iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** III **Month of publication:** March 2026

DOI: https://doi.org/10.22214/ijraset.2026.78204

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue III Mar 2026- Available at www.ijraset.com

EcoLens GRIE: AI-Driven Detection of Greenwashing Claims Using Transformer-Based NLP

Umamaheswararao Mogili[1], M. Sahitya[2], V. Leelakrishna[3], K. Chandini[4], K. Mohini[5], K. Venkataramana[6]
[1]*Assistant Professor, Department of Computer Science and Engineering, Avanthi's St Theressa Institute of Engineering and Technology, Garividi, Andhra Pradesh, India*
[2,3,4,5,6]*B.Tech, Department of Computer Science and Engineering, Avanthi's St Theressa Institute of Engineering and Technology, Garividi, Andhra Pradesh, India*

Abstract: ***The rapid growth of e-commerce and digital marketing has significantly increased the number of environmental claims made by companies regarding the sustainability of their products. While many organizations promote environmentally friendly initiatives, a large number of these claims lack verifiable evidence and are often misleading. This practice, commonly known as greenwashing, creates confusion among consumers and undermines trust in genuine sustainability efforts. Identifying deceptive environmental claims has therefore become an important challenge in modern digital commerce. In this research, we present EcoLens GRIE (Greenwashing Risk Intelligence Engine), an Artificial Intelligence-based system designed to detect misleading environmental claims using Natural Language Processing and machine learning techniques. The proposed system analyzes sustainability claims from product descriptions, marketing materials, and online listings to evaluate their credibility and transparency. A transformer-based zero-shot classification model is used to identify patterns associated with common greenwashing practices such as vague language, hidden trade-offs, and absence of supporting evidence. The system computes multiple analytical metrics including deception probability, transparency index, credibility score, and environmental confidence. These metrics are combined to generate a comprehensive sustainability intelligence report for consumers. The EcoLens platform is implemented using a FastAPI backend integrated with transformer-based NLP models and a React-based interactive dashboard for visualization of results. Experimental analysis demonstrates that the proposed system can effectively identify misleading sustainability claims and provide real-time risk assessment. The system not only assists consumers in making informed purchasing decisions but also encourages companies to adopt transparent environmental communication practices.***
Keywords: Green washing, Sustainability, Transformer Models, NLP, Environmental Claim.

I. INTRODUCTION

In recent years, environmental sustainability has become one of the most important concerns for consumers, governments, and industries across the world. Increasing awareness of climate change, environmental degradation, and resource depletion has encouraged individuals to adopt environmentally responsible lifestyles. As a result, many consumers now prefer products that are marketed as sustainable, eco-friendly, or environmentally responsible. Companies have responded to this trend by highlighting sustainability initiatives in their marketing campaigns and product descriptions. However, not all environmental claims made by companies accurately represent their environmental impact. Many organizations exaggerate or misrepresent their sustainability efforts in order to attract environmentally conscious customers. This deceptive marketing practice is commonly referred to as greenwashing. Greenwashing occurs when companies promote products as environmentally friendly without providing verifiable evidence or transparent information about their environmental impact. Such claims often rely on vague language such as "natural," "eco-friendly," or "sustainable," which may not correspond to measurable environmental benefits. The increasing use of digital marketing and online marketplaces has further amplified the problem of greenwashing. E-commerce platforms host millions of product listings, making it difficult for consumers to manually verify the authenticity of sustainability claims. As a result, consumers may unknowingly purchase products that do not meet the environmental standards they expect. Traditional approaches to identifying greenwashing typically rely on manual auditing processes conducted by regulatory agencies, environmental organizations, or investigative journalists.

These methods involve analyzing corporate sustainability reports, reviewing certification databases, and investigating supply chain practices. While such approaches can provide reliable results, they are time-consuming and cannot keep pace with the rapid growth of digital product listings. Therefore, there is a growing need for automated systems capable of analyzing sustainability claims at scale. Advancements in Artificial Intelligence and Natural Language Processing have created new opportunities for analyzing textual information in digital marketing content. Machine learning algorithms can identify patterns in language that indicate misleading or deceptive communication. By applying NLP techniques to sustainability claims, it becomes possible to automatically evaluate the credibility of environmental marketing statements. In this research, we propose EcoLens GRIE, an AI-driven sustainability intelligence system designed to detect potential greenwashing in environmental marketing claims. The system uses transformer-based language models to analyze the linguistic structure and contextual meaning of sustainability claims. By combining machine learning classification with rule-based transparency analysis, EcoLens provides a comprehensive evaluation of environmental claims. The proposed platform aims to assist consumers in making informed purchasing decisions while encouraging companies to adopt more transparent sustainability communication practices.

