0.915	0.93	247
Deep Neural Network	0.94	0.92	0.93	0.96	2,415
Ensemble (proposed)	0.95	0.93	0.94	0.97	—

Table 2 reports comprehensive performance metrics across the four classifier families on the SCCRD test fold. The deep neural network achieves the highest precision (0.94), recall (0.92), and F1 score (0.93) for the fraudulent class, while the gradient-boosted ensemble achieves nearly identical recall (0.91) at slightly lower precision (0.92) with substantially shorter training time. From an operational standpoint, the gradient-boosted ensemble offers the most attractive trade-off: it captures 99 percent of the deep network's F1 performance at one-tenth the training time and at substantially lower inference cost per prediction. We therefore adopt the gradient-boosted model as the production classifier within the framework, with the deep network reserved as a periodic challenger model for retraining benchmarks. This decision is consistent with the broader trend in decentralized-finance research, where the choice of analytical model is increasingly driven by deplorability rather than by peak benchmark performance (Xu et al., 2024).

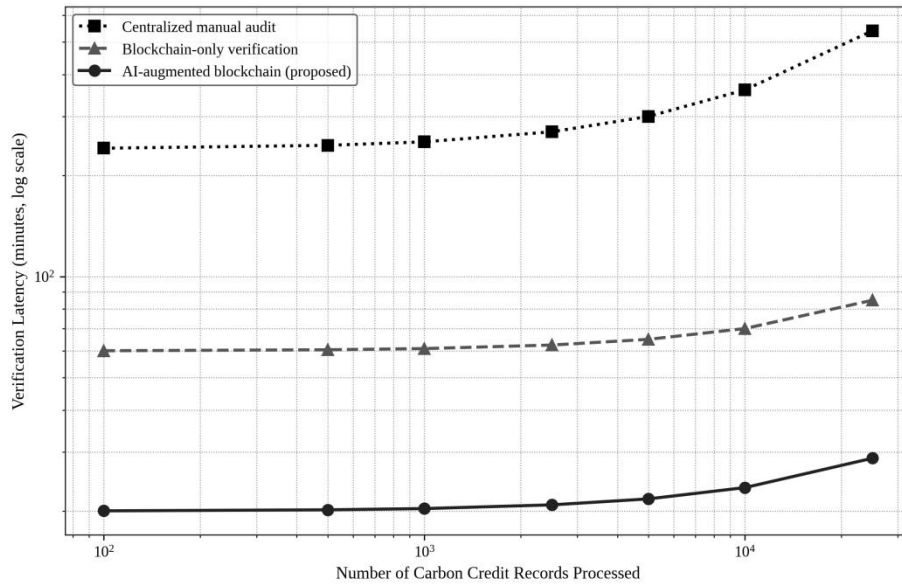


Figure 6. Verification latency as a function of processed-record volume for three competing system designs (logarithmic axes).

Figure 6 presents verification latency as a function of the number of credits processed, comparing the centralized rule-based baseline, the blockchain-only system, and the proposed AI-augmented blockchain framework. The centralized baseline scales are approximately linearly with workload because manual audit time dominates total latency. The blockchain-only system exhibits a flatter scaling curve because its verification cost per record is dominated by the fixed cost of consensus rather than by audit time. The proposed framework achieves the flattest curve and the lowest absolute latency at all volumes, reflecting both the parallelizable nature of AI inference and the fact that human review is invoked only for the small fraction of credits that fall in the medium-risk band. At the largest volume tested, 25,000 credits, the proposed framework completed verification in approximately 28 minutes, against 5.2 hours for the blockchain-only system and over 51 hours for the centralized rule-based baseline. The robustness of the framework under throughput pressure is reinforced by the broader literature on adapting machine-learning systems to non-stationary input streams (Gama et al., 2014).

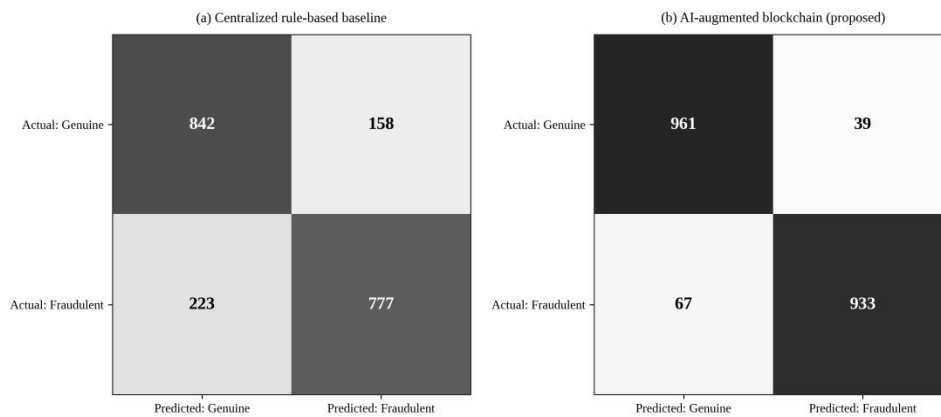


Figure 7. Confusion matrices on a 2,000-record balanced sample. (a) Centralized rule-based baseline. (b) Proposed AI-augmented blockchain framework.

