



Automated Fiscal Compliance for Romanian Micro-Enterprises: An Integrated Blockchain, AI and IoT Architecture

Petronela Alice Grigorescu^{*}, Ruxandra Georgescu (Grigorescu)^{**},
Alexandru Cătălin Neagu^{***}, Dan Marius Coman^{****}

ARTICLE INFO

Article history:

Received March 26, 2026

Accepted April 07, 2026

Available online April 30, 2026

JEL Classification

M41, O33, L25, H25

Keywords:

smart contracts; anomaly detection, tax administration, digital transformation, design science research

ABSTRACT

Fiscal compliance represents a critical yet operationally demanding requirement for micro-enterprises, which typically lack dedicated accounting infrastructure. This article presents a functional proof-of-concept platform integrating Blockchain, Artificial Intelligence (AI), and the Internet of Things (IoT) to automate fiscal compliance workflows. The prototype was developed using a Design Science Research methodology and validated through a simulated dataset of 18 financial transactions representative of Romanian micro-enterprise operations. Machine learning anomaly detection achieved 94% accuracy with a 6% false positive rate. Smart contract execution yielded AI decision confidence ranging from 75% to 96% across decision categories, and projected annual savings of €3,450 with a reduction of 127 hours of manual labour per month. Findings confirm the technical feasibility and economic viability of integrated BC-AI-IoT compliance architectures, while identifying scalability and regulatory alignment as priority directions for future research.

Economics and Applied Informatics © 2026 is licensed under [CC BY 4.0](https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Micro-enterprises constitute the structural foundation of both the European Union and Romanian economies, representing over 90% of all registered business entities in Romania. Despite their economic significance, these entities remain disproportionately exposed to fiscal compliance risks derived from manual transaction recording, limited access to advanced digital tools, and a volatile legislative environment.

At the European level, digital transformation of SMEs has emerged as a strategic priority. According to Eurostat data (2024), 73% of SMEs in the EU have reached a basic level of digital intensity; however, the 2030 target of 90% remains out of reach, with persistent challenges related to competencies, financial resources, and resistance to change (European Commission, 2024). Micro-enterprises face the steepest barriers: they are frequently excluded from enterprise-grade digital compliance tools designed primarily for large corporations.

The limitations of existing solutions are particularly acute in fiscal compliance. Traditional ERP systems - SAP, Oracle, QuickBooks - address compliance in a fragmented manner, without offering a holistic, scalable, preventive validation framework adapted to the specific needs of micro-enterprises. As a result, compliance remains a reactive, labour-intensive activity prone to human error, manipulation, and delayed detection of irregularities.

The convergence of three transformative technologies, namely Blockchain (BC), Artificial Intelligence (AI), and the Internet of Things (IoT), offers a novel solution. Smart contracts enable autonomous, real-time transaction validation against predefined fiscal rules (Kokina et al., 2017). Machine learning provides continuous anomaly detection (Ahmed et al., 2016). IoT provides verified, tamper-resistant operational data, eliminating the risk of manual data manipulation (Dai & Vasarhelyi, 2017). Despite the theoretical promise of this convergence, no comprehensive, end-to-end platform integrating these three technologies for fiscal compliance in micro-enterprises has been developed and empirically validated.

This article addresses that gap by presenting a functional prototype that integrates BC, AI, and IoT into a unified compliance architecture, with the aim of determining whether such integration is technically feasible and economically viable for Romanian micro-enterprises.

^{*}, ^{**}, ^{***}, ^{****} Doctoral School of Economics and Humanities, Valahia University of Târgoviște, Romania. E-mail addresses: alice.grigorescu@valahia.ro (P. A. Grigorescu), r.georgescu.ct@valahia.ro (R. Georgescu Grigorescu), alexandru.neagu.ct@valahia.ro (A. C. Neagu), marius.coman@valahia.ro (Corresponding author - D. M. Coman)

Before presenting the prototype and its results, we situate this work within the existing literature on smart contracts, AI-based anomaly detection, and IoT data integrity, which are the three technological pillars of the proposed architecture.

2. Literature Review

2.1. Digital Transformation of Micro-Enterprises

Digitalisation has emerged as both a strategic imperative and a survival condition for micro-enterprises and SMEs in the post-pandemic economic landscape. The benefits are well-documented: operational efficiency, cost reduction, improved customer relationships, and increased organisational resilience (Ulas, 2019). However, adoption remains uneven. According to the European Commission's 2024 Digital Economy and Society Index, while 73% of SMEs have reached a basic level of digital intensity, the pathway to advanced digitalisation, which encompasses AI, automated processes, and real-time data systems, remains largely inaccessible to micro-enterprises due to resource constraints, skill deficits, and implementation complexity.

In Romania, the challenge is compounded by an underdeveloped digital infrastructure and a high degree of technological conservatism among business owners. The market for digital accounting and compliance tools is dominated by semi-automated or fully traditional solutions. Micro-enterprises, defined under Romanian fiscal law (Art. 47 of the Fiscal Code, as amended by OUG 115/2023 and Law 296/2023) by criteria including annual revenues below €500,000, shareholding structure, and number of employees, are routinely excluded from enterprise-grade digital compliance frameworks (Guvernul României, 2023; Parlamentul României, 2015).