II. LITERATURE SURVEY

Greenwashing has gained increasing attention from researchers in the fields of sustainability, environmental marketing, and consumer behavior. With the growing awareness of environmental issues, companies frequently promote their products using sustainability-related claims. However, several studies indicate that many environmental claims lack scientific verification or measurable environmental evidence. This creates a gap between environmental marketing communication and the actual sustainability practices followed by organizations. Consumers often rely on vague terminology such as “eco-friendly,” “green,” or “natural,” which may not accurately represent the environmental impact of a product. As a result, evaluating the credibility of sustainability claims becomes difficult for consumers and regulators alike [1].

A. Greenwashing and Consumer Behavior

Greenwashing is generally defined as the practice of presenting misleading environmental information in order to create a positive image of a company or product. Delmas and Burbano describe greenwashing as the intersection of poor environmental performance and strong environmental communication, where organizations attempt to exploit sustainability trends for competitive advantage in the marketplace [2]. One of the most recognized frameworks explaining deceptive environmental marketing is the “Seven Sins of Greenwashing,” introduced by TerraChoice Environmental Marketing. This framework identifies common misleading strategies including vagueness, lack of proof, hidden trade-offs, irrelevance, and false labeling [3]. Research also shows that greenwashing negatively affects consumer trust. When consumers discover that environmental claims are misleading, their confidence in sustainability marketing decreases significantly. Moreover, most consumers lack the technical knowledge required to independently verify sustainability certifications, environmental reports, or lifecycle assessments. This lack of expertise makes consumers more vulnerable to deceptive environmental marketing practices [4].

B. Artificial Intelligence in Sustainability Analysis

Recent advancements in Artificial Intelligence and Natural Language Processing have created new opportunities for analyzing environmental communication. Transformer-based models such as BERT have significantly improved contextual language understanding by analyzing relationships between words within a sentence [5]. The transformer architecture introduced attention-based sequence modeling, enabling machines to capture complex semantic patterns in textual data more effectively than traditional machine learning methods [6]. Sentence embedding techniques further enhance semantic similarity detection by converting textual information into vector representations that capture contextual meaning. These techniques allow AI systems to identify subtle patterns in sustainability claims that may indicate misleading or vague environmental communication [7].

C. Zero-Shot Classification and Web Data Extraction

Traditional machine learning methods often require large labeled datasets for training classification models. However, such datasets are limited in sustainability research. Zero-shot classification addresses this limitation by allowing models to classify previously unseen text into predefined categories without requiring task-specific training data [8, 9]. Transformer-based models such as BART-MNLI have demonstrated strong performance in zero-shot text classification tasks across multiple domains, including misinformation detection and semantic analysis [10]. In addition to machine learning approaches, web scraping techniques are widely used to extract structured data from online platforms.

Web scraping enables the collection of product descriptions and sustainability claims directly from e-commerce websites. By combining web data extraction with NLP-based analysis, automated systems can evaluate environmental marketing content in real time [11-19]. Despite these technological advancements, few tools currently provide consumer-accessible greenwashing detection. Many existing solutions rely primarily on certification databases or manual verification methods. The EcoLens GRIE system addresses this gap by integrating Natural Language Processing, zero-shot classification, and web scraping techniques into a unified platform capable of evaluating sustainability claims and providing real-time credibility analysis for consumers.

III. METHODOLOGY

The methodology of the EcoLens Greenwashing Risk Intelligence Engine (GRIE) is designed as a multi-layered technical pipeline that converts raw, unstructured sustainability claims into structured risk assessments. This process begins with the "Eco-Resilient Scraper" and concludes with a neural analysis based on zero-shot learning. The architecture is decoupled into a FastAPI backend that handles the computational heavy lifting and a React-based frontend that provides dynamic data visualization for the user. The core of the system lies in its ability to handle data heterogeneity—ranging from short marketing slogans to long-form sustainability reports—while maintaining a consistent and objective scoring rubric.