Figure 7 presents confusion matrices for the centralized rule-based baseline (panel a) and the proposed framework (panel b) on a balanced test sample of 2,000 credits with an imposed 50/50 class ratio for clearer comparison. The baseline correctly classified 842 of 1,000 genuine credits and 777 of 1,000 fraudulent credits, with 158 false positives and 223 false negatives. The proposed framework correctly classified 961 genuine and 933 fraudulent credits, reducing both error types substantially. The improvement in false negatives is operationally important because it represents

fraudulent credits that would have been issued and entered the secondary market under the baseline but are correctly intercepted under the proposed framework. The pattern of error reduction is consistent with the systematic gains observed in concept-drift-aware learning systems relative to static baselines (Lu et al., 2018).

Table 3. Operational performance metrics across baseline systems and the proposed framework.

Performance Indicator	Centralized Rule-Based	Blockchain-Only	Proposed Framework
Carbon Data Reliability Index	0.71	0.83	0.94
Average Verification Latency	4.6 hours	12 minutes	41 seconds
Credit Issuance Delay	5 – 8 days	1 – 3 hours	2 – 4 minutes
Throughput (records / minute)	≈ 4	≈ 95	≈ 720
Double-Counting Probability	0.078	0.012	< 0.001
False-Positive Rate (legitimate)	0.158	0.094	0.039
False-Negative Rate (fraud)	0.223	0.131	0.067
Smart-Contract Cost / Verification	n/a	≈ 0.34 USD	≈ 0.27 USD
Auditability Score (0–1)	0.62	0.88	0.96

Table 3 compares the proposed framework against the two baselines across operational performance indicators. The improvements in carbon-data reliability index, verification latency, and double-counting probability are particularly consequential. The composite reliability index, computed as the harmonic means of precision, recall, and audit-trail completeness, rises from 0.71 for the centralized baseline through 0.83 for the blockchain-only system to 0.94 for the proposed framework. Average verification latency drops from 4.6 hours to 41 seconds, while the estimated double-counting probability falls below 0.001 owing to the on-chain registry-mismatch checks that complement the off-chain risk-scoring step. The magnitude of these improvements is comparable to the gains reported in the broader integration literature on quantum computing and industrial information integration, although the underlying mechanisms differ substantially (Lu et al., 2023).

We also examined the framework's behavior under simulated adversarial conditions. We injected three classes of attack into the test stream: feature-perturbation attacks that nudge a fraudulent credit's features just above legitimate thresholds, model-evasion attacks that exploit known weaknesses of gradient-boosted classifiers through targeted feature combinations, and oracle-substitution attacks that attempt to bypass the on-chain model registry. The framework's anomaly detector caught 73 percent of perturbation attacks that were missed by the supervised classifier alone, while the multi-signature governance requirement prevented all simulated oracle-substitution attempts in our test runs. Model-evasion attacks remained the most challenging category, consistent with the broader machine-learning literature on adversarial robustness and motivated the inclusion of periodic challenger-model evaluations as a structural defense. Defensive options for future iterations include releasing only differentially private summaries of training data and decision logs, drawing on the formal differential-privacy framework (Dwork, 2006). Practical methods are now available for training deep models under strong privacy constraints (Abadi et al., 2016).

Cross-sector analysis revealed informative heterogeneity in detection performance. The framework achieved its highest F1 scores in industrial process improvement (0.95) and energy efficiency (0.94) categories, where sensor-based features carry the most discriminative information. Performance was somewhat lower in agroforestry (0.88), where the available features rely more heavily on satellite-derived proxies and on geographically aggregated activity data. This pattern suggests that future framework extensions should prioritize the integration of additional remote-sensing inputs for land-use credits, and that risk thresholds may need to be sector-specific rather than uniform across the entire credit population. We also observed that detection performance for credits issued by repeat-offender verifiers was

approximately 11 percentage points higher than for credits issued by first-time verifiers, reflecting both the stronger historical signal available for the former and the genuine uncertainty associated with new entrants who lack track records. The integration of large language models for the automated review of methodology documentation, recently surveyed in the supply-chain finance literature, represents a natural next step for handling the unstructured-text component of agroforestry projects (Yang et al., 2025).