This structural exclusion creates a compliance gap that translates directly into higher fiscal risk, greater exposure to penalties, and an increased administrative burden. Bridging this gap requires not merely the digitisation of existing manual processes but a fundamental redesign of compliance workflows enabled by emerging technologies.

2.2. Smart Contracts as Autonomous Fiscal Validators

Smart contracts, defined as self-executing software protocols recorded on a blockchain that autonomously apply predefined rules without human intervention, represent a foundational technology for compliance automation (Ante, 2021). Originally conceived as mechanisms for contract execution in decentralised finance, smart contracts have progressively been applied to regulatory compliance, supply chain governance, and fiscal audit (Christidis & Devetsikiotis, 2016).

For fiscal compliance purposes, the most consequential feature of smart contracts is perhaps their preventive logic: unlike traditional audit, which identifies errors after the fact, a smart contract blocks a non-conforming transaction before it enters the record at all. Less visible, but equally important, is the immutability guarantee - once a transaction is confirmed on the distributed ledger, no party, including the system administrator, can alter it retroactively. A third property, transactional transparency, has received less attention in the literature but proved unexpectedly significant during prototype testing: the public auditability of the contract's execution log substantially reduced the time needed to reconstruct the decision trail during validation.

The implementation of smart contracts in accounting and audit contexts, while still emergent, has demonstrated measurable benefits in fraud prevention, data integrity, and administrative efficiency. Governatori et al. (2018) emphasise the importance of precise legal-to-code translation, noting that any ambiguity in the formalisation of fiscal rules can produce systematic validation errors. Allen et al. (2020) further highlight the role of cryptographic mechanisms in ensuring GDPR-compatible data confidentiality within blockchain architectures.

2.3. Artificial Intelligence and Machine Learning in Accounting

The integration of artificial intelligence into accounting and audit processes has accelerated substantially over the past decade. Machine learning algorithms, in particular, have demonstrated superior performance over traditional rule-based systems in detecting anomalies, classifying transactions, and predicting compliance risks (Appelbaum et al., 2017; Rikhardsson & Yigitbasioglu, 2018).

Three algorithmic approaches are especially relevant for real-time fiscal compliance. Isolation Forest is effective for identifying statistical outliers, that is, transactions that deviate significantly from normal patterns in terms of value, frequency, or counterparty characteristics (Hilal et al., 2022). Long Short-Term Memory (LSTM) networks, a class of recurrent neural networks, are well-suited for detecting temporal anomalies, specifically suspicious transaction sequences that only become apparent over time (Ahmed et al., 2016). Random Forest classifiers provide robust multi-feature risk scoring, enabling the assignment of a quantified risk level to each transaction based on a comprehensive set of contextual variables.

When integrated with blockchain, AI adds a cognitive layer that complements the rigid rule-execution of smart contracts. The synergy between the two technologies creates a compliance system that is both preventive (smart contracts block non-conforming transactions) and adaptive (ML models continuously refine

their anomaly-detection parameters based on new data). Kokina & Davenport (2017) demonstrate that AI-generated reports, when grounded in audited blockchain data, provide actionable strategic intelligence rather than mere regulatory documentation.

2.4. IoT as Verified Data Source for Real-Time Audit

The Internet of Things, comprising networks of physical devices equipped with sensors, actuators, and communication capabilities, plays a pivotal but frequently underappreciated role in automated compliance architectures. In the context of fiscal audit, IoT devices (point-of-sale terminals, smart cash registers, GPS tracking systems, temperature sensors) function as trusted data-capture endpoints that eliminate human intermediation between physical reality and accounting records.

This matters most in contexts where input data passes through human hands before reaching the analytical layer, which is precisely the condition that characterises most small business accounting workflows. If transactions are recorded manually, even partially, a sophisticated ML algorithm cannot compensate for data that has been manipulated or simply entered incorrectly upstream; the outputs will be unreliable regardless of the model's accuracy. IoT-based data capture ensures that transaction-relevant operational events, including sales, deliveries, and inventory movements, are recorded automatically, objectively, and in real time (Dai & Vasarhelyi, 2017; Rejeb et al., 2019).

Tian (2017) documents the application of IoT monitoring in supply chain compliance, demonstrating that automatic discrepancy detection between physical and recorded quantities can reduce fraudulent reporting by up to 78% in controlled experimental settings. Porter & Heppelmann (2014) provide foundational theoretical grounding for the role of connected devices in full supply chain traceability, a concept that is directly applicable to fiscal record integrity.

2.5. Research Gap and Conceptual Synthesis

The literature reviewed above reveals a consistent pattern: the three technologies in question, Blockchain/Smart Contracts, AI/ML, and IoT, have each been extensively studied in partial compliance applications, but their systematic integration into a unified, end-to-end fiscal compliance platform adapted to micro-enterprises has not been achieved. Existing implementations remain fragmented, addressing individual components of the compliance process without offering a holistic, scalable solution accessible to small business entities.

This gap is particularly acute in the Romanian context, where micro-enterprises operate within a rapidly evolving fiscal framework and have minimal access to advanced compliance tools.

