

Integrating Artificial Intelligence into Blockchain-Based Systemic Report Verification: Towards an AI Protocol for Systemic Analysis

Original

Integrating Artificial Intelligence into Blockchain-Based Systemic Report Verification: Towards an AI Protocol for Systemic Analysis / Liboni, M., Mucchetti, F., Peruccio, P.P.. - ELETTRONICO. - RSD14:(In corso di stampa), pp. 1-22. (ARCS OF IMPACT Toronto, ON (CAN) October 3 - 21, 2025).

Availability:

This version is available at: 11583/3006262 since: 2026-01-05T09:31:08Z

Publisher:

Systemic Design Association

Published

DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)



ARCS OF IMPACT

**RELATING SYSTEMS THINKING & DESIGN
RSD14 | OCTOBER 2025**

Integrating Artificial Intelligence into Blockchain-Based Systemic Report Verification: Towards an AI Protocol for Systemic Analysis

Martina Liboni, Francesca Mucchetti, and Pier Paolo Peruccio

Integrity and transparency of systemic reports are fundamental elements for ensuring the reliability of sustainability assessments, circular economy evaluations, and preventing instances of project greenwashing. In the research presented at RSD13 ("Applying Blockchain to Systemic Design: Ensuring Project Authenticity via Quantitative Report Verification"), it was demonstrated through case studies how blockchain technologies—utilizing notarization, smart contracts, and NFTs—can provide an immutable verification system for systemic reports.

Building on this premise, this study introduces an evolution in the verification protocol by integrating Artificial Intelligence (AI), aimed at enhancing systemic design methodologies through the synergy of AI and blockchain verification systems.

The proposed protocol leverages Machine Learning (ML) algorithms to identify anomalies, inconsistencies, and potential distortions in systemic reports, automating the verification process through real-time data analysis and pattern recognition via Optical Character Recognition (OCR). With advancements in computational technologies towards "Light Large Language Model (LLM)" generative AI models, which are less complex and fully open source compared to traditional LLMs, the project aims to secure the verification process within a completely safe environment. This approach allows training neural models exclusively with specific data, avoiding the need to upload sensitive information online, thereby ensuring greater data security. This new frontier enables harnessing all benefits of AI without compromising cybersecurity, a critical concern for safeguarding sensitive information.

This research aims to develop a hybrid AI-blockchain model for report validation, where AI evaluates data integrity and blockchain ensures decentralized certification. It implements Natural Language Processing (NLP) techniques for analysing textual components of reports, ensuring compliance with major sustainability frameworks such as Global Reporting Initiative (GRI) and ISO14001, without being centralized under a single control.

The objective is to define a standardized AI protocol for systemic design verification applicable across various sectors, ensuring more accurate and reliable sustainability reporting.

Integrating AI into systemic design methodologies aims to introduce secondary verification processes complementary to blockchain, currently absent yet necessary to strengthen protection for designers and involved companies. Furthermore, it aims to offer a scalable, automated approach resistant to data validation fraud. The use of advanced AI tools not only optimizes process efficiency but also ensures greater compliance with sustainability standards, thereby enhancing stakeholder confidence and reducing the risk of misleading information dissemination.

This research contributes to the fields of systemic design, data verification, and digital trust, proposing an innovative AI-based framework to ensure the authenticity of systemic design documentation.

KEYWORDS: Sustainability documents, AI verification, Blockchain, Sustainability compliance, Digital trust, Data integrity

RSD: Methods & Methodology.

Context and Evolution of the Research Proposed to RSD13

Current literature demonstrates the effectiveness of blockchain technology in ensuring the immutability of systemic reports through notarization, smart contracts, and Non-Fungible Token (NFT). This has been highlighted in case studies concerning industrial circularity assessments, sustainability balances, and systemic reports (1,2), where major global certifiers have been examined, highlighting issues of centralization,

control, and associated costs (3). The aim of current research is to overcome the limitations of attestation solely based on decentralized certification, as this approach has revealed a significant challenge: the absence of predictive analytic tools capable of proactively identifying semantic inconsistencies or methodological discrepancies in certification preparatory documents (4). This gap is particularly evident in the context of Global Reporting Initiative (GRI) (5) and ISO14001 (6) frameworks, where 27% of discrepancies reported in 2024 concerned categorical classification errors that could not be detected through purely quantitative checks (7).

The primary objective of this research is to overcome this limitation through a hybrid AI-blockchain system that integrates:

- The use of advanced Optical Character Recognition (OCR) to automatically extract data from reports, with cross-validation comparing the extracted data with other sources (ILCD, US LCI Database, Ecoinvent, etc.), thus ensuring the reliability and consistency of the information (8).
- The application of Machine Learning (ML) algorithms to analyse the metrics in the reports, identifying anomalies, inconsistencies, and potential errors in time series and measurement units, to ensure that the data are accurate and logical (9).
- The use of Natural Language Processing (NLP) techniques to examine the textual content of reports, checking compliance with regulatory frameworks and detecting misleading or excessively promotional statements, to avoid the risk of greenwashing and ensure transparency (10).