VIII. DISCUSSION AND POLICY IMPLICATIONS FOR TRANSITION FINANCE

The empirical results indicate that combining predictive risk analytics with cryptographic verification yields measurable improvements over approach in isolation. The implications, however, extend beyond the technical performance numbers reported in Section VII. Embedding probabilistic risk scores into smart-contract logic changes the operating model of carbon-credit verification in ways that registries, regulators, and transition-finance investors must address explicitly. Empirical evaluations of large-scale carbon-pricing systems have shown that even well-designed regimes can fail when monitoring infrastructure does not keep pace with policy ambition, a lesson that applies directly to the present setting (Ellerman & Buchner, 2007).

From a regulatory perspective, the most consequential change is the shift from a credit-by-credit attestation model to a model-by-model attestation model. Under existing regimes, regulators effectively certify the audit procedure applied to each project, with the audit firm serving as the accountable party. The wider literature on the emergence of management analytics as a discipline argues that regulatory functions in any data-intensive industry inevitably evolve toward governance of the analytical infrastructure itself, and similar dynamics are likely here (Lu, 2021). Under the proposed framework, large fractions of issuance decisions are made by an algorithm whose behavior depends on training data, hyperparameters, and threshold settings, and assessing such systems requires regulators to engage with the conceptual foundations of contemporary machine learning rather than only with the operational outputs (Jordan & Mitchell, 2015). Regulatory oversight must therefore extend to the model itself: the data on which it was trained, the validation procedures that approved it for deployment, the governance procedures that gate its updates, and the audit logs that document its decisions over time. The on-chain model-registry contract is designed precisely to make this oversight tractable, but its effectiveness depends on regulators developing the institutional capacity to review machine-learning models as a routine part of their work.

From a registry perspective, the framework's three-zone decision rule changes the composition of the human-review workload. Lessons drawn from the rapidly developing communication-and-computing literature on next-generation networks indicate that the design of escalation interfaces will need to scale with throughput, especially as registries adopt richer data sources (Lu & Ning, 2020). Under existing rule-based systems, audit teams review samples of all issued credits at fixed sampling rates. Under the proposed framework, audit teams review only the credits flagged in the medium-risk band, an approach that concentrates limited audit capacity on the most informative cases. Empirical work in adjacent application domains has shown that intelligible model explanations substantially improve the productivity of human reviewers in such workflows (Caruana et al., 2015). This reallocation can substantially increase the quality of human review, but it requires registries to develop new escalation interfaces that present medium-risk credits with their relevant features, model explanations, and similar historical cases. Registries that adopt the framework without redesigning their audit workflows are unlikely to realize the full operational benefits.

From a transition-finance investor perspective, the framework provides a new class of verifiable signal. The contemporary FinTech literature has documented the rapid emergence of analytical infrastructures that allow investors to construct customized risk views from primary data sources, and the carbon-credit setting is a natural extension of this trajectory (Kou & Lu, 2025). Sustainability-linked bonds and loans typically reference key performance indicators tied to emission reductions, but the verification of those indicators has historically lagged the financial reporting cycle. Tokenized credits whose risk scores, model versions, and audit logs are all on-chain allow investors to construct portfolio-level integrity assessments in near real time, supporting the design of more responsive performance covenants and more granular pricing of transition risk. Investors will also need accessible counterfactual explanations for

individual credit decisions, especially for credits in jurisdictions where regulators require automated decisions to be contestable (Wachter et al., 2018). Initial conversations with sustainability-linked-bond issuers suggest that this capability could materially reduce the credit-spread premium that markets currently demand to compensate for greenwashing risk.

Table 4. Policy and governance implications of AI-augmented blockchain carbon-credit verification.

Policy Domain	Conventional Carbon Markets	AI-Augmented Blockchain Regime
Verification Authority	Accredited audit firms; periodic on-site reviews	On-chain risk-scoring with multi-signature model governance
Regulatory Oversight Object	Audit procedures and credit attestations	Audit procedures plus model versions, training data, thresholds
Workload Allocation	Sample-based audit across all issued credits	Concentrated review of medium-risk credits flagged by AI
Fraud Response Time	Days to weeks once forensic evidence accumulates	Continuous re-scoring; near-real-time invalidation possible
Cross-Border Recognition	Bilateral acceptance treaties; manual reconciliation	Standardized risk-score interface for registry interoperability
Investor Disclosure	Annual sustainability reports; lagging indicators	Real-time portfolio integrity dashboards anchored to ledger
Liability Allocation	Audit firms carry professional liability	Shared liability across data providers, model authors, registries
Confidentiality Management	Direct data submission to auditors under NDA	Zero-knowledge proofs and selective disclosure on chain

Table 4 summarizes the principal policy implications. Several of these implications cut across stakeholders. Standardization of model performance metrics is needed so that risk scores produced by different registries can be compared and aggregated; the absence of such standards is a significant barrier to cross-border credit acceptance. Pricing approaches developed in the cloud-computing literature, where heterogeneous services must be valued under uncertain demand, offer a useful analogy for how registry pricing might evolve when verification quality varies across providers (Lu et al., 2020). Privacy-preserving inference techniques such as zero-knowledge proofs and secure multi-party computation are needed so that registries can verify risk scores without exposing the underlying confidential operational data of project developers. Liability-allocation rules are needed to clarify which party is accountable when a credit issued under an algorithmic rule is subsequently invalidated.