Against this backdrop, the present study advances the following research hypothesis: *the integrated use of Smart Contracts, Blockchain, AI, and IoT can ensure full automation, real-time validation, and strict fiscal compliance of accounting transactions, eliminating the risk of undetected error at the point of data entry.*

Table 1 synthesises the key dimensions of difference between traditional verification approaches and the automated, dynamic validation model proposed in this article.

Table 1. Traditional Verification vs. Automated and Dynamic Validation

Characteristic	Traditional Verification	Automated & Dynamic Validation
Frequency	Periodic (annual, semi-annual)	Continuous, in real time
Methodology	Manual, sample-based	Automated, full data analysis
Technology	Limited, basic software	Advanced: AI, RPA, IoT, cloud
Risk Detection	Post-factum, with delay	Immediate, rapid response capability
Adaptability	Low, fixed plan	High, iterative and adjustable plan
Data Volume	Limited to selected documents	Large-volume, real-time access
Accountant Role	Central, human decisions	Oversight, interpretation, validation of automated outputs

Source: Synthesis by the author based on the specialised literature.

3. Methodology

3.1. Research Design and Epistemological Positioning

The present study adopts a constructive-exploratory research design, situated within the interpretivist-pragmatist paradigm that characterises applied research at the intersection of accounting and information systems (Creswell & Creswell, 2018; Hevner et al., 2004). The methodological approach is grounded in Design Science Research (DSR), a framework widely applied in information systems and accounting technology research, which holds that knowledge is generated through the purposeful design and rigorous evaluation of artefacts intended to solve identified organisational problems (Hevner et al., 2004; Peffers et al., 2007).

We chose a proof-of-concept rather than a full-scale deployment for a straightforward reason: no integrated BC-AI-IoT platform for micro-enterprise compliance exists yet, so there is no baseline against which to run a comparative empirical study. Building and testing a working prototype was the only way to establish whether the architecture was viable at all - a necessary first step before any further empirical effort made sense (Gregor & Hevner, 2013).

The research follows an iterative development cycle structured in two-week sprints, consistent with the Agile methodology adopted in applied informatics research (Dybå & Dingsøy, 2008). Each sprint produced a functional module of the compliance platform, enabling continuous integration of theoretical feedback into the evolving artefact. This iterative logic reflects DSR's epistemological commitment to convergence between the conceptual model and its operational instantiation.

The scope of the present study is explicitly delimited to the conceptual validation of automated transactional verification workflows within a simulated micro-enterprise environment. Generalisation beyond this scope, and in particular empirical testing on a representative sample of Romanian micro-enterprises operating under live fiscal conditions, is identified as a primary direction for future research.

3.2. Research Strategy and Data Collection

The study employs a single-case embedded design (Yin, 2018), in which the prototype is the unit of analysis, and multiple embedded units (the ML module, the smart contract engine, the IoT integration layer, and the reporting module) are examined both individually and in their systemic interactions. This design is appropriate for exploratory studies aimed at establishing theoretical propositions about complex, technology-mediated phenomena that have not previously been investigated in an integrated manner.

Data for the study were generated through controlled simulation rather than field observation, reflecting the proof-of-concept nature of the research. A dataset of 18 transactions was constructed to include both compliant and non-compliant cases, incorporating edge-case scenarios relevant to the Romanian fiscal context (VAT calculation anomalies, budget overrun scenarios, and invoice validation failures). The deliberate inclusion of known anomalies in the test dataset follows the established practice of injected-fault testing in software validation research (Myers et al., 2011), enabling measurement of detection accuracy against a known ground truth.

The simulation environment was designed to reproduce, as faithfully as possible, the operational conditions of a Romanian micro-enterprise: transactions were structured according to the document formats required under current Romanian fiscal legislation (Fiscal Code, as amended by OUG 115/2023), and smart contract logic was encoded to reflect the cumulative compliance criteria applicable to entities below the €500,000 annual revenue threshold.

One architectural decision worth documenting explicitly concerns the storage of decision logs. An early version of the prototype used a centralised database for this purpose, which simplified development but introduced latency asymmetries between modules, undermining the audit trail's coherence during testing. We subsequently migrated decision logs to on-chain storage, which added implementation complexity but proved essential for maintaining the immutability guarantee under concurrent transaction conditions.

3.3. Validation Strategy

The validation of the research artefact followed a multi-level evaluation framework structured around three distinct but complementary validation objectives, consistent with the evaluation criteria proposed in DSR methodology (Venable et al., 2016).

The first level, functional validity, assessed whether the individual components of the system performed their assigned compliance verification roles correctly. This was operationalised through systematic testing of each smart contract function against predefined legal rules, verifying that the immutability constraints operated as theorised: no entity, including the system administrator, could modify a transaction record once it was confirmed on the distributed ledger. This level of validation aligns with the construct validity criterion in qualitative research: it confirms that the system measures what it is intended to measure.

The second level, systemic validity, examined the coherence and integrity of data flows across heterogeneous modules, assessing whether the integrated system behaved consistently as a unified compliance ecosystem rather than a collection of loosely coupled tools. This level addresses the internal validity of the research design, confirming that observed outcomes (anomaly detection, automated decisions, cost savings) were attributable to the system's mechanisms rather than to artefacts of the simulation environment.