Innovations

This research introduces an evolutionary paradigm in the validation of systemic reports, sustainability reports, and other sustainability-related documents such as Life Cycle Assessment (LCAs), which will hereafter be generically referred to as "document or reports". It is characterized by the synergistic integration of three key innovations: the adoption of open-source light LLM (Large Language Model) models for advanced textual

analysis, real-time synchronization between ML-based anomaly detection algorithms, OCR analysis, anti-greenwashing verification, and subsequent blockchain notarization.

The integration of open-source light LLM models represents a significant breakthrough compared to conventional architectures. Unlike traditional LLMs, which are often proprietary and centralized, light LLM models offer algorithmic transparency and the possibility of customization through fine-tuning on domain-specific datasets (e.g., sustainability reports, circularity assessments). This customization maximizes the precision of semantic analysis of document while minimizing the risk of exposure of sensitive data, as training occurs in an isolated and controlled environment (11). The architecture is designed to operate in compliance with privacy-by-design principles, ensuring that data are not transmitted to external servers or used for purposes other than report validation.

The synchronization between anomaly detection and blockchain notarization introduces a dynamic feedback mechanism in the verification process (12). ML algorithms, trained to identify statistical anomalies, logical inconsistencies, and potential biases in the data, generate a "reliability" score for each document analysed. This score directly influences the conditions for notarization. Reports with scores above 7.5 are automatically notarized, while those with lower scores undergo manual review or require further checks (13). This layered approach optimizes process efficiency by focusing human resources on the most problematic areas and reducing the risk of false positives or negatives.

The implementation of pattern recognition models for alignment with well-known standards such as GRI/ISO14001 automates a process that has traditionally been labour-intensive and subject to subjective interpretation (14). These models, based on NLP and text mining techniques, analyse the textual content of the document, identifying the presence of key indicators, compliance with guidelines, and completeness of information required by the standards, as well as the presence of reliable sources such as Ecoinvent (15). The system generates a detailed compliance result-a table designed to highlight any gaps or areas for improvement. This not only facilitates the preparation of reports in the event of a required review but also ensures greater consistency between analyses conducted at different times.

Methodology: Hybrid Architecture

The proposed model is defined as a "hybrid" architecture because it combines and integrates different technologies and disciplinary approaches, creating a cohesive workflow that enhances the effectiveness and reliability of the validation process. Its hybrid nature is rooted in systems thinking, an approach that considers each element as part of a complex interconnected system, where the various components work together dynamically. Systems thinking is the foundation that enables the development of a model that not only addresses individual problems in isolation but also recognises the importance of integrating them into a broader context, where data, technologies, and regulations interact to ensure accurate and multidimensional validation (16).

This model is structured into three distinct but closely interconnected layers: the Data Acquisition Layer, the AI Verification Layer, and the Blockchain Certification Layer. Each layer has a fundamental and complementary role in the process, but together they form a system that allows for the systemic and holistic validation of reports (Fig.1).




Figure 1. Summary table of the Hybrid AI-Blockchain Verification Workflow. Source: Authors.

Document Verification Methodology

The methodology adopted for the validation of systemic reports is structured as a sequential verification pipeline, consisting of four main phases: data extraction via advanced OCR with cross-validation, dimensional analysis via ML for metric consistency control, textual validation with NLP for alignment with regulatory frameworks, and the assessment of eligibility for blockchain certification of the documents. At the end of each phase, a score is assigned which is summed and divided across the three layers.

Data Extraction via Advanced OCR with Cross-Validation

The initial phase of the pipeline involves the automatic extraction of data contained in the reports, using advanced OCR technology optimized to handle a wide range of document formats (PDF, TIFF, JPEG) and languages. The implemented OCR leverages deep learning techniques to improve character recognition accuracy, minimizing transcription errors (17). A crucial aspect of this phase is the cross-validation of extracted data. The system compares the data obtained from OCR with other available information sources, such as public databases, previous corporate reports, or integrated IoT sensors. Any discrepancies are flagged for manual review. For example, if a report declares a certain level of CO₂ emissions, the system checks whether this value is consistent with data from environmental sensors installed at the production facility. Furthermore, the system applies syntactic and semantic validation rules to verify the formal correctness of the data (e.g., date formats, measurement units, permissible value ranges). The output of this phase is a structured and validated dataset, ready for dimensional analysis.

Dimensional Analysis via ML (Metric Consistency Control)

The second phase of the pipeline consists of the dimensional analysis of the extracted data, using ML algorithms to verify the consistency of metrics. This phase focuses on identifying statistical anomalies, logical inconsistencies, and potential biases in numerical data. The system employs time series analysis techniques to identify unusual trends or sudden changes in key metrics (e.g., energy consumption, waste production, recycling rate). Clustering algorithms are implemented to group similar reports and identify outliers that deviate significantly from the norm (18). Additionally, the system applies regression models to verify consistency between different metrics. For example, if a document reports an increase in production, the system checks whether this

increase is consistent with the declared energy consumption. Particular attention is paid to the dimensional consistency of measurement units. The system verifies that all metrics are expressed in the correct units and that conversions are performed correctly. Any anomalies or inconsistencies are flagged for manual review. The output of this phase is a further validated dataset, with a flag and a score for each metric indicating the level of reliability.