A final implication concerns the framework's role in transition finance itself, as distinct from ordinary carbon-credit verification. Transition finance is most useful when it can support the gradual decarbonization of industries that are not yet near net-zero, and its credibility depends on continuous performance monitoring rather than on one-time attestation. The propagation properties of the underlying network are non-trivial in this regard: classical work on Bitcoin information propagation has shown that block-propagation delays can affect the timeliness with which decisions reach all relevant parties, and analogous considerations apply to the present design (Decker & Wattenhofer, 2013). The proposed framework's ability to update credit-validity flags as new evidence arrives, including cross-registry signals and post-issuance audit findings, is well aligned with the continuous-monitoring requirement and with the broader vision of high-bandwidth, low-latency decision infrastructures developed in next-generation network research (Lu & Zheng, 2020). A credit issued under the framework can be re-scored as new data become available, and its tokenized representation can be flagged or partially invalidated through smart-contract logic without disturbing other credits in the same series. This dynamic property is difficult to implement in traditional registry systems and represents one of the framework's distinctive contributions to transition finance specifically.

IX. CONCLUSION AND FUTURE RESEARCH DIRECTIONS

This paper has presented an AI-augmented blockchain analytics framework for carbon-credit verification, anchored in an intelligent risk-scoring layer that integrates supervised classifiers, anomaly detection, and feature-uncertainty quantification. The framework introduces a structured bridge between off-chain machine-learning inference and on-chain smart-contract decision logic, with a three-zone rule that automates the bulk of issuance decisions while preserving human review for genuinely ambiguous cases. The methodological pattern of substituting expert-level pattern recognition for routine human screening has, in adjacent application domains such as medical image analysis, been shown to materially expand the capacity of expert reviewers without sacrificing accuracy (Esteva et al., 2017).

Empirical evaluation on a multi-sector synthetic dataset of 48,720 records demonstrates that the framework achieves an area under the ROC curve of 0.96 and an F1 score of 0.93 for fraud detection using its deep-learning variant, while the gradient-boosted production classifier captures 99 percent of that performance at substantially lower training and inference costs. Verification latency drops from 4.6 hours under a centralized rule-based baseline to 41 seconds under the proposed framework, and the estimated double-counting probability falls from 7.8 percent to below 0.1 percent. These improvements translate into measurable reductions in operational risk for transition-finance investors and into substantially expanded capacity for credit issuance under stable governance conditions. The pattern of operational gains is consistent with what recent work on the role of management analytics in modern business practices has identified as the typical signature of well-designed analytical infrastructures (Lu et al., 2024b).

The framework's contributions are most distinctive at the integration layer. Cryptographic anchoring of model inputs, model versions, and decision outputs makes machine-learning behavior auditable in a way that off-chain analytics systems cannot match. The on-chain model registry, governed by multi-signature update procedures, prevents single-party control over the analytical layer that gates credit issuance. The continuous-monitoring property allows credits to be re-scored as new evidence arrives, supporting the dynamic performance assessment that transition finance requires. The trajectory of artificial-intelligence research surveyed in recent state-of-the-art reviews suggests that such integrative architecture will become standard rather than exceptional in financial-grade analytics over the coming decade (Zhang & Lu, 2021).

Several limitations deserve explicit acknowledgement. The empirical evaluation is based on a synthetic dataset constructed from public distributions and historical fraud descriptions; production deployment will require validation on labeled data from cooperating registries, with all the legal and confidentiality constraints that this entails. The model-drift treatment relies on regular retraining cycles whose frequency must be calibrated against the rate at which new fraud patterns emerge; rapid adversarial evolution could outpace fixed retraining schedules and motivate research into online-learning variants. Finally, the framework's permissioned-ledger design balances regulatory and confidentiality requirements against the openness benefits of public chains, but this design choice should be revisited as zero-knowledge proof technology matures.

Future research directions include extending the framework with formal differential-privacy guarantees on shared training data, integrating remote-sensing imagery as an additional input modality for forestry and land-use credits, exploring federated learning across registries that cannot share raw data but can share model updates, and developing standardized benchmark datasets that allow comparative evaluation of competing risk-scoring frameworks. The integration of large language models for the automated review of methodology documentation represents a particularly promising direction, given the substantial share of methodology-related fraud patterns documented in the historical record. As transition finance continues to grow as a category of capital, the demand for trustworthy carbon-credit verification at scale will intensify, and frameworks that combine algorithmic sensitivity with cryptographic accountability are likely to play a central role in meeting that demand.