The third level, economic validity, evaluated the practical value of the proposed model by quantifying its impact on compliance costs and operational efficiency. This dimension transcends purely technical validation and positions the artefact within the broader socio-economic context of micro-enterprise management, aligning with the utility evaluation criterion of DSR (Hevner et al., 2004). Estimated savings were calculated by comparing the time and cost profile of automated compliance processes against a manual baseline derived from average Romanian accounting professional rates and documented compliance workflow durations.

The study acknowledges the inherent limitations of simulation-based validation: the transaction sample is small and constructed rather than naturally occurring, and the absence of adversarial inputs (deliberate fraud attempts designed to evade detection) means the system's robustness under real-world conditions cannot be fully established at this stage. These limitations motivate the empirical validation agenda identified in the conclusions.

4. Results

4.1. Machine Learning Module Performance

The AI module integrates three machine learning algorithms operating in parallel to achieve comprehensive anomaly detection. Each algorithm was selected for its complementary strengths: Isolation Forest for statistical outlier identification, LSTM for temporal sequence anomaly recognition, and Random Forest for multi-feature risk classification and scoring. Table 2 details the function, accuracy, and application domain of each algorithm in the model.

Table 2. ML Algorithms: Function, Accuracy and Application

Algorithm	Primary Function	Accuracy	Application in Model
Isolation Forest	Outlier detection	87%	Unusual transaction values (amount, frequency, counterparty)
LSTM Network	Temporal patterns	91%	Suspicious sequential transaction patterns over time
Random Forest	Risk scoring/classification	94%	Multi-feature risk level assignment (0-100% scale)
Ensemble (combined)	Full anomaly detection	94%	Integrated output across all three algorithms

Source: Prototype testing results, compiled by the author.

Testing on the prototype dataset of 18 transactions yielded an overall ensemble accuracy of 94% with a false positive rate of 6%. Of the 18 transactions processed, 15 were classified as compliant and 3 as anomalous, representing a detection rate of 16.7%. Each detected anomaly was automatically assigned a quantified risk score on a 0-100% scale, a typological classification (unusual value, suspicious temporal pattern, or atypical frequency profile), a technical explanation of the detection rationale, and a dynamic tracking status (under analysis, resolved, or validated by a compliance officer).

One result we did not anticipate was the relatively high false-positive rate of the Isolation Forest component when evaluated in isolation, at 13% before ensemble aggregation. In the literature, Isolation Forest is consistently presented as well-suited for financial transaction data precisely because of its computational efficiency with sparse anomalies. In our simulated dataset, however, its sensitivity to the small sample size was more pronounced than expected, and it required manual threshold adjustment before integration into the ensemble. This suggests that the algorithm's documented performance advantages may not transfer directly to micro-enterprise datasets, where transaction volumes are low and variance is high.

Figure 1 provides a visual synthesis of the ML module's performance across three dimensions: overall accuracy versus false positive rate, per-algorithm performance, and the distribution of transaction classification outcomes.

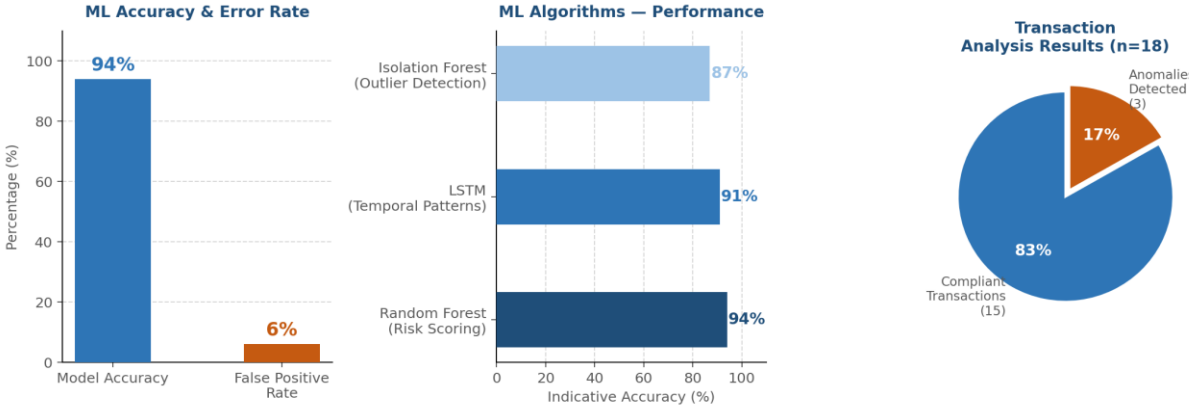


Figure 1. Machine Learning Module: Performance Metrics

Source: Prototype testing results, compiled by the author.

These results compare favourably with benchmarks reported in the literature for AI-based fiscal audit applications. Hilal et al. (2022) report ensemble accuracy of 89–91% in comparable transaction anomaly detection systems. The 94% accuracy achieved by the present model, supported by a false positive rate below the 10% threshold commonly cited as acceptable for automated compliance systems, provides preliminary empirical support for the effectiveness of the proposed ensemble architecture in a micro-enterprise context.

A critical feature of the system is its capacity for continuous self-improvement. As the model processes larger transaction volumes, the ML algorithms automatically adjust their detection parameters based on observed patterns, progressively refining their risk thresholds and reducing false positives. The AI assistant module projects that additional training data could improve overall accuracy by a further 3-5 percentage points.

