IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, January 2026
Impact Factor: 7.67

A TSP Framework for Verifying Cosmetic Sustainability Claims on ONDC

Maitri More, Vaishnavi Rajput, Snehal Naik, Shruti Wawage, Prof. Parag Jambhulkar
Department of Computer Engineering
SCTR's Pune Institute of Computer Technology, Pune, India
maitrimoreofficial@gmail.com, vaishnavirajput273@gmail.com
snehalnaik139@gmail.com, shrutiwawagecm@gmail.com, pjambhulkar@pict.edu

Abstract: *The global cosmetics industry is experiencing a surge in sustainability and safety claims, often leading to greenwashing—the practice of misleading consumers with unverifiable labels like "organic," "paraben-free," or "eco-friendly." This paper presents a comprehensive survey of AI methodologies applicable to verifying environmental and safety claims. Our findings are structured to define the core technical requirements for a Technology Service Provider (TSP) aiming for network integration.We thoroughly review and synthesize literature concerning three critical areas: first, techniques involving Natural Language Processing (NLP) for extracting, classifying, and assessing textual claims; second, approaches leveraging Computer Vision (CV) for the automated detection and validation of visual certifications and logos; and third, methods for integrating chemical informatics and regulatory databases to verify ingredient-level safety. Furthermore, we propose a unified framework for synthesizing these verification outputs into an objective, quantifiable metric, which we term the Trust Score. This framework is discussed with a focus on its integration feasibility within the Open Network for Digital Commerce (ONDC) architecture.*

Keywords*:* AI Verification, Natural Language Processing, Computer Vision, Greenwashing, Cosmetics, ONDC, Trust Score

I. INTRODUCTION

The rapid growth of the global cosmetics and personal care industry, valued at hundreds of billions of dollars, has been paralleled by a dramatic shift in consumer behavior towards sustainability, ethical sourcing, and product safety. As a result, market demand for "clean," "natural," "organic," and "free-from" products has surged, driving companies to heavily market these attributes. This movement, while positive, has inadvertently created fertile ground for greenwashing—the practice of deceptively promoting products or policies as environmentally sound or safer than they actually are. Such practices, often manifested through vague, unsubstantiated, or misleading claims on packaging and digital storefronts, erode consumer trust and undermine genuine sustainability efforts. The core challenge lies in the information asymmetry between sophisticated corporate marketing and the average consumer's ability to verify complex claims, which often span chemistry, regulatory compliance, and supply chain ethics. Traditional manual verification is slow, costly, and cannot scale to meet the volume of products traded on digital platforms. Therefore, the market urgently requires a scalable, objective, and automated solution to scrutinize product claims at the point of sale.

This survey is motivated by the critical need for a scalable, objective, and automated verification mechanism to restore transparency in digital commerce. Specifically, the rise of open-network initiatives like the Open Network for Digital Commerce (ONDC) in India has created an immediate, practical demand for such a service. For a potential Technology Service Provider (TSP) to successfully integrate a claim verification service into the ONDC architecture, a clear, classified roadmap of the current state-of-the-art AI methodologies (NLP, CV) is essential. This paper serves as that roadmap, identifying proven techniques and defining the necessary integration architecture to transition verification from a manual, brand-dependent chore into an automated, network-standard service.

DOI: 10.48175/IJARSCT-31039

ISSN 2581-9429 IJARSCT

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, January 2026
Impact Factor: 7.67

II. RELATED WORK

The comprehensive verification framework proposed in this survey is built upon the synthesis of three distinct research domains: the theoretical understanding of market deception, the technical application of AI in unstructured data analysis, and the architectural context of open digital networks. This section reviews the foundational literature that underpins our analysis. We first establish the conceptual boundaries of the problem by reviewing existing taxonomies of greenwashing and the established link between unverified claims and trust erosion. Subsequently, we examine the technical precedents for using multi-modal AI in product verification and compliance within e-commerce. Finally, we situate our service within the strategic context of the Open Network for Digital Commerce (ONDC), reviewing the existing literature on the Beckn Protocol and the role of Technology Service Providers (TSPs) in such decentralized architectures.