Textual Validation with NLP to Limit Greenwashing

The third phase of the pipeline is dedicated to the textual validation of reports, using NLP techniques to verify alignment with regulatory frameworks (e.g., GRI, ISO14001) and identify potential greenwashing. The system uses text classification models to identify the type of report (e.g., sustainability report, life cycle assessment, social balance sheet) and the thematic areas addressed (e.g., emissions, energy consumption, human rights). Named Entity Recognition (NER) algorithms (19) are implemented to identify key entities mentioned in the document (e.g., companies, products, standards, regulations). The system uses sentiment analysis models to assess the emotional tone of the text and identify any excessively positive or self-congratulatory statements that may indicate greenwashing. The system analyses the report text, identifying the presence of key indicators, the completeness of required information, and consistency with regulatory guidelines, as well as the presence of verified sources. The automated Compliance Check (20) generates a result on the document's compliance status, highlighting any gaps or areas for improvement. Additionally, the system checks for the presence of disclaimers and warnings that mitigate the company's liability in the event of inaccurate statements. The output of this phase is a validation table indicating the report's level of alignment with regulatory frameworks and the risk of greenwashing.

Blockchain Certification

The final phase of the pipeline involves the blockchain certification of reports, applied exclusively to datasets that have successfully passed the previous validation phases. The system uses smart contracts (21) to automate the certification process. Reports that pass all checks are automatically notarized and recorded on the blockchain. A unique

certificate is generated for each certified document, containing a hash of the document and the certification timestamp.

The generation of a certifying NFT is currently being implemented, which, thanks to the metadata associated with the NFT, could include information on the validation process, such as the reliability score generated by the dimensional analysis, the textual validation report, and the Compliance Check. Blockchain certification is differential, meaning that the level of validation, with the possibility of a corresponding badge on the NFT, reflects the degree of reliability of the document. This approach incentivizes companies to improve the quality of their reports and reduce the risk of greenwashing.

Innovative Components Compared to the Previous Verification Proposed in the RSD13 Article

Compared to the model outlined in the RSD13 paper (22), the current method for verifying systemic reports introduces a series of substantial innovations that enhance both analytical depth and the ability to ensure data transparency and reliability. The previously described approach was primarily based on the use of blockchain technology as a tool for notarization and automation via smart contracts, aiming to render report data immutable and traceable, and to certify the authenticity of documents through the issuance of NFTs, thereby protecting the interests of designers and companies involved. This model represented a significant advancement over centralized verification systems, offering a decentralized and tamper-resistant solution capable of directly involving stakeholders and adapting to various operational contexts. However, the RSD13 methodology mainly focused on the security and integrity of existing data within the documents, without systematically addressing the quality and consistency of the reported information or its adherence to relevant regulatory frameworks.

By contrast, the current method is characterized by its “hybrid” and systemic nature, integrating AI, ML, and blockchain technologies within a structured, multi-layered pipeline (23). This model goes beyond ensuring data immutability by introducing a verification sequence that begins with intelligent information extraction using advanced OCR and cross-validation with external sources, continues with automated dimensional analysis of metrics through machine learning algorithms, and culminates in textual

validation using NLP techniques aimed at both checking regulatory compliance and detecting potential greenwashing practices. The output of each phase is quantified and incorporated into a scoring system that determines the overall reliability level of the report.

The most significant innovation of this approach lies in its capacity to validate reports in a multidimensional manner: not only does it verify the internal coherence of numerical data and the formal accuracy of the information, but it also evaluates alignment with major international standards and detects possible narrative inconsistencies or overly self-congratulatory statements, which are often indicative of greenwashing. This level of analysis, entirely absent from the RSD13 model, enables each document to be assigned a transparent and well-substantiated reliability score, which is subsequently embedded into the blockchain certification process through differential NFTs. In this way, blockchain serves not only as an immutable ledger but also as a mechanism for incentivization: the quality and transparency of reports are rewarded with badges and certification levels that reflect the actual degree of compliance and trustworthiness achieved.

The integration of these two processes enables both a precise verification of document contents and mutual protection of the involved parties through blockchain-based certification, functioning as an immutable registry.

Computational Security

Computational security (24) constitutes a fundamental element of the proposed methodology, considering the sensitive nature of the data involved and the necessity of safeguarding the entire system against external threats. Security is built upon three main pillars, coherently integrated within the technological framework: the adoption of federated learning (25) for training artificial intelligence models, the use of Zero-Knowledge Proofs (26) for privacy-preserving result sharing, and the decentralized storage of models on IPFS with blockchain-based hash notarization (27).

In the context of AI model training—particularly for lightweight models based on LLMs—the federated learning strategy allows for the preservation of data confidentiality (28) by avoiding the need to aggregate data on a single central node. Each participating

node, typically representing an organisation or institution, locally trains its own model using proprietary data. The weights of these local models are then securely aggregated to form a shared global model, without the original data ever being transferred. To further enhance privacy and prevent inferences about local data, differential privacy techniques are applied by introducing statistical noise into the models prior to aggregation (29). This approach not only mitigates risks associated with data centralization but also allows the model to better adapt to the specific contexts of individual organizations, thereby improving validation effectiveness.