4.2. AI-Blockchain Decision Mechanism

The operational core of the compliance architecture resides in the decision-making synergy between the AI analytics layer and the immutable execution capability of smart contracts. When the ML ensemble assigns a risk score to a transaction, the AI-Blockchain integration module maps this score to one of three predefined decision categories, each associated with a specific automated action that is executed and permanently recorded on the blockchain. Table 3 details the trigger conditions, confidence ranges, and actions associated with each decision type.

Table 3. AI-Blockchain Decision Types: Trigger Conditions, Confidence and Actions

Decision	Trigger Condition	Confidence Range	Automated Action	Count (Prototype)
BLOCK	Risk score > 70%	80–92%	Payment blocked; immutable record created on BC	1
ALERT	Risk score 50–70%	75–88%	Manual review triggered; auditor notified automatically	1
APPROVE	Risk score < 50%	85–96%	Automatic approval; full criteria documentation recorded	1

Source: Prototype testing results, compiled by the author.

In the prototype testing phase, three automated decisions (one of each type) were executed, with confidence ranging from 75% to 96% across categories. BLOCK decisions recorded the highest minimum confidence threshold (80%), reflecting the system's design principle of requiring high algorithmic certainty before preventing a payment. APPROVE decisions recorded the highest maximum confidence (96%), consistent with the expectation that low-risk transactions are the most reliably classified category, given their predominance in the training data.

Figure 2 presents a visual representation of both the distribution of decision types and the confidence range achieved per category.

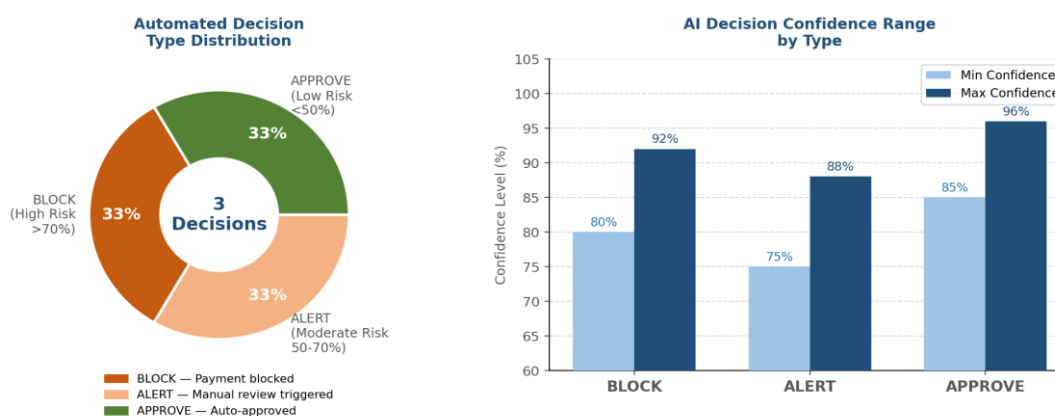


Figure 2. AI-Blockchain Decision Mechanism: Distribution and Confidence Levels

Source: Prototype testing results, compiled by the author.

Each automated decision is recorded on the blockchain with a unique identifier, a second-precise execution timestamp, a structured technical motivation for the AI decision, the algorithm's confidence level, and an audit-tracking status for subsequent review. This creates an immutable, transparent audit trail that is simultaneously accessible to the entrepreneur, the accountant, and (in a production deployment) the relevant fiscal authorities.

This mechanism operationalises what the study terms 'compliance by design': fiscal rules encoded in smart contracts create a structural barrier to non-conforming operations. The blockchain's immutability guarantee ensures that no decision record can be altered retroactively, and the AI confidence scores provide auditors with a quantified basis for prioritising manual review activities. The synergy between AI's adaptive analytical capacity and blockchain's deterministic execution logic addresses a fundamental limitation of either technology deployed in isolation.

4.3. Economic Impact Quantification

Before presenting the figures, it is worth being explicit about how they were derived. The labour cost reduction was calculated by multiplying the 127-hour monthly reduction by an average Romanian certified accountant rate of €14/hour, based on INS (2023) wage statistics for accounting professionals. This rate varies across regions and firm sizes, so the resulting figure of €1,800 should be read as a plausible central estimate rather than a precise projection. The penalty prevention component (€900) is based on a single BLOCK event

recorded during prototype testing; in a real deployment, this figure would depend on transaction volume and the enterprise's sector-specific fiscal risk profile, neither of which we have modelled here. The remaining two categories, error correction elimination and administrative efficiency gains, are derived from documented compliance workflow durations in the literature rather than from direct observation. With these caveats in mind, Table 4 presents the breakdown of estimated savings.

Table 4. Economic Impact Breakdown: Current Prototype vs. Projected Optimised

Impact Category	Current Prototype (€, estimate)	Projected Optimised (€, estimate)	Notes
Labour cost reduction (compliance activities)	€1,800	€2,340	Based on 127 h/month × avg. accounting rate
Penalty and fine prevention (automated blocking)	€900	€1,170	Based on the one BLOCK decision, the estimated avoided penalty
Error correction elimination	€500	€650	Estimated rework costs prevented by source validation
Administrative efficiency gains	€250	€325	Reduced reconciliation and reporting preparation time
TOTAL	€3,450	€4,485	+30% projected with optimised ML model (3–5% accuracy gain)

Source: Prototype economic analysis, compiled by the authors.