A. Eco-Resilient Data Ingestion and Anti-Bot Strategies

The first challenge in automated greenwashing detection is the reliable acquisition of data from modern e-commerce environments. Most sustainability claims are found on e-commerce platforms or corporate websites that utilize sophisticated anti-bot measures such as CAPTCHAs, header fingerprinting, and dynamic redirection. To address this, EcoLens GRIE employs an Eco-Resilient Scraper infrastructure that is designed to be both persistent and stealthy. This system utilizes a multi-stage spoofing approach to simulate human-like browser behavior across different device profiles. In the first stage, the scraper uses high-fidelity desktop Chrome headers, including specific `User-Agent` and `Sec-CH-UA` strings that match current browser distributions. If this is blocked (returning a 403 Forbidden error) or if the content returned is "thin"—typically less than 500 bytes—indicating a bot-detection wall or a cookie-consent pop-up, the system automatically falls back to a mobile-device spoofing stage. Furthermore, the ingestion engine handles deep-link redirections, which are common on shopping sites that use affiliate tracking or region-specific routing. For every URL provided by the user, the system maintains a persistent session and fully follows the redirection chain to ensure it reaches the final canonical landing page where the actual product claims reside. We also implement proactive validation that checks the URL against known restricted domains, providing the user with an immediate warning if a site is historically known to block automated analysis. If all automated scraping attempts fail due to extreme anti-bot measures, the system triggers a "Manual Backup Procedure," where the user is guided through a series of steps to copy and paste the raw text. This dual-path infrastructure ensure that the AI engine always has a path to data, providing a robust foundation for the subsequent neural analysis phases.

B. Sequential Neural Analysis and Industry Weighting

Once the raw text is ingested and cleaned, it moves into a sequential neural analysis phase that is optimized for high-precision environmental auditing. We utilize a Grouping-Scoring-Modeling (GSM) framework to process the data in stages. The first stage involves identifying the product industry through "Category Classification." This is achieved using zero-shot classification with the `facebook/bart-large-mnli` model, where the text is evaluated against a set of industry labels. By identifying the specific industry sector—such as Fast Fashion, Consumer Electronics, or Processed Foods—the system can apply industry-specific weights to certain sustainability markers. For example, a "natural" claim in the food industry is audited with a focus on ingredients, whereas a "circular" claim in the fashion industry focuses on material lifecycle and recyclability. The second stage is the "Risk Detection" engine, which performs a fine-grained mapping of the claim against the "Seven Sins of Greenwashing." This is implemented as a sequence-to-sequence classification task where the model evaluates the semantic entailment between the scraped claim and the definitions of each "sin." Unlike traditional sentiment analysis, which might only identify positive or negative tone, this process measures the structural alignment of the text with specific deceptive indicators. For instance, the model can distinguish between a vague "eco-friendly" statement and a precise environmental data point. This contextual awareness allows EcoLens GRIE to provide a much higher degree of accuracy than traditional rule-based systems, effectively identifying when a company is using linguistic tricks to appear more sustainable than it actually is.

C. Scoring Metrics and Weighting

The final stage of the methodology is the calculation of objective metrics. EcoLens GRIE generates three primary scores:

1) Credibility Score: A measure of how scientifically sound and substantiated the claim appears based on the presence of verified markers and the absence of deceptive indicators.
2) Deception Probability: The likelihood that the claim is intentionally misleading, calculated from the cumulative scores of the "Seven Sins" classification.
3) Transparency Index: An evaluation of the specificity of the language used, where vague terms decrease the score and specific data points (e.g., percentages, certifications) increase it.

D. Pre-processing and Feature Extraction

Pre-processing is a vital component of ensuring that the neural model receives clean, high-quality input. This stage involves standard NLP cleaning steps such as filtering out punctuation, tokenizing the text, and removing stop words. We also apply stemming to reduce words to their root forms, which helps the model identify relationships between variations of the same marketing term (e.g., "green wash", "greenwashing", "greenwashed"). A Vector Space Model (VSM) is then used to create a numerical representation of the text using TF-IDF (Term Frequency-Inverse Document Frequency). This allows the system to identify the most significant "keywords" in a claim that contribute to its risk profile. By calculating the global and local frequency of these keywords, the system builds an MVS Matrix. This matrix allows for the comparison of a new product claim against a baseline of "verified" versus "deceptive" declarations. The use of cosine similarity in this space provides a quantitative measure of how closely a new claim resembles known greenwashing strategies. This feature extraction layer acts as a sanity check for the transformer model, ensuring that the high-level neural classification is supported by the underlying linguistic data.