Furthermore, at the level of result sharing, a cryptographic logic based on Zero-Knowledge Proofs (ZKPs) is implemented, allowing the verification of accuracy and compliance without disclosing the underlying data (30). ZKPs enable one party (the "prover") to demonstrate to another party (the "verifier") that a specific condition has been met without revealing the information that supports this condition. This mechanism proves particularly useful in contexts where it is necessary to demonstrate compliance with Environmental, Social and Governance (ESG) criteria (31) or GRI standards without making sensitive data public. The proofs generated are cryptographically verifiable by third parties, ensuring system transparency and auditability without compromising confidentiality.

Taken together, these technologies contribute to establishing a robust, distributed, and privacy-oriented security architecture capable of supporting a transparent, reliable validation ecosystem in line with the most advanced standards of cybersecurity.

Practical Implementation: Sector-Specific Case Studies

To assess the effectiveness and adaptability of the AI-blockchain hybrid framework, two case studies were conducted in different sectors, each characterized by specific challenges and validation requirements: one focused on LCA analysis (32), and the other on sustainability reporting (33). The documents, originating from distinct organizational contexts, were subjected to the analytical process detailed below.

Life Cycle Assessment Analysis — Case Study: GSI S.r.l., Gestione Servizi Integrati

In the environmental services and facility operations sector, the LCA methodology is a crucial tool for objectively evaluating the environmental impact of operational

processes, materials used and adopted technological choices. However, the collection and validation of the data required for a comprehensive LCA analysis often prove complex, especially when dealing with articulated production systems or comparative assessments of alternative operational solutions (34).

In the specific case study concerning GSI S.r.l., the hybrid framework was applied to the validation of an LCA report related to a professional cleaning service performed over a representative surface area of 26,907 m². The objective was to compare two operational methods: the traditional immersion system and a more efficient pre-impregnation approach. The overall analysis process is summarized in the corresponding table.

The Data Acquisition Layer operated by extracting and normalizing data from documentary sources, such as internal reports, product certifications, and technical datasheets provided by manufacturers. Although no Internet of Things (IoT) sensors or real-time data sources were employed, the data collection was methodologically sound due to the use of advanced OCR tools, which ensured effective traceability and adequate standardization of relevant environmental data.

The AI-based Verification Layer analysed the entire life cycle structure—divided into upstream, corestream, and downstream stages—confirming consistency between numerical data and qualitative statements within the report. The calculated environmental impact reductions between the two systems were accurate and aligned with absolute data, and the contributions of each phase (e.g., the impact of electricity consumption in the corestream phase or the importance of waste disposal in the downstream phase) were correctly identified and described. No statistical anomalies or semantic inconsistencies were detected. Moreover, textual analysis using a lightweight LLM model, trained on a specific corpus of LCA reports, confirmed the absence of greenwashing both lexically and conceptually.

The Blockchain Certification Layer was not implemented in the original document; however, a readiness assessment for digital notarization was conducted. Based on internal coherence, data completeness, and narrative transparency, the document scored 8.3 out of 10—above the minimum threshold of 7.5 required for potential NFT-based cryptographic certification and blockchain registration (Fig.2).

Framework Module	Evaluation GSI_LCA	Summary Description	Score (0–10)
1. Data Acquisition Layer	⚠ Adequate	The report is well structured and uses representative sources for LCA (Upstream, Corestream, Downstream), with declared data and comparisons between two production systems. However, they do not include data acquired from IoT sensors, nor direct real-time sources.	6,5
2. AI Verification Layer	✓ Excellent	LCA analysis is coherent and complete; environmental claims are backed by quantitative data.	9
3. NLP & Greenwashing Check	✓ Excellent	No vague or misleading claims detected; no semantic or symbolic greenwashing present.	9,5
Total Score / Notarization Status		✓ ELIGIBLE	8,33
4. Blockchain Readiness	✓ Compliant	SUITABLE for notarization	

Figure 2. Summary table of the verification phases with positive (Compliant) evaluation of the document under analysis. Source: Authors.

This case study confirms that applying the AI-based hybrid framework can enhance the efficiency and reliability of LCA report validation, even in complex sectors such as environmental services. In the broader context of environmental, social, and economic reporting, sustainability reports are increasingly important tools for transparently and verifiably communicating organizational commitments to responsible practices (35). However, the growing scrutiny of stakeholders—combined with regulatory complexity and the multiplicity of data sources—necessitates structured methods to verify the reliability and consistency of reported information.

Sustainability Report Analysis — Case Study: Ferretti Group 2023 Social Report

In this case study, the hybrid framework was applied to validate the 2023 Sustainability Report of Ferretti Group (36), aiming to assess the methodological quality of the document, the transparency of ESG disclosures, and the potential presence of inconsistencies or greenwashing practices, as shown in the summary table.

The Data Acquisition Layer detected extensive use of structured internal data from certified environmental management systems (ISO 14001), operational reports, consumption monitoring, and technical datasheets. However, no integration of real-time data sources via IoT sensors or external or georeferenced datasets was identified elements that would have enhanced the robustness of the validation. Nonetheless, the data collection was methodically well-curated and thoroughly documented. The AI Verification Layer confirmed a high level of internal consistency. The report adheres to the structure and principles of the GRI (37), with precise quantitative indicators related to energy consumption, resource use, waste

management, social performance, and governance. Semantic and numerical analysis revealed no significant contradictions between reported data and narrative statements. The framework recognised a well-structured information architecture capable of adequately representing the complexity of the Group's activities.