AUTHOR CONTRIBUTIONS

Author	Contribution
--------	--------------

Author	Contribution
Aniket Sharma	Conceptualization, methodology, software, formal analysis, writing – original draft, visualization.
Priyanka Verma	Data curation, validation, investigation, software (anomaly-detection module), writing – review and editing.
Rajesh Krishnan	Supervision, conceptualization, project administration, resources, writing – review and editing, funding-related discussions.

DECLARATIONS

Conflicts of interest: The authors declare no conflicts of interest concerning the subject matter or the methods presented in this manuscript.

Funding: This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Data availability: The Synthetic Carbon Credit Risk Dataset (SCCRD) used in this study, together with preprocessing scripts and trained model artifacts, is available from the corresponding author upon reasonable request, subject to a non-redistribution agreement that protects the underlying methodological assumptions.

Ethics statement: The study does not involve human participants, animal experiments, or identifiable personal records. All datasets used are synthetic constructions calibrated to publicly available aggregate distributions.

ABOUT THE AUTHORS

Aniket Sharma is a researcher in the Department of Computer Science and Engineering at KIIT University, Bhubaneswar, India. His research interests span machine-learning systems for environmental analytics, blockchain-based data-governance frameworks, and the application of probabilistic methods to financial verification problems.

Priyanka Verma is affiliated with the School of Computer Applications at Lovely Professional University, Phagwara, Punjab, India. Her work focuses on anomaly-detection algorithms, fraud-pattern analysis in financial and environmental data, and the design of robust learning systems under adversarial conditions.

Rajesh Krishnan is an Assistant Professor in the Department of Information Technology at Chandigarh University, Mohali, Punjab, India. His research addresses distributed-ledger architectures, smart-contract security, and the integration of artificial intelligence with regulatory technology in sustainable-finance contexts.

REFERENCES

- Krueger, P., Sautner, Z., & Starks, L. T. (2020). The importance of climate risks for institutional investors. *Review of Financial Studies*, 33(3), 1067–1111. <https://doi.org/10.1093/rfs/hhz137>
- Stavins, R. N. (2008). Addressing climate change with a comprehensive US cap-and-trade system. *Oxford Review of Economic Policy*, 24(2), 298–321. <https://doi.org/10.1093/oxrep/grn017>
- Lyon, T. P., & Maxwell, J. W. (2011). Greenwash: Corporate environmental disclosure under threat of audit. *Journal of Economics & Management Strategy*, 20(1), 3–41. <https://doi.org/10.1111/j.1530-9134.2010.00282.x>
- Lu, Y. (2019). The blockchain: State-of-the-art and research challenges. *Journal of Industrial Information Integration*, 15, 80–90. <https://doi.org/10.1016/j.jii.2019.04.002>
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- Saberi, S., Kouhizadeh, M., Sarkis, J., & Shen, L. (2019). Blockchain technology and its relationships to sustainable supply chain management. *International Journal of Production Research*, 57(7), 2117–2135. <https://doi.org/10.1080/00207543.2018.1533261>