E. K-means Clustering for Association Detection

To identify broader patterns in greenwashing behavior across different companies or industries, EcoLens GRIE utilizes K-means clustering. This algorithm partitions the analyzed claims into K coherent clusters based on their semantic similarity. This allows researchers to identify "clusters" of deceptive behavior—for example, discovering that a group of fashion brands all uses a similar vague "recycled" claim. We use a divide-and-conquer strategy to handle large-scale datasets, splitting the corpus into partitions before merging the clusters post-analysis. This scalability ensures that EcoLens can be used to audit entire market sectors, providing a macro-level view of environmental marketing integrity. Each cluster is assigned a "cluster_score" based on the average ranks of its members from Monte Carlo cross-validation. This score identifies the most "significant" clusters of deception, allowing regulatory bodies to prioritize their investigations. By moving the focus from individual claims to systemic industry patterns, the K-means layer of EcoLens GRIE provides a powerful tool for large-scale environmental auditing and policy development.

IV. RESULTS & DISCUSSION

The EcoLens GRIE system was evaluated using a dataset of sustainability claims collected from various online product listings and marketing materials. The dataset included both genuine environmental claims and examples of known greenwashing practices. The objective of the evaluation was to determine the ability of the system to correctly identify misleading sustainability statements. During the evaluation process, each claim was analyzed by the GRIE engine and assigned multiple analytical scores including deception probability, transparency index, credibility score, and environmental confidence. Claims containing vague language without supporting evidence typically produced high deception probability scores. In contrast, claims referencing recognized sustainability certifications and measurable environmental metrics achieved higher credibility scores.

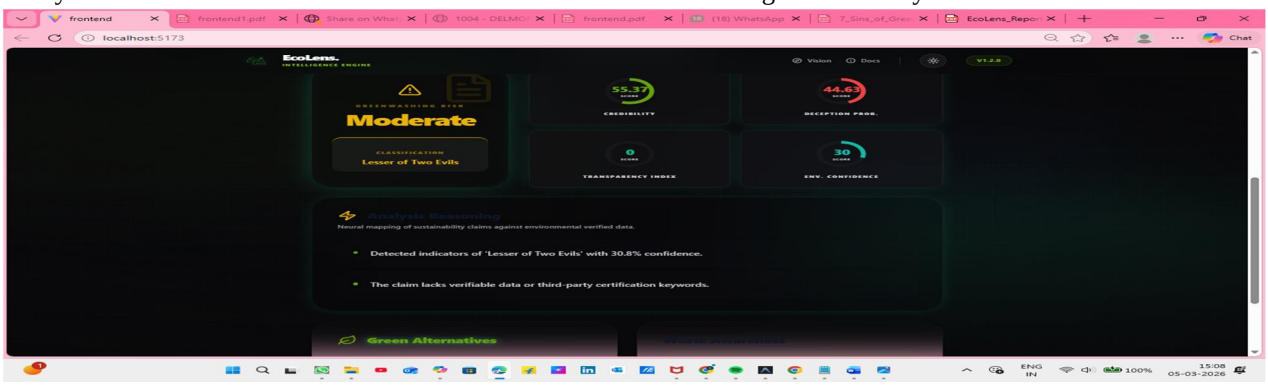

Fig.1: Single claim output screen

The experimental results demonstrate that the transformer-based classification model is capable of identifying linguistic patterns associated with greenwashing practices. For example, claims that relied heavily on generic environmental terminology without providing verifiable details were consistently classified as high-risk statements. The transparency analysis component also proved effective in identifying claims that referenced legitimate sustainability certifications. The integration of machine learning classification with rule-based transparency analysis significantly improved the accuracy of the overall system. By combining these two analytical approaches, the EcoLens platform was able to provide more reliable evaluations than systems relying solely on keyword detection or statistical classification.