The NLP-based textual verification and greenwashing detection module found overall coherent communication aligned with the underlying data. However, it identified some promotional wording (e.g., "environmental leadership," "sustainable innovation") that, although not inaccurate, occasionally lacked specific numerical evidence. While common in sustainability reports, such language may reduce perceived objectivity and should be monitored to prevent symbolic greenwashing.

The cryptographic certification module was not applied to the document: the report contains no hashes, digital signatures, or references to blockchain infrastructures. However, based on the three phases of analysis, the document received an average score of 7.3 (repeating) out of 10, slightly below the 7.5 threshold required to be deemed ready for automated notarization via NFT or DLT systems (Fig.3).

Framework Module	Evaluation FG_sustainability-report	Summary Description	Score (0–10)
1. Data Acquisition Layer	☒ Adequate	The report uses well-organized internal data and certifications, but lacks real-time IoT or external link for verifiable sources.	6,5
2. AI Verification Layer	✓ Excellent	Data consistency is high; the report follows GRI standards and presents coherent sustainability KPIs, particularly around emissions, materials, and energy.	8,5
3. NLP & Greenwashing Check	☒ Adequate	Statements about sustainability and innovation are generally aligned with data, but some marketing-style language lacks quantifiable support.	7
Total Score / Notarization Status			7,33
4. Blockchain Readiness	✗ Non-compliant	The document does not pass the minimum threshold of 7.5/10 required for certification. However, it is very close. Integration of primary data from objective sources or sensors (phase 1). Strengthening quantitative support for sustainability claims (phase 3), avoiding unfounded "green" claims.	

Figure 3. Summary table of the verification phases with negative (Non-Compliant) evaluation of the document under analysis. Source: Authors.

This case study demonstrates that, even in the absence of blockchain, the application of an AI-based hybrid framework can provide in-depth and objective analysis of the informational and narrative quality of sustainability reports. The integration of semantic analysis tools, cross-checking of data, and automated evaluation of ESG claims enables

the identification of improvement areas, reinforces document credibility, and enhances transparency for stakeholders. In the future, the adoption of cryptographic certification mechanisms could further complete this process, ensuring traceable, immutable, and increasingly trustworthy sustainability reporting.

Discussion: Impact on Systemic Design

The integration of a hybrid AI-blockchain framework for systemic report validation extends far beyond the realm of technical verification; it reconfigures the structural, procedural, and epistemic foundations of systemic design practice. By coupling AI-driven anomaly detection with decentralized blockchain certification, the framework introduces a distributed, tamper-resistant knowledge infrastructure that can reorient stakeholder incentives, reshape governance mechanisms, and inform the emergence of new design paradigms.

One of the most profound systemic implications lies in the attenuation of information asymmetries among heterogeneous stakeholder groups, facilitated by the enhanced transparency, traceability, and reliability of validated data. Presently, investors, consumers, and regulatory authorities operate in environments characterised by heterogeneous reporting practices, limited standardization, and the absence of independent verification mechanisms (38). The proposed framework operationalises a unified and auditable “source of truth” that enables both individual and collective decision-making processes to be grounded in evidence that is simultaneously accurate, verifiable, and immutable. For instance, investors may more precisely assess ESG-related risk profiles, consumers may discriminate between genuinely sustainable and merely nominally sustainable products, and regulators may monitor policy outcomes with greater precision. In systemic design terms, this constitutes a recalibration of feedback loops, reducing informational “noise” and enabling more adaptive, responsive, and evidence-driven interventions. Blockchain immutability further safeguards the integrity of validated datasets, mitigating risks of greenwashing and speculative misrepresentation, while providing robust intellectual property protection for designers’ methodological contributions.

A second transformative dimension is the articulation of a “verifiable design” paradigm for the circular economy. Conventional circular design has prioritised waste minimisation, reuse, and repairability, yet often in the absence of rigorous, independent verification of the actual environmental efficacy of these practices (39). The hybrid framework embeds objective validation—spanning material provenance tracing, lifecycle performance monitoring, and recovery or recycling rate verification—directly into the design process. This engenders a more accountable design ecosystem in which environmental claims are substantiated by empirical evidence. The paradigm thus reframes designers not merely as creators of artefacts but as active stewards of system-level material and value flows, incentivising design-for-disassembly, supplier traceability enhancements, and the proliferation of product-service systems demonstrably aligned with circularity metrics (40).

Finally, the framework possesses considerable normative potential in the evolution of ISO standardisation for ESG reporting (41). While current ISO instruments—such as ISO 14001 (environmental management) and ISO 26000 (social responsibility)—offer comprehensive procedural guidance, they lack an embedded mechanism for independent, technology-enabled ESG data verification. The hybrid framework may serve as a catalytic complement, introducing a globally standardised, cross-sector validation protocol capable of integration within existing and emergent ISO structures. Such integration would foster regulatory convergence, augment cross-border comparability, and reinforce transnational trust architectures—systemic prerequisites for the large-scale deployment of circular economy and sustainability strategies.

In sum, the hybrid AI-blockchain framework transcends its immediate function as a validation tool to act as an architectural intervention in systemic design, fostering a data ecosystem in which transparency, accountability, and verifiability are embedded by default. By institutionalising high-integrity feedback mechanisms and harmonising verification with design and governance processes, it contributes to the cultivation of adaptive, resilient, and ethically robust socio-technical systems.