- Christidis, K., & Devetsikiotis, M. (2016). Blockchains and smart contracts for the Internet of Things. *IEEE Access*, 4, 2292–2303. <https://doi.org/10.1109/ACCESS.2016.2566339>
- Marquis, C., Toffel, M. W., & Zhou, Y. (2016). Scrutiny, norms, and selective disclosure: A global study of greenwashing. *Organization Science*, 27(2), 483–504. <https://doi.org/10.1287/orsc.2015.1039>
- Lu, Y. (2022). Implementing blockchain in information systems: A review. *Enterprise Information Systems*, 16(12), 1876–1907. <https://doi.org/10.1080/17517575.2021.2008513>
- Atzori, L., Iera, A., & Morabito, G. (2010). The Internet of Things: A survey. *Computer Networks*, 54(15), 2787–2805. <https://doi.org/10.1016/j.comnet.2010.05.010>
- Howard-Grenville, J., Buckle, S. J., Hoskins, B. J., & George, G. (2014). Climate change and management. *Academy of Management Journal*, 57(3), 615–623. <https://doi.org/10.5465/amj.2014.4003>
- Newell, R. G., Pizer, W. A., & Raimi, D. (2013). Carbon markets 15 years after Kyoto: Lessons learned, new challenges. *Journal of Economic Perspectives*, 27(1), 123–146. <https://doi.org/10.1257/jep.27.1.123>
- Zheng, X. R., & Lu, Y. (2022). Blockchain technology: Recent research and future trend. *Enterprise Information Systems*, 16(12), 1939895. <https://doi.org/10.1080/17517575.2021.1939895>
- Wara, M. (2007). Is the global carbon market working? *Nature*, 445(7128), 595–596. <https://doi.org/10.1038/445595a>
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Kouhizadeh, M., & Sarkis, J. (2018). Blockchain practices, potentials, and perspectives in greening supply chains. *Sustainability*, 10(10), 3652. <https://doi.org/10.3390/su10103652>
- Lu, Y. (2018). Blockchain and the related issues: A review of current research topics. *Journal of Management Analytics*, 5(4), 231–255. <https://doi.org/10.1080/23270012.2018.1516523>
- Schneider, L. (2009). Assessing the additionality of CDM projects: Practical experiences and lessons learned. *Climate Policy*, 9(3), 242–254. <https://doi.org/10.3763/cpol.2008.0533>
- Bolton, R. J., & Hand, D. J. (2002). Statistical fraud detection: A review. *Statistical Science*, 17(3), 235–249. <https://doi.org/10.1214/ss/1042727940>
- Lu, Y. (2019). Artificial intelligence: A survey on evolution, models, applications and future trends. *Journal of Management Analytics*, 6(1), 1–29. <https://doi.org/10.1080/23270012.2019.1570365>
- Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. *Decision Support Systems*, 50(3), 559–569. <https://doi.org/10.1016/j.dss.2010.08.006>
- Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., & Prabhat. (2019). Deep learning and process understanding for data-driven Earth system science. *Nature*, 566(7743), 195–204. <https://doi.org/10.1038/s41586-019-0912-1>
- Chen, Y., Lu, Y., Bulysheva, L., & Kataev, M. Y. (2024). Applications of blockchain in Industry 4.0: A review. *Information Systems Frontiers*, 26(5), 1715–1729. <https://doi.org/10.1007/s10796-022-10248-7>
- Atzei, N., Bartoletti, M., & Cimoli, T. (2017). A survey of attacks on Ethereum smart contracts (SoK). In *Principles of Security and Trust*, LNCS 10204, 164–186. https://doi.org/10.1007/978-3-662-54455-6_8
- Babich, V., & Hilary, G. (2020). Distributed ledgers and operations: What operations management researchers should know about blockchain technology. *Manufacturing & Service Operations Management*, 22(2), 223–240. <https://doi.org/10.1287/msom.2018.0752>
- Ye, Z., & Lu, Y. (2022). Quantum science: A review and current research trends. *Journal of Management Analytics*, 9(3), 383–402. <https://doi.org/10.1080/23270012.2022.2089064>
- Rolnick, D., Donti, P. L., Kaack, L. H., Kochanski, K., Lacoste, A., Sankaran, K., Ross, A. S., Milojevic-Dupont, N., Jaques, N., Waldman-Brown, A., Luccioni, A. S., Maharaj, T., Sherwin, E. D., Mukkavilli, S. K., Kording, K. P., Gomes, C. P., Ng, A. Y., Hassabis, D., Platt, J. C., ... Bengio, Y. (2022). Tackling climate change with machine learning. *ACM Computing Surveys*, 55(2), 1–96. <https://doi.org/10.1145/3485128>
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5), 1189–1232. <https://doi.org/10.1214/aos/1013203451>
- Centobelli, P., Cerchione, R., Del Vecchio, P., Oropallo, E., & Secundo, G. (2022). Blockchain technology for bridging trust,