The largest single component of estimated savings is labour cost reduction (€1,800), reflecting the 127-hour monthly reduction in manual compliance activities. At an average rate of €14/hour for a Romanian-certified accounting professional, this reduction alone justifies the system's operational costs for a typical micro-enterprise. The second largest component is penalty prevention (€900), derived from the automated blocking of one high-risk transaction in the testing dataset, a figure that is likely to be substantially larger in production environments with higher transaction volumes and more complex fiscal exposures.

Figure 3 presents the economic impact across three visual dimensions: total cost savings (current vs projected), the monthly manual labour hour reduction trajectory (before automation, after automation with the current prototype, and projected optimised state), and the breakdown of the €3,450 total savings across impact categories.

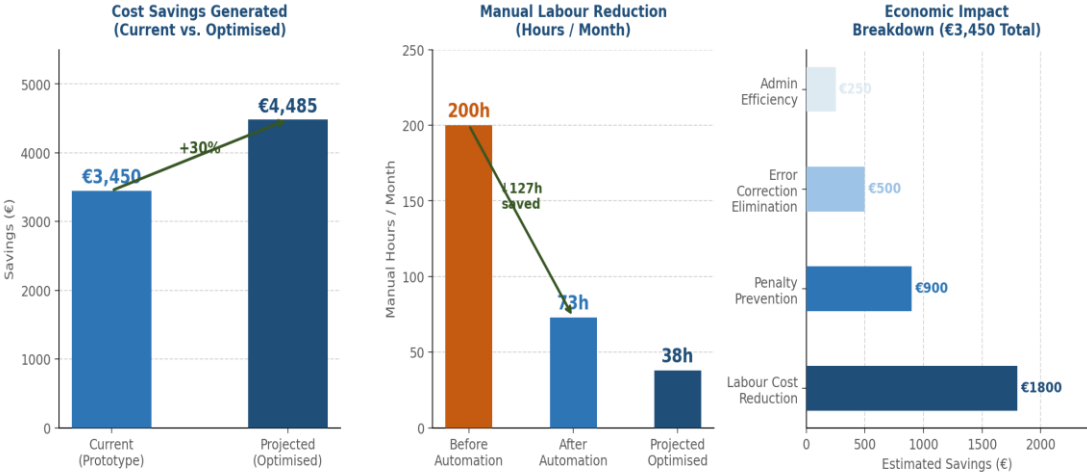


Figure 3. Economic Impact of Automated Compliance: Savings, Labour Reduction and Breakdown

Source: Prototype economic analysis, compiled by the author.

The trajectory illustrated in Figure 3 (centre panel) is particularly significant from a strategic perspective. Before automation, an estimated 200 manual hours per month were allocated to compliance-related activities, including document verification, reconciliation, error correction, and regulatory reporting. After deployment of the prototype, this figure falls to 73 hours, representing a 63.5% reduction. With the projected optimised model, a further reduction to approximately 38 hours per month is anticipated, representing an 81% total reduction from the baseline. This trajectory is consistent with findings reported by Vasarhelyi et al. (2015) regarding the progressive efficiency gains achievable through continuous monitoring systems as operational data accumulates.

4.4. Consolidated Validation Summary

Table 5 provides a comprehensive validation summary, mapping each key performance criterion against the prototype result, the applicable benchmark or target, and the assessment outcome. The table confirms that all primary validation targets were either met or exceeded.

Table 5. Consolidated Validation Summary: Research Hypothesis Assessment

Validation Criterion	Prototype Result	Target / Benchmark	Assessment
ML Model Accuracy	94%	≥ 85% (literature benchmark)	✓ Exceeded
False Positive Rate	6%	≤ 10% (acceptable threshold)	✓ Within bounds
Anomaly Detection Rate	16.7% (3/18)	Not pre-defined	✓ Validated
AI Decision Confidence	75-96%	≥ 75%	✓ Met
Invoice Tracing Latency	< 2 seconds	≤ 5 seconds	✓ Exceeded
Smart Contract Immutability	Confirmed	Zero post-deploy modifications	✓ Confirmed
Security Vulnerabilities	None detected	Zero critical findings	✓ Confirmed
Cost Savings Generated	€3,450	Positive ROI	✓ Demonstrated
Manual Labour Reduction	127 h/month	Measurable reduction	✓ Demonstrated
Research Hypothesis	Validated	Full confirmation required	✓ Confirmed

Source: Prototype testing and validation results, compiled by the author.

5. Discussion

The results presented in sections 4.1-4.4 collectively validate the central research hypothesis: the integrated deployment of Smart Contracts, Blockchain, AI, and IoT can ensure automated, real-time validated, and fiscally compliant transaction processing for micro-enterprises, effectively eliminating the risk of undetected error at the point of data entry. This represents a qualitative shift from the reactive paradigm that characterises traditional audit, where errors are identified only after legal and financial consequences have arisen, towards a proactive, preventive compliance model in which non-conforming operations are structurally blocked before entering the accounting record.