Limitations and Challenges

Despite its numerous advantages, the implementation of the hybrid AI-blockchain framework also presents some limitations and challenges that must be addressed.

Among these are biases in lightweight LLM models for niche sectors, the issue of interoperability among enterprise blockchain platforms, and the computational costs of real-time validations.

Biases in lightweight LLM models represent a significant challenge, especially in niche sectors with limited training data availability. While lightweight LLMs offer benefits in terms of transparency and customization, they may be less accurate than larger, more generalized models—especially when trained on small or unrepresentative datasets (42). This can lead to classification errors, misinterpretations, and ultimately, incorrect report validation. To mitigate this risk, it is necessary to invest in the creation of high-quality training datasets that are representative of various sectoral niches. Additionally, transfer learning techniques should be used to adapt large pre-trained models to specific sectoral datasets. Continuous model performance monitoring and bias correction are also essential.

Interoperability among enterprise blockchain platforms represents a barrier to the widespread adoption of the hybrid framework. The lack of interoperability makes data sharing and collaboration between different organizations difficult (43). Overcoming this challenge requires promoting the development of open standards for blockchain interoperability and encouraging companies to adopt platforms that support such standards. Middleware solutions must also be developed to connect different blockchain platforms and enable secure, transparent data exchange.

The computational costs of real-time validations are a limitation to implementing the hybrid framework, particularly in applications requiring rapid data validation. Analysing data using AI algorithms and recording transactions on blockchain can demand significant computing resources, resulting in delays and high costs.

Alongside these technical and economic challenges, it is also crucial to consider the environmental impact inherent in adopting computation-intensive technologies like AI and blockchain (44). Training complex machine learning models requires considerable energy consumption, which may partially offset the environmental benefits of improved sustainability report validation. Similarly, some blockchain consensus mechanisms—such as Proof-of-Work (PoW)—are notoriously energy-intensive and significantly contribute to greenhouse gas emissions. Therefore, the development and

implementation of the hybrid AI-blockchain framework must be guided by strong awareness of its environmental implications, prioritizing energy-efficient solutions such as optimized AI models and blockchain systems based on alternative consensus mechanisms (e.g., Proof-of-Stake, Delegated Proof-of-Stake). It is also essential to offset unavoidable emissions through investments in renewable energy projects and certified carbon offsetting practices. Only through a holistic and responsible approach can the benefits of the hybrid framework be maximized while minimizing its environmental impact.

Conclusions and Future Developments

This research has proposed a unified hybrid AI-blockchain certification framework aimed at improving the reliability and transparency of systemic reports. The case studies conducted in various sectors have demonstrated the framework's potential to reduce information asymmetries, promote verifiable design for the circular economy, and influence ISO standardization in ESG reporting. Despite the identified limitations and challenges, the framework represents a significant step towards a more robust and trustworthy validation system.

Future developments of this research will focus on three main areas: integration with digital product passport systems, extension to materiality analysis validation, and the development of a unified framework for hybrid AI-blockchain certification.

Integration with Digital Product Passport (DPP) systems represents a natural evolution of the hybrid framework. DPPs are digital documents containing detailed product information, such as material origin, production process, environmental impact, and recycling instructions (45). Integrating the framework with DPPs would enable automatic validation of the information they contain, ensuring accuracy and reliability. This would facilitate product traceability, promote the circular economy, and allow consumers to make more informed decisions. It would also enhance the credibility of sustainability and systemic reports, life cycle assessments (LCA), and all initiatives related to the circular economy and CO₂ or waste impact reduction—allowing companies to focus their efforts on certified and recognised documents.

Finally, the long-term goal is to develop a unified framework for hybrid AI-blockchain certification applicable to a wide range of documents and sectors. This framework should define open standards for AI and blockchain integration, data validation, security management, and system governance. It should be flexible and adaptable to different industries and organizations, while also guaranteeing a minimum level of quality and reliability. The development of a unified framework would accelerate the adoption of hybrid AI-blockchain certification and foster greater transparency and accountability in sustainable reporting worldwide.

References

1. Shilling, A. F., & Liboni, M. (2023). La blockchain al servizio del design sistemico: Come una tecnologia può trasformare le validazioni dei report sistemici nelle aziende vinicole spagnole. In P. P. Peruccio, G. Pedone, & F. Bruschi (Eds.), *Relazioni di progetto* (pp. 250–262). Politecnico di Torino.
2. Reber. (2023, ottobre). *Il bilancio di sostenibilità Reber è nella blockchain*. <https://www.reber.it/news/block/>
3. Thistlethwaite, J., & Menzies, M. (2016). *Assessing the governance practices of sustainability reporting*. C. Hurst & Company.
4. Gold, N. (2020). Impact of corporate sustainability reporting practice on corporate performance: A review of literature. *International Journal of Industrial Management*, 8, 19–34. <https://doi.org/10.15282/ijim.8.0.2020.5760>
5. Brown, H. S., de Jong, M., & Levy, D. L. (2009). Building institutions based on information disclosure: Lessons from GRI's sustainability reporting. *Journal of Cleaner Production*, 17(6), 571–580. <https://doi.org/10.1016/j.jclepro.2008.12.009>
6. Nawrocka, D., & Parker, T. (2009). Finding the connection: Environmental management systems and environmental performance. *Journal of Cleaner Production*, 17(6), 601–607. <https://doi.org/10.1016/j.jclepro.2008.10.003>
7. Tariq, A., & Rahim, M. R. U. (2024). Role of artificial intelligence in enhancing sustainability reporting and green accounting in Industry 4.0. *Acta Marisiensis Seria Technologica*, 21(2), 31–38. <https://doi.org/10.62838/amset-2024-0015>