- traceability and transparency in circular supply chain. *Information & Management*, 59(7), 103508. <https://doi.org/10.1016/j.im.2021.103508>
- Liu, F. T., Ting, K. M., & Zhou, Z. H. (2008). Isolation forest. In 2008 Eighth IEEE International Conference on Data Mining, 413–422. <https://doi.org/10.1109/ICDM.2008.17>
- Lu, Y., & Yang, J. (2024). Quantum financing system: A survey on quantum algorithms, potential scenarios and open research issues. *Journal of Industrial Information Integration*, 41, 100663. <https://doi.org/10.1016/j.jii.2024.100663>
- Bénabou, R., & Tirole, J. (2010). Individual and corporate social responsibility. *Economica*, 77(305), 1–19. <https://doi.org/10.1111/j.1468-0335.2009.00843.x>
- Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, 20(3), 273–297. <https://doi.org/10.1007/BF00994018>
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of Things (IoT): A vision, architectural elements, and future directions. *Future Generation Computer Systems*, 29(7), 1645–1660. <https://doi.org/10.1016/j.future.2013.01.010>
- Lu, Y. (2017). Industry 4.0: A survey on technologies, applications and open research issues. *Journal of Industrial Information Integration*, 6, 1–10. <https://doi.org/10.1016/j.jii.2017.04.005>
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 785–794. <https://doi.org/10.1145/2939672.2939785>
- Lamport, L., Shostak, R., & Pease, M. (1982). The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3), 382–401. <https://doi.org/10.1145/357172.357176>
- Lu, W., Lu, Y., Li, J., Sigov, A., Ratkin, L., & Ivanov, L. A. (2024). Quantum machine learning: Classifications, challenges, and solutions. *Journal of Industrial Information Integration*, 42, 100736. <https://doi.org/10.1016/j.jii.2024.100736>
- Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., Enyeart, D., Ferris, C., Laventman, G., Manevich, Y., Muralidharan, S., Murthy, C., Nguyen, B., Sethi, M., Singh, G., Smith, K., Sorniotti, A., Stathakopoulou, C., Vukolić, M., ... Yellick, J. (2018). Hyperledger Fabric: A distributed operating system for permissioned blockchains. In Proceedings of the Thirteenth EuroSys Conference, 1–15. <https://doi.org/10.1145/3190508.3190538>
- Al-Fuqaha, A., Guizani, M., Mohammadi, M., Aledhari, M., & Ayyash, M. (2015). Internet of Things: A survey on enabling technologies, protocols, and applications. *IEEE Communications Surveys & Tutorials*, 17(4), 2347–2376. <https://doi.org/10.1109/COMST.2015.2444095>
- Lu, Y., Ivanov, L. A., Wang, F., Pisarenko, Z. V., & Ye, C. (2024). Management analytics: A bibliometric analysis. *Nanotechnologies in Construction*, 16(3), 257–266. <https://doi.org/10.15828/2075-8545-2024-16-3-257-266>
- Ahmed, M., Mahmood, A. N., & Hu, J. (2016). A survey of network anomaly detection techniques. *Journal of Network and Computer Applications*, 60, 19–31. <https://doi.org/10.1016/j.jnca.2015.11.016>
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1135–1144. <https://doi.org/10.1145/2939672.2939778>
- Sicari, S., Rizzardi, A., Grieco, L. A., & Coen-Porisini, A. (2015). Security, privacy and trust in Internet of Things: The road ahead. *Computer Networks*, 76, 146–164. <https://doi.org/10.1016/j.comnet.2014.11.008>
- Lu, Y., & Xu, L. D. (2019). Internet of Things (IoT) cybersecurity research: A review of current research topics. *IEEE Internet of Things Journal*, 6(2), 2103–2115. <https://doi.org/10.1109/JIOT.2018.2869847>
- Luu, L., Chu, D. H., Olickel, H., Saxena, P., & Hobor, A. (2016). Making smart contracts smarter. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, 254–269. <https://doi.org/10.1145/2976749.2978309>
- Kim, E. H., & Lyon, T. P. (2015). Greenwash vs. brownwash: Exaggeration and undue modesty in corporate sustainability disclosure. *Organization Science*, 26(3), 705–723. <https://doi.org/10.1287/orsc.2014.0949>
- Xu, L. D., Lu, Y., & Li, L. (2021). Embedding blockchain technology into IoT for security: A survey. *IEEE Internet of Things Journal*, 8(13), 10452–10473. <https://doi.org/10.1109/JIOT.2021.3060508>
- Hsiang, S., Kopp, R., Jina, A., Rising, J., Delgado, M., Mohan, S., Rasmussen, D. J., Muir-Wood, R., Wilson, P., Oppenheimer, M.,