A. Conceptual Foundations of Greenwashing and Trust

The fundamental challenge addressed by this survey is the mitigation of greenwashing—the deceptive practice of misleading consumers with unverifiable environmental or safety claims. The nature of this deception is often framed by the well-established taxonomy of the ’Six Sins of Greenwashing’ [10], which systematically categorize misleading practices ranging from the hidden trade-off (e.g., sustainable packaging but toxic ingredients) to the sin of outright lying (i.e., making false claims). Understanding these discrete categories is not merely academic; it is crucial for developing targeted AI detection models by establishing the specific linguistic and visual patterns that NLP and CV modules must be trained to identify.The ultimate consequence of persistent greenwashing is the erosion of consumer trust within the digital marketplace. Our proposed Trust Score framework directly addresses the academic need for a ’better conceptual trust definition’ in e-commerce by shifting the focus from subjective human factors to objective, verifiable metrics. As defined by GrabnerKra¨uter and Kaluscha [11], trust in e-commerce requires both system trust and assurance-based trust building measures. Our framework integrates technology-driven metrics—specifically, Claim NLP (for textual veracity) and Visual CV (for certification authenticity)—to instantiate these concepts. The AI system acts as a transparent, auditable assurance mechanism (based on verifiable evidence) to build trust in the digital system itself, independent of the vendor’s reputation. This synthesis of technical verification with foundational trust theory forms the basis of our quantifiable Trust Score. Finally, the entire verification structure must be engineered to enforce legal mandates, including the stringent US FTC Green Guides, the mandatory E.U. Green Claims Directive, and India’s CCPA Greenwashing Guidelines[12]. These regulations collectively establish the high bar for evidence, substantiation, and transparency that our proposed AI verification system is designed to meet at scale.

B. Precedents in AI-Driven Product Verification

While a unified, ONDC-integrated AI system specifically for combating cosmetics greenwashing remains a novel pursuit, the core technical components required for such a service are well-established in academic literature concerning product compliance, safety, and sustainability assessment. This section reviews the precedents for using AI techniques across the three necessary verification layers: ingredient safety, claim analysis, and visual recognition.

1) AI for Ingredient Safety and Regulatory Analysis:

The most critical component of the verification framework—checking claims against ingredient lists—has strong precedents in in silico modeling. Researchers have successfully applied AI and predictive modeling to address complex safety and regulatory issues in product formulation [1]. Specifically, predictive modeling has been demonstrated as a viable technique for evaluating ingredient safety and skin compatibility, often replacing traditional in vivo or physical testing [2]. This body of work validates the use of AI to automate decisions regarding regulatory status and toxicology, forming the technical basis for the ingredient-level checks within our Trust Score. Furthermore, recent literature has focused on building functional AI tools, such as toxicity analyzers and recommenders, which integrate analytical results into consumer-facing metrics [5], [6].

2) Deep Learning for Claim Extraction and Product Analysis

The use of advanced AI, particularly Deep Learning, for analyzing product-related text is established. While few studies focus specifically on greenwashing, precedents exist for utilizing these techniques to process complex product information. Lee et al. [3] demonstrated the successful application of Deep Learning to skincare product recommendations based on comprehensive cosmetic ingredient analysis. This work shows the capacity of modern NLP models to accurately extract, process, and interpret detailed textual product features—a capability that is directly transferable to classifying environmental claims and their associated qualifying language for the Claim NLP module.

3) Sustainability and Life Cycle Assessment

Beyond ingredient safety, AI has been utilized to verify broader sustainability claims, particularly those related to manufacturing and packaging. Research has integrated AI into Life Cycle Assessment (LCA) methodologies to map and analyze circularity for packaging materials [4]. This precedent is crucial for the framework's ability to verify non-ingredient claims, such as ”sustainable packaging” or ”zero-waste,” by linking textual claims to verifiable process data, a necessity for preventing the ”Sin of Hidden Trade-Off” identified in the greenwashing literature [10].

In summary, the literature confirms the feasibility of the proposed system's core technical requirements, from using AI for regulatory safety modeling to applying Deep Learning for detailed product analysis. The primary research gap is therefore not in the individual AI components, but in the synthesis of these multi-modal inputs into a unified, quantifiable Trust Score for deployment on open digital networks.