8. Subramani, N., Matton, A., Greaves, M., & Lam, A. (2020). *A survey of deep learning approaches for OCR and document understanding*.
<https://doi.org/10.48550/arXiv.2011.13534>
9. Rundo, F., & Di Stallo, A. L. (2020). *Predictive justice: Algorithms and deep learning technology*. ResearchGate.
https://www.researchgate.net/publication/338450142_GIUSTIZIA_PREDITTIVA_ALGORITMI_E_DEEP_LEARNING_TECHNOLOGY
10. Vinella, A., Capetz, M., Pattichis, R., Chance, C., & Ghosh, R. (2023). *Leveraging language models to detect greenwashing*. <https://doi.org/10.48550/arXiv.2311.01469>
11. Carammia, M., Iacus, S. M., & Porro, G. (2024). *Rethinking scale: The efficacy of fine-tuned open-source LLMs in large-scale reproducible social science research*.
<https://doi.org/10.48550/arXiv.2411.00890>
12. Gebrezgiher, Y. T., Camacho, D., & Park, J. H. A. (2025). *Blockchain-based framework for secure anomaly detection and trust management in IoV*. SSRN.
<http://dx.doi.org/10.2139/ssrn.5229668>
13. Hasan, M., Rahman, M. S., Janicke, H., & Sarker, I. H. (2024). Detecting anomalies in blockchain transactions using machine learning classifiers and explainability analysis. *Blockchain: Research and Applications*, 5(3), 100207.
<https://doi.org/10.48550/arXiv.2401.03530>
14. International Organization for Standardization. (2015). *ISO 14001: Environmental management systems – Requirements with guidance for use*.
<https://www.iso.org/iso-14001-environmental-management.html>
15. Frischknecht, R., Jungbluth, N., Althaus, H. J., et al. (2005). The ecoinvent database: Overview and methodological framework. *International Journal of Life Cycle Assessment*, 10, 3–9. <https://doi.org/10.1065/lca2004.10.181.1>
16. Meadows, D. H. (2008). *Thinking in systems: A primer*. Chelsea Green Publishing.
17. Sharma, P. (2023). Advancements in OCR: A deep learning algorithm for enhanced text recognition. *International Journal of Inventive Engineering and Sciences*, 10(8), 1–7.
<https://doi.org/10.35940/ijies.f4263.0810823>
18. Veronesi, V. (2024). *On the effects of measurement error and misclassification on clustering algorithms for mixed-type data* (Tesi di dottorato, Università degli Studi di

Milano-Bicocca). Università degli Studi di Milano-Bicocca.

<https://www.boa.unimib.it/handle/10281/465020>

19. Monaikul, N., Castellucci, G., Filice, S., & Rokhlenko, O. (2021). Continual learning for named entity recognition. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15), 13570–13577. <https://doi.org/10.1609/aaai.v35i15.17600>
20. Chen, N., Lin, X., Jiang, H., & An, Y. (2024). Automated building information modeling compliance check through a large language model combined with deep learning and ontology. *Buildings*, 14(7), 1983. <https://doi.org/10.3390/buildings14071983>
21. Mukhopadhyay, M. (2018). *Ethereum smart contract development: Build blockchain-based decentralized applications using Solidity* (1st ed.). Packt.
22. Liboni, M. (2024, 12–26 ottobre). Applying blockchain to systemic design: Ensuring project authenticity via quantitative report verification [Conference presentation]. In *Proceedings of Relating Systems Thinking and Design (RSD13)*, Oslo School of Architecture and Design, Oslo & online. <https://rsdsymposium.org/ai-blockchain-and-social-media>
23. Xuan, T. R., & Ness, S. (2023). Integration of blockchain and AI: Exploring application in the digital business. *Journal of Engineering Research and Reports*, 25(8), 20–39. <https://doi.org/10.9734/jerr/2023/v25i8955>
24. Squillero, G., & Marcelli, A. (2019). *Machine learning and other computational-intelligence techniques for security applications*. Politecnico di Torino.
25. Jin, Y. (2023). *Federated learning: Fundamentals and advances*. Springer.
26. Shashidhara, R., Nair, R. C., & Panakalapati, K. P. (2025). Promise of zero-knowledge proofs (ZKPs) for blockchain privacy and security: Opportunities, challenges, and future directions. *Security and Privacy*, 8, e461. <https://doi.org/10.1002/spy2.461>
27. Nalina, V., Navaneeth, S., Nayak, R. A., & Prakash, N. (2024). Decentralized file storage platform using IPFS and blockchain. In *2024 International Conference on Emerging Technologies in Computer Science for Interdisciplinary Applications (ICETCS)*, Bengaluru, India (pp. 1–6). <https://doi.org/10.1109/ICETCS61022.2024.10543705>
28. Patel, M., Dayan, I., Fishman, E. K., Flores, M., Gilbert, F. J., Guindy, M., Koay, E. J., Rosenthal, M., Roth, H. R., & Linguraru, M. G. (2023). Accelerating artificial intelligence: How federated learning can protect privacy, facilitate collaboration, and