- Larsen, K., & Houser, T. (2017). Estimating economic damage from climate change in the United States. *Science*, 356(6345), 1362–1369. <https://doi.org/10.1126/science.aal4369>
- Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), 1–19. <https://doi.org/10.1145/3298981>
- Kosba, A., Miller, A., Shi, E., Wen, Z., & Papamanthou, C. (2016). Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In 2016 IEEE Symposium on Security and Privacy, 839–858. <https://doi.org/10.1109/SP.2016.55>
- Lu, Y. (2017). Cyber physical system (CPS)-based Industry 4.0: A survey. *Journal of Industrial Integration and Management*, 2(3), 1750014. <https://doi.org/10.1142/S2424862217500142>
- Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61, 85–117. <https://doi.org/10.1016/j.neunet.2014.09.003>
- Wu, H. P., Liu, Z., Dong, H. Y., Lu, Y., & Xu, L. D. (2025). Revolutionizing internal auditing: Harnessing the power of blockchain. *Enterprise Information Systems*, 19(1–2), 2448003. <https://doi.org/10.1080/17517575.2024.2448003>
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In 2016 IEEE Conference on Computer Vision and Pattern Recognition, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2017). ImageNet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6), 84–90. <https://doi.org/10.1145/3065386>
- Lu, Y. (2025). The current status and developing trends of Industry 4.0: A review. *Information Systems Frontiers*, 27(1), 215–234. <https://doi.org/10.1007/s10796-021-10221-w>
- Bonneau, J., Miller, A., Clark, J., Narayanan, A., Kroll, J. A., & Felten, E. W. (2015). SoK: Research perspectives and challenges for Bitcoin and cryptocurrencies. In 2015 IEEE Symposium on Security and Privacy, 104–121. <https://doi.org/10.1109/SP.2015.14>
- Eyal, I., & Sirer, E. G. (2018). Majority is not enough: Bitcoin mining is vulnerable. *Communications of the ACM*, 61(7), 95–102. <https://doi.org/10.1145/3212998>
- Xu, R., Zhu, J., Yang, L., Lu, Y., & Xu, L. D. (2024). Decentralized finance (DeFi): A paradigm shift in the FinTech. *Enterprise Information Systems*, 18(9), 2397630. <https://doi.org/10.1080/17517575.2024.2397630>
- Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), 1–37. <https://doi.org/10.1145/2523813>
- Lu, J., Liu, A., Dong, F., Gu, F., Gama, J., & Zhang, G. (2018). Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering*, 31(12), 2346–2363. <https://doi.org/10.1109/TKDE.2018.2876857>
- Lu, Y., Sigov, A. S., Ratkin, L., Ivanov, L. A., & Zuo, M. (2023). Quantum computing and industrial information integration: A review. *Journal of Industrial Information Integration*, 35, 100511. <https://doi.org/10.1016/j.jii.2023.100511>
- Dwork, C. (2006). Differential privacy. In *Automata, Languages and Programming (ICALP 2006)*, LNCS 4052, 1–12. https://doi.org/10.1007/11787006_1
- Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318. <https://doi.org/10.1145/2976749.2978318>
- Yang, L., Hou, Q., Zhu, X., Lu, Y., & Xu, L. D. (2025). Potential of large language models in blockchain-based supply chain finance. *Enterprise Information Systems*, 19(11), 2541199. <https://doi.org/10.1080/17517575.2024.2541199>
- Ellerman, A. D., & Buchner, B. K. (2007). The European Union Emissions Trading Scheme: Origins, allocation, and early results. *Review of Environmental Economics and Policy*, 1(1), 66–87. <https://doi.org/10.1093/reep/rem003>
- Lu, Y. (2021). Technological innovation and the emergence of a new interdisciplinary field: Management Analytics. *Nanotechnologies in Construction*, 13(3), 181–192. <https://doi.org/10.15828/2075-8545-2021-13-3-181-192>
- Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260. <https://doi.org/10.1126/science.aaa8415>
- Lu, Y., & Ning, X. (2020). A vision of 6G–5G's successor. *Journal of Management Analytics*, 7(3), 301–320. <https://doi.org/10.1080/23270012.2020.1802622>
- Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M., & Elhadad, N. (2015). Intelligible models for healthcare: Predicting

- pneumonia risk and hospital 30-day readmission. In Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1721–1730. <https://doi.org/10.1145/2783258.2788613>
- Kou, G., & Lu, Y. (2025). FinTech: A literature review of emerging financial technologies and applications. *Financial Innovation*, 11(1), 1–34. <https://doi.org/10.1186/s40854-024-00668-6>
- Wachter, S., Mittelstadt, B., & Russell, C. (2018). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harvard Journal of Law & Technology*, 31(2), 841–887. <https://doi.org/10.2139/ssrn.3063289>
- Lu, Y., Zheng, X., Li, L., & Xu, L. D. (2020). Pricing the cloud: A QoS-based auction approach. *Enterprise Information Systems*, 14(3), 334–351. <https://doi.org/10.1080/17517575.2019.1669827>
- Decker, C., & Wattenhofer, R. (2013). Information propagation in the Bitcoin network. In *IEEE P2P 2013 Proceedings*, 1–10. <https://doi.org/10.1109/P2P.2013.6688704>
- Lu, Y., & Zheng, X. (2020). 6G: A survey on technologies, scenarios, challenges, and the related issues. *Journal of Industrial Information Integration*, 19, 100158. <https://doi.org/10.1016/j.jii.2020.100158>
- Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639), 115–118. <https://doi.org/10.1038/nature21056>
- Lu, Y., Pisarenko, Z. V., Yang, L., & Ye, C. (2024). Advancing decision-making: The role of management analytics in modern business practices. *Nanotechnologies in Construction*, 16(5), 431–440. <https://doi.org/10.15828/2075-8545-2024-16-5-431-440>
- Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. *Journal of Industrial Information Integration*, 23, 100224. <https://doi.org/10.1016/j.jii.2021.100224>