C. The Open Network Architecture (Beckn/ONDC) and Sandbox Validation

The final domain of this survey is the strategic and architectural context of deployment. Our entire framework is situated within India's paradigm of Digital Public Infrastructure (DPI), specifically leveraging the Open Network for Digital Commerce (ONDC). ONDC is built on the open-source Beckn Protocol [13], which aims to unbundle the traditional e-commerce value chain to enable interoperability and foster competition, thereby democratizing digital commerce in India [14]. Our proposed system is designed to fulfill the role of a Technology Service Provider (TSP), an independent entity providing specialized software and services to enhance the network's function, particularly in areas requiring trust and integrity.

1) Validating Architecture via the ONDC Sandbox

Crucially, the verification framework's architectural feasibility is validated by its ability to integrate with the ONDC environment. The ONDC Sandbox is a dedicated testing environment that accurately simulates the network's production characteristics and facilitates the development and validation of TSP solutions [15]. By leveraging the Sandbox, we can demonstrate the system's compliance with the necessary Beckn protocol specifications (e.g., standardized APIs, schema validation, and authorization) [16]. This approach allows us to confirm that our AI service can successfully perform the key functions of a verification TSP—specifically, handling transaction requests and securely returning the calculated Trust Score—without requiring full production onboarding, thereby proving the architectural readiness of the framework [17].

2) The TSP Role as an Assurance Mechanism

The AI verification service is positioned as a critical TSP by addressing a core challenge identified in the growth of decentralized networks: trust-building.By offering transparent, auditable compliance checks against regulatory baselines (as discussed in Section II.A), the AI TSP fills a vital need for strong assurance measures. This empowers MSMEs by providing a credible, network-wide mechanism to verify their product claims, thereby enhancing their visibility and competitiveness [12]. The utilization of the Sandbox serves to technologically anchor this strategic TSP role within the existing ONDC ecosystem specifications.

III. SYSTEM ARCHITECTURE

The AI-Driven Claim Verification Framework is engineered as a modular, three-tiered microservice designed for interoperability with the Open Network for Digital Commerce (ONDC) via the Beckn Protocol. This architecture is validated against the ONDC Sandbox specifications to prove its real-world integration capability. The system is designed to fulfill the role of a Technology Service Provider (TSP), providing a critical, auditable assurance layer to the decentralized network.The architecture comprises three distinct and independent layers optimized for high-throughput, secure, and transparent verification cycles (Fig. 1).

A. Ingestion Layer (ONDC Endpoint)

This layer serves as the external interface and the TSP's primary API endpoint, responsible for secure and compliant data exchange with the ONDC network.

Beckn Protocol Compliance: The system is configured to receive standardized product data packets (containing product title, description, image URL, and ingredient list) via Beckn transaction requests (e.g., on_search or on_confirm) from Buyer or Seller Network Participants.

Data Validation: The incoming payload undergoes immediate schema validation against the Beckn specification and functional checks to ensure the presence of all required data fields (text, image URL, ingredients) necessary for complete multi-modal verification.

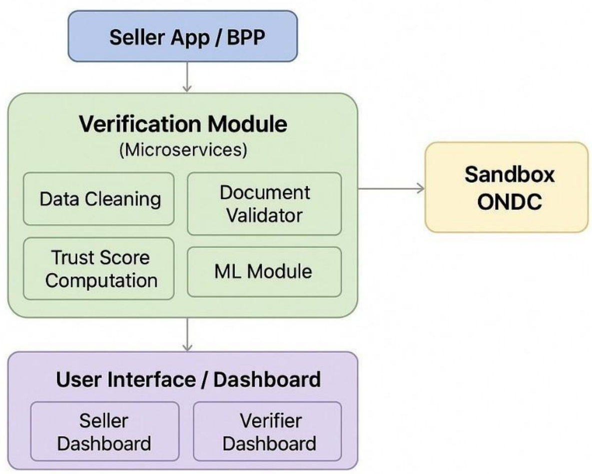

Fig. 1. Architecture of the AI-Driven Claim Verification Framework.

Security: This layer handles the necessary digital signature verification of the incoming Beckn message payload to confirm the identity of the sending network participant.