- improve outcomes. *Health Informatics Journal*, 29(4).
<https://doi.org/10.1177/14604582231207744>
29. Cheng, Y., Liu, Y., Chen, T., & Yang, Q. (2020). Federated learning for privacy-preserving AI. *Communications of the ACM*, 63(12), 33–36.
<https://doi.org/10.1145/3387107>
30. Sun, X., Yu, F. R., Zhang, P., Sun, Z., Xie, W., & Peng, X. (2021). A survey on zero-knowledge proof in blockchain. *IEEE Network*, 35(4), 198–205.
<https://doi.org/10.1109/mnet.011.2000473>
31. Net0. (2024). *AI-powered ESG compliance: How anomaly detection helps sustainability reporting*. <https://www.net0.com/resources>
32. Hauschild, M. Z., Rosenbaum, R. K., & Olsen, S. I. (Eds.). (2018). *Life cycle assessment*. Springer International Publishing. <https://doi.org/10.1007/978-3-319-56475-3>
33. Machado, B. A. A., Dias, L. C. P., & Fonseca, A. (2020). Transparency of materiality analysis in GRI based sustainability reports. *Corporate Social Responsibility and Environmental Management*, 28(2), 570–580. <https://doi.org/10.1002/csr.2066>
34. Famiglietti, J., Toosi, H. A., Dénarié, A., & Motta, M. (2022). Developing a new data-driven LCA tool at the urban scale: The case of the energy performance of the building sector. *Energy Conversion and Management*, 256, 115389.
<https://doi.org/10.1016/j.enconman.2022.115389>
35. Boiral, O., Heras-Saizarbitoria, I., & Brotherton, M.-C. (2017). Assessing and improving the quality of sustainability reports: The auditors' perspective. *Journal of Business Ethics*, 155(3), 703–721. <https://doi.org/10.1007/s10551-017-3516-4>
36. Ferretti Group. (2023). *Sustainability report 2023*.
<https://www.ferrettigroup.com/it-it/Sustainability/sust-archive>
37. Global Reporting Initiative. (2021). *GRI 3000: Environmental standards series*.
<https://www.globalreporting.org/standards/>
38. Higgins, C., Tang, S., & Stubbs, W. (2020). On managing hypocrisy: The transparency of sustainability reports. *Journal of Business Research*, 114, 395–407.
<https://doi.org/10.1016/j.jbusres.2019.08.041>
39. Geissdoerfer, M., Savaget, P., Bocken, N. M. P., & Hultink, E. J. (2017). The circular economy – A new sustainability paradigm? *Journal of Cleaner Production*, 143, 757–768. <https://doi.org/10.1016/j.jclepro.2016.12.048>

40. Kibret, N., Edmonson, W., & Gebreyohannes, S. (2019). Category theoretic based formalization of the verifiable design process. In *2019 IEEE International Systems Conference (SysCon)* (pp. 1–8). IEEE. <https://doi.org/10.1109/syscon.2019.8836804>
41. Ronalter, L. M., Poltronieri, C. F., & Gerolamo, M. C. (2023). ISO management system standards in the light of corporate sustainability: A bibliometric analysis. *The TQM Journal*, 35(9), 256–298. <https://doi.org/10.1108/tqm-09-2022-0279>
42. Sheng, E., Chang, K.-W., Natarajan, P., & Peng, N. (2019). The woman worked as a babysitter: On biases in language generation (Version 2). *arXiv*. <https://doi.org/10.48550/arXiv.1909.01326>
43. Kotey, S. D., Tchao, E. T., Ahmed, A., Agbemenu, A. S., Nunoo-Mensah, H., Sikora, A., Welte, D., & Keelson, E. (2023). Blockchain interoperability: The state of heterogeneous blockchain-to-blockchain communication. *IET Communications*, 17(8), 891–914. <https://doi.org/10.1049/cmu2.12594>
44. Heston, T. F. (2024). A blockchain AI solution to climate change. *International Journal of Science and Research Archive*, 11(2), 450–454. <https://doi.org/10.30574/ijsra.2024.11.2.0471>
45. Voulgaridis, K., Lagkas, T., Angelopoulos, C. M., Boulogeorgos, A.-A. A., Argyriou, V., & Sarigiannidis, P. (2024). Digital product passports as enablers of digital circular economy: A framework based on technological perspective. *Telecommunication Systems*, 85(4), 699–715. <https://doi.org/10.1007/s11235-024-01104-x>

Authors

Martina Liboni, PhD Fellow, Politecnico di Torino,

<https://www.polito.it/personale?p=martina.liboni>, martina.liboni@polito.it

Francesca Mucchetti, PhD Fellow, Politecnico di Torino,

<https://www.polito.it/personale?p=francesca.mucchetti>, francesca.mucchetti@polito.it

Pier Paolo Peruccio, Full Professor, Politecnico di Torino,

<https://www.polito.it/personale?p=003150>, pierpaolo.peruccio@polito.it