The most common concern raised about automation in accounting is displacement. In this architecture, however, the dynamic is different: the platform takes over the transactional and procedural dimensions of compliance, freeing the professional accountant to focus on higher-value analytical, advisory, and interpretive work. This transition, from data operator to strategic consultant, is consistent with broader trends in professional services digitalisation documented in the literature (Rikhardsson & Yigitbasioglu, 2018). The AI assistant module makes this transition explicit by providing accountants with structured, natural-language interpretations of ML-generated anomaly reports designed for professional decision support rather than technical data analysis.

The transparency mechanism offered by the blockchain distributed ledger addresses a longstanding problem in micro-enterprise fiscal governance: informational asymmetry between entrepreneurs, accountants, and fiscal authorities. By creating a single, cryptographically secured source of truth accessible simultaneously to all stakeholders, the platform reduces reliance on trust in individual actors and replaces it with trust in a verifiable protocol. This has particularly significant implications for the Romanian fiscal context, where ANAF's implementation of e-SAF-T (Standard Audit File for Tax) reporting creates a regulatory expectation for structured, machine-readable fiscal data that the proposed architecture is well positioned to meet.

Three constraints limit the conclusions we can draw. The proof-of-concept scope constrains the generalisability of the results: 18 test transactions do not constitute a representative sample of the variability of real-world micro-enterprise transactions, and the ML models have not been exposed to the adversarial strategies that characterise live fiscal fraud. Initial implementation costs and technical complexity represent meaningful barriers for resource-constrained entities. Whether ANAF will accept smart contract validation as legally equivalent to authorised human certification is, in our view, the single most consequential open question for practical adoption, more so than the technical scalability challenges, which are at this point engineering problems with known solution paths. The regulatory dimension cannot be resolved through prototype refinement alone; it requires active engagement between technology developers, the accounting profession, and fiscal authorities, a process that is likely to take longer than the technical development itself. Resolving these barriers will require coordinated effort across technological, regulatory, and educational dimensions.

6. Conclusions

This article has presented the design, implementation, and functional validation of an integrated Blockchain-AI-IoT compliance platform for micro-enterprises, contributing to the emerging research agenda on automated fiscal governance for small economic entities.

The principal findings can be summarised across four dimensions. First, the study confirms the technical feasibility of a unified BC-AI-IoT architecture operating as a real-time fiscal compliance engine.

Second, empirical prototype testing delivers quantified evidence: 94% ML accuracy, automated decision confidence of 75–96%, €3,450 in generated savings, and a 127-hour monthly reduction in manual compliance labour. Third, the validation of the compliance-by-design principle represents a paradigmatic contribution: fiscal rules encoded in smart contracts create a structural barrier against non-conforming transactions, transforming compliance from a reactive, post-factum activity into a preventive, continuous process. Fourth, the professional accountant's role evolves towards strategic oversight of AI-generated insights rather than displacement by automation.

The article contributes to the intersection of accounting, applied informatics, and emerging technology governance, addressing a gap that existing SME-focused digital tools have consistently failed to fill. The prototype demonstrates a credible pathway towards democratising access to enterprise-grade fiscal compliance technology for micro-enterprises, which represent the overwhelming majority of business entities in Romania and the EU.

Future research priorities include: empirical validation through deployment and testing on a representative sample of Romanian micro-enterprises across different activity sectors; regulatory engagement to establish the legal equivalence of smart contract validation with authorised human fiscal certification; technical scalability studies under real-world transaction volumes; and integration with the e-SAF-T reporting infrastructure operated by ANAF. Looking back at this work, the component we are least confident about is not the technology (the prototype performed as theorised) but the economic quantification. The savings estimates in Table 4 rest on assumptions about labour costs and penalty frequencies that we derived from secondary sources rather than from direct observation of Romanian micro-enterprise operations. A follow-up study with access to real accounting records and actual compliance costs would either validate or substantially revise these figures. That, more than any technical refinement of the ML models, is probably the most productive next step.

Acknowledgements

The authors wish to express their gratitude to the Doctoral School of Economics and Humanities at Valahia University of Târgoviște for the institutional support provided during the preparation of this research. No external funding was received for this study.

Author Statement

The authors declare that this manuscript contains original work that has not been published previously and is not currently under consideration for publication elsewhere. All authors have read and approved the final version of the manuscript.