B. Verification Layer (Modular AI Microservices)

This layer represents the computational core where data is processed in parallel by specialized AI models, corresponding directly to the Verification Taxonomy (detailed in Section IV). Parallel Processing: Validated input is simultaneously routed to three independent microservices: the NLP Claim Module, the CV Visual Module, and the Ingredient Data Integration Module. This concurrent design minimizes latency. Microservice Isolation: The modular nature ensures that model updates or retraining in one domain (e.g., CV recognition) do not impact the stability or output of the other verification processes.

Quantifiable Output: Each module generates a normalized, quantifiable score based on its unique methodology, reflecting the level of compliance found in its domain.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, January 2026
Impact Factor: 7.67

C. Synthesis Layer (Trust Score Engine)

This final layer aggregates the results and computes the single, authoritative assurance metric returned to the ONDC network.

Score Aggregation: The normalized scores are fed into the Trust Score Engine.

Weighted Calculation: The engine applies a pre-defined, regulatory-informed weighting scheme to compute the final Composite Trust Score (T), ensuring that higher-priority regulatory or safety concerns carry appropriate influence.

Secure Response: The final Trust Score is digitally signed by the TSP's key, packaged according to the Beckn response protocol, and returned to the ONDC network (e.g., via the on_search or on_confirm response).

IV. KEY TECHNOLOGIES AND METHODOLOGIES

The AI-Driven Claim Verification Framework is architected as a set of dedicated microservices deployed as a Technology Service Provider (TSP) on the ONDC platform, leveraging the Beckn Protocol for data exchange. The core methodology employs three distinct AI modules that process multi-modal data concurrently to generate the final Composite Trust Score (T), enabling the sophisticated API-driven services listed below.

A. Textual Claim Verification (NLP Module)

This module is responsible for verifying the accuracy and authenticity of textual claims made in product descriptions, forming the foundation of the Claim Verification service. It uses a fine-tuned Transformer-based language model (such as BERT or RoBERTa) for two key tasks — sequence classification and Named Entity Recognition (NER). These allow the system to detect both explicit and implied "green" claims in the product title and description provided through the ONDC item object. Once identified, each claim is categorized into one of three types: Certified (C), Generic (G), or False (F).

The model also outputs a confidence score ($conf \in [0, 1]$) for every classification, which quantifies how certain the model is about its decision. This confidence value directly contributes to the Claim Validity Score (Sc) — a normalized metric that reflects how trustworthy the textual claim is and forms part of the response shared via the verification API.

B. Visual Claim Verification (CV Module)

This module employs computer vision to authenticate third-party certifications and assess product circularity, ensuring that visual data supports regulatory compliance. A high-performance Convolutional Neural Network (CNN) model—such as YOLO or ResNet—is used for object detection and image classification on the product image URL. In its certification focus, the system detects and verifies the presence and authenticity of eco-labels, third-party certifications, and regulatory markings (e.g., Green Dot, Leaping Bunny), providing visual evidence that either supports or contradicts the textual claims identified by the NLP module. In its circularity focus, the CNN identifies recycling symbols and resin identification codes ($k \in 1, 2, 3, \ldots, 7$) on packaging. Each detected code is mapped to a predefined recyclability factor (Rk), and the Circular Economy Score (CE) is computed using the highest recyclability factor found. This approach ensures that the score represents the strongest sustainability attribute that can be visually verified.

C. Ingredient Safety and Toxicity Assessment

This module performs compliance checks on product ingredients and powers two key services: Ingredient Risk Alerts and Sustainable Alternative Recommendations. It operates on a robust data processing pipeline that integrates a relational database engine with authoritative external data sources for verification. The input INCI (International Nomenclature of Cosmetic Ingredients) list is first standardized using stringmatching algorithms to ensure accurate and consistent ingredient identification across datasets. For hazard scoring, the module leverages the EWG Skin Deep database, which assigns each ingredient a hazard score from 1 (low risk) to 10 (high risk). These values are normalized into advisory hazard weights (wEWG,i). Simultaneously, the system crossreferences each ingredient with regulatory databases such as the E.U. CosIng list to assign a regulatory weight (wEU,i), where a score of 1 denotes a restricted or banned ingredient. The Safety Index (SI) is then computed by combining these two weights, applying a higher coefficient (λEU) to the regulatory compliance factor to emphasize mandatory safety constraints over advisory hazard data. The resulting SI value quantifies the overall ingredient safety level, while detailed risk factors and raw EWG hazard data are made accessible through a dedicated API endpoint for transparency and user insight.