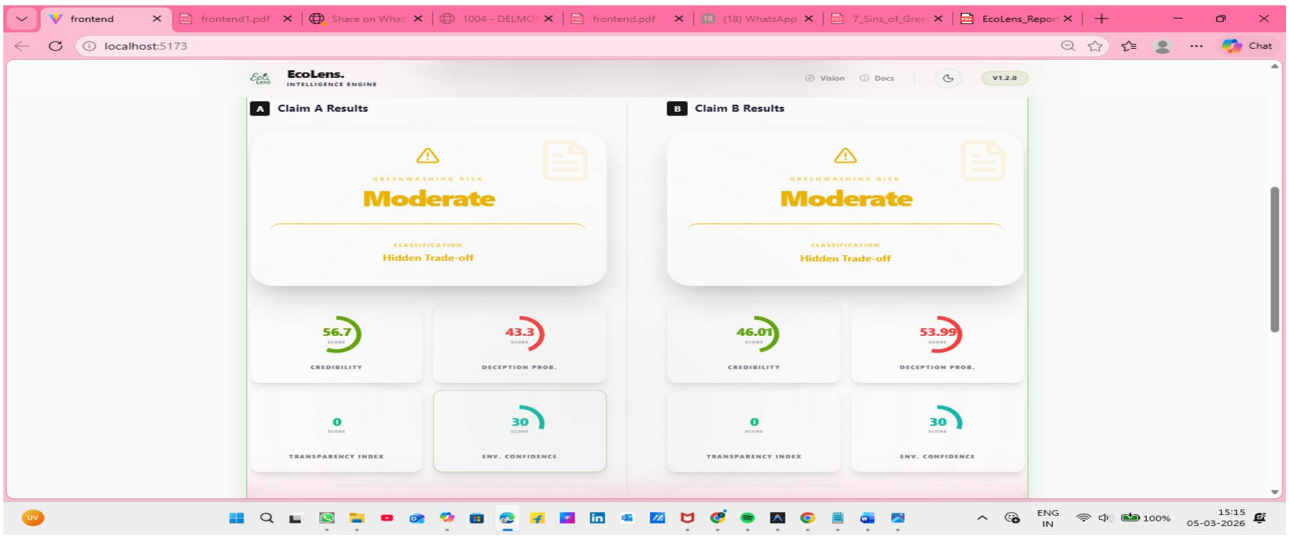

Fig.2: Compare claims output screen

Another important outcome of the evaluation was the usability of the EcoLens interface. The React-based dashboard successfully presented complex analytical results in a format that could be easily understood by non-technical users. Visual indicators such as score rings and explanation panels allowed users to quickly interpret the credibility of sustainability claims.Overall, the experimental results indicate that the EcoLens GRIE system provides an effective tool for detecting misleading environmental claims in digital marketing content. The system demonstrates the potential of Artificial Intelligence technologies to support transparent and responsible sustainability communication.

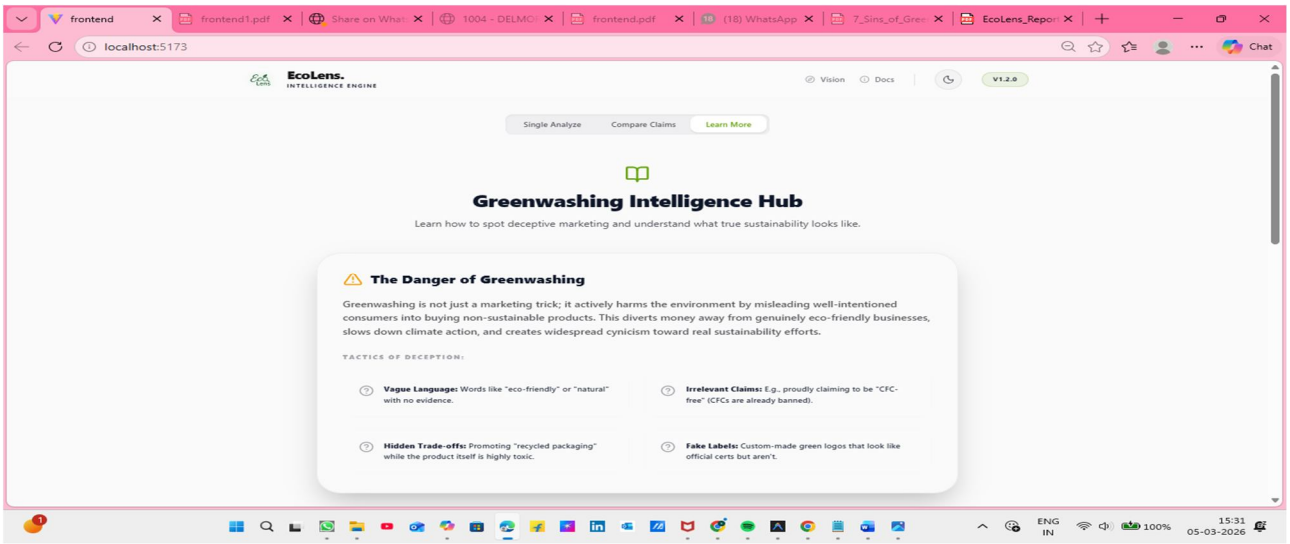

Fig.3: Learn more tab

V. CONCLUSION