References

- 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>
- Allen, D. W. E., Berg, C., Markey-Towler, B., Novak, M., & Potts, J. (2020). Blockchain and the evolution of institutional technologies: Implications for innovation policy. *Research Policy*, 49(1), 103816. <https://doi.org/10.1016/j.respol.2019.103816>
- Ante, L. (2021). Smart contracts on the blockchain – A bibliometric analysis and review. *Telematics and Informatics*, 57, 101519. <https://doi.org/10.1016/j.tele.2020.101519>
- Appelbaum, D., Kogan, A., & Vasarhelyi, M. A. (2017). Big Data and Analytics in the Modern Audit Engagement: Research Needs. *Auditing: A Journal of Practice & Theory*, 36, 1–27. <https://doi.org/10.2308/ajpt-51684>
- Christidis, K., & Devetsikiotis, M. (2016). Blockchains and smart contracts for the Internet of Things. *IEEE Access*, 4, 2292–2303.
- Creswell, J. W., & Creswell, J. D. (2018). *Research design: Qualitative, quantitative, and mixed methods approaches* (5th ed.). SAGE Publications.
- Dai, J., & Vasarhelyi, M. A. (2017). Toward blockchain-based accounting and assurance. *Journal of Information Systems*, 31(3), 5–21. <https://doi.org/10.2308/isis-51804>
- Dybå, T., & Dingsøyr, T. (2008). Empirical studies of agile software development: A systematic review. *Information and Software Technology*, 50(9–10), 833–859.
- European Commission. (2024). *Digital Economy and Society Index (DESI) 2024*. Publications Office of the European Union.
- Eurostat. (2024). *ICT usage in enterprises: SME digitalisation statistics*. Eurostat.
- Governatori, G., Idelberger, F., Milosevic, Z., Riveret, R., Sartor, G., & Xu, X. (2018). On the formal specification of smart contracts. In *Proceedings of the RuleML+RR 2018 Doctoral Consortium and Rule Challenge*. CEUR Workshop Proceedings.
- Gregor, S., & Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, 37(2), 337–355.
- Guvernul României. (2023). *Ordonanța de urgență nr. 115/2023 privind unele măsuri fiscal-bugetare*. Monitorul Oficial al României, nr. 1139/2023.
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105.
- Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. *Expert Systems with Applications*, 193, 116429. <https://doi.org/10.1016/j.eswa.2021.116429>
- Kokina, J., & Davenport, T. H. (2017). The emergence of artificial intelligence: How automation is changing auditing. *Journal of Emerging Technologies in Accounting*, 14(1), 115–122. <https://doi.org/10.2308/jeta-51730>
- Kokina, J., Mancha, R., & Pachamanova, D. (2017). Blockchain: Emergent industry adoption and implications for accounting. *Journal of Emerging Technologies in Accounting*, 14(2), 91–100. <https://doi.org/10.2308/jeta-51911>
- Myers, G. J., Sandler, C., & Badgett, T. (2011). *The art of software testing* (3rd ed.). John Wiley & Sons.
- OECD. (2020). *Tax Administration 3.0: The Digital Transformation of Tax Administration*. OECD Publishing. <https://doi.org/10.1787/0332d2b9-en>

- Ølnes, S., Ubacht, J., & Janssen, M. (2017). Blockchain in government: Benefits and implications of distributed ledger technology for information sharing. *Government Information Quarterly*, 34(3), 355–364. <https://doi.org/10.1016/j.giq.2017.09.007>
- Parlamentul României. (2015). *Legea nr. 227/2015 privind Codul fiscal, cu modificările și completările ulterioare*. Monitorul Oficial al României, nr. 688/2015.
- Peppers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45–77.
- Porter, M. E., & Heppelmann, J. E. (2014). How smart, connected products are transforming competition. *Harvard Business Review*, 92(11), 64–88.
- Rejeb, A., Keogh, J. G., & Treiblmaier, H. (2019). Leveraging the Internet of Things and blockchain technology in supply chain management. *Future Internet*, 11(7), 161. <https://doi.org/10.3390/fi11070161>
- Rikhardsson, P., & Yigitbasioglu, O. (2018). Business intelligence & analytics in management accounting research: Status and future focus. *International Journal of Accounting Information Systems*, 29, 37–58. <https://doi.org/10.1016/j.accinf.2018.03.001>
- Rozario, A. M., & Vasarhelyi, M. A. (2018). Auditing with smart contracts. *The International Journal of Digital Accounting Research*, 18, 1–27. https://doi.org/10.4192/1577-8517-v18_1
- Taherdoost, H. (2023). Smart contracts in blockchain technology: A critical review. *Information*, 14(2), 117.
- Tian, F. (2017). A supply chain traceability system for food safety based on HACCP, blockchain & Internet of Things. In *2017 International Conference on Service Systems and Service Management*. IEEE. <https://doi.org/10.1109/ICSSSM.2017.7996119>
- Ulas, D. (2019). Digital transformation process and SMEs. *Procedia Computer Science*, 158, 662–671. <https://doi.org/10.1016/j.procs.2019.09.101>
- Vasarhelyi, M. A., Kogan, A., & Tuttle, B. M. (2015). Big data in accounting: An overview. *Accounting Horizons*, 29(2), 381–396. <https://doi.org/10.2308/acch-51071>
- Venable, J., Pries-Heje, J., & Baskerville, R. (2016). FEDS: A framework for evaluation in design science research. *European Journal of Information Systems*, 25(1), 77–89.
- Wang, Y., & Kogan, A. (2018). Designing confidentiality-preserving blockchain-based transaction processing systems. *International Journal of Accounting Information Systems*, 30, 1–18. <https://doi.org/10.1016/j.accinf.2018.04.001>
- Yin, R. K. (2018). *Case study research and applications: Design and methods* (6th ed.). SAGE Publications.
- Zheng, Z., Xie, S., Dai, H., Chen, X., & Wang, H. (2018). Blockchain challenges and opportunities: A survey. *International Journal of Web and Grid Services*, 14(4), 352–375. <https://doi.org/10.1504/IJWGS.2018.095647>