D. Composite Trust Score Synthesis and API Delivery

The three normalized scores are combined into the final, weighted Composite Trust Score (T), which serves as the primary API payload.

$$T = WCSc + WISI + WP\,CE \quad \text{where } WC + WI + WP = 1$$

A higher weighting is applied to WI to ensure that safety concerns identified by the Data Integration Module immediately penalize the final T score, reflecting their critical regulatory importance.

V. CHALLENGES AND FUTURE DIRECTIONS

Despite significant progress, several challenges remain in developing robust and widely applicable AI interview systems.

A. Data and Technical Limitations

1) Training Data Scarcity and Diversity: The accuracy of the NLP and CV models is fundamentally dependent on the quality and diversity of training data. High-quality, accurately labeled datasets for specialized green claims, particularly those reflecting the diversity of Indian regional cosmetic markets and the vast variety of MSME product imagery, are currently limited. Biases in the training datasets can lead to algorithmic bias, potentially discriminating against products from underrepresented manufacturers or those with non-standard packaging.

2) Noisy Multimodal Data: Integrating heterogeneous data (unstructured text, noisy JPG images, and structured ingredient lists) introduces technical complexity. Issues like lowresolution product images, inconsistent label placement, and variability in material definitions can lead to alignment issues and reduced accuracy in calculations.

3) Beckn Latency Management: While the architecture is designed for parallel processing, the real-time constraints of an e-commerce transaction require the TSP to return the Trust Score (T) within strict time limits dictated by the Beckn protocol's request/response timing. Maintaining this low latency while executing three complex AI models remains a significant architectural challenge.

B. Ethical and Governance Concerns

1) Algorithmic Transparency and Explainability: For the score to function as a trust mechanism within the ONDC ecosystem, the factors contributing to non-compliance must be transparently communicated to Network Participant. This requires implementing Explainable AI (XAI) techniques within the Synthesis Layer to justify the assigned score.

2) Accountability in Dispute Resolution: As the TSP provides the assurance layer, clear accountability must be established regarding the Trust Score's role in online dispute resolution (ODR) within ONDC. Procedures must define the recourse available to a seller who contests the AI's verification outcome, necessitating human-in-the-loop oversight.

C. Future Scope

1) Cross-Domain ONDC Services: The methodology can be extended beyond cosmetics to other domains within ONDC, such as food, where claims related to organic certification or nutritional facts require similar multi-modal verification across text, image, and ingredient data, fulfilling ONDC's goal of fostering specialized TSP innovations.

2) Recommendation Engine Development: A significant future direction involves leveraging the calculated Composite Trust Score (T) to power a Trust-Based Recommendation Engine. This engine would use the compliance metrics as primary features to filter and suggest alternative products with higher scores when a user attempts to purchase a lowscoring item. This elevates the TSP's function from passive verification to active consumer guidance, directly reinforcing ethical purchasing behavior across the ONDC ecosystem.

VI. CONCLUSION

This survey addressed the growing challenge of greenwashing within the rapidly evolving digital commerce ecosystem. We examined the issue not merely as a matter of deceptive marketing, but as a deeper systemic problem that undermines consumer trust and regulatory integrity. Our analysis established that, although existing research supports isolated AIbased verification components, there remains a lack of a unified, multi-modal framework engineered for deployment within decentralized public digital infrastructures.

The core contribution of this work is the conceptualization and architectural validation of an AI-Driven Claim Verification Framework, implemented as a Technology Service Provider (TSP) integrated with the ONDC Sandbox. We proposed a three-part Verification Taxonomy—comprising Textual NLP, Visual CV, and Ingredient Data Integration—which together form the structural foundation of our verification methodology.

Furthermore, we introduced the Composite Trust Score (T), a quantifiable assurance metric derived through a weighted synthesis of the three AI modules. This score is explicitly designed to align with regulatory priorities and to mitigate the deceptive patterns encapsulated in the “Six Sins of Greenwashing.”

By grounding the framework in mandatory global regulatory baselines and aligning it with the interoperable standards of the Beckn protocol, we demonstrate a technically viable and strategically scalable pathway to embed verifiable transparency and trust within digital commerce ecosystems.

REFERENCES

[1] M. Di Guardo et al., “Artificial Intelligence in Cosmetic Formulation: Predictive Modeling for Safety, Tolerability, and Regulatory Perspectives,” Cosmetics, vol. 12, no. 4, 2025.

[2] M. Juncan et al., “Evaluation of the Safety of Cosmetic Ingredients and Their Skin Compatibility through In Silico and In Vivo Assessments,” Toxics, vol. 12, no. 7, 2024.

[3] H. Lee et al., “Deep Learning-Based Skin Care Product Recommendation: A Focus on Cosmetic Ingredient Analysis,” J. Cosmetic Dermatology, 2024.

[4] R. Ragava Raja and P. Khanna, “Life Cycle-Based Sustainability Assessment and Circularity Mapping for Packaging Materials: Integrating Artificial Intelligence,” J. AI and Industry, vol. 7, no. 1, 2025.

[5] “Product Ingredient Toxicity Analyzer and Recommender,” Int. J. Adv. Res. Comput. Commun. Eng. (IJARCCE), vol. 13, no. 2, 2024.

[6] “Product Ingredient Analysis,” Int. J. Recent Adv. Sci. Eng. Technol. (IJRASET), vol. 12, no. 3, 2024.

[7] EWG, “Understanding Skin Deep Ratings,” [Online]. Available: https: //www.ewg.org/skindeep/understanding skin deep ratings/. [Accessed: Oct. 6, 2025].

[8] EWG, “EWG’s Skin Deep: Helping Consumers and Scientists Understand the Connection Between Personal Care Products and Health,” [Online]. Available: https://www.ewg.org/news-insights/news/2025/09/ ewgs-skin-deepr-helping-consumers-and-scientists-understand-connection. [Accessed: Oct. 6, 2025].

[9] Think Dirty, “Our Methodology — How We Rate Ingredients,” [Online]. Available: https://www.thinkdirtyapp.com/methodology/. [Accessed: Oct. 6, 2025].

[10] S. Dimitrieska, A. Stankovska, and T. Efremova, “The Six Sins of Greenwashing,” Economics and Management, South-West University “Neofit Rilski”, Blagoevgrad, vol. 13, no. 2, pp. 82–89, 2017.

[11] S. Grabner-Kra¨uter and A. E. Kaluscha, “Consumer Trust in Electronic Commerce: Conceptualization and Classification of Trust Building Measures,” in Handbook of Research on Electronic Commerce in the Digital Economy, Edward Elgar Publishing, 2008. doi: 10.4337/9781848445086.00009.

[12] K. C. L. V. S. Subhash and V. Singh, “ESG/Regulatory: Green Claims with Real Stakes,” Legal Business Online, Oct. 2024. [Online].

[13] “Imagining with Beckn,” Beckn Protocol, [Online]. Available: https:// becknprotocol.io/imagining-with-beckn/. [Accessed: Nov. 1, 2025].

[14] A. S. George and A. George, “Open Network for Digital Commerce (ONDC): Democratizing Digital Commerce and Curbing Digital Monopolies in India,” in Proc. Int. Conf. Innov. Manag., Eng. Sci., 2022, pp. 92–102. doi: 10.5281/zenodo.6799694.

[15] ONDC, "ONDC — Mock & Sandbox," ONDC Official Website, [Online]. Available: https://mock.ondc.org/. [Accessed: Oct. 2, 2025].

[16] ONDC, "ONDC ONEST Sandbox User Guide," ONDC Official Website, [Online]. Available: https://onest-mock-service.ondc.org/ user-guide. [Accessed: Oct. 2, 2025].

[17] M. Aggarwal and B. Madan, "ONDC Unpacked: Deep Dive into India's Open Network Revolution," ResearchGate, Apr. 2025. [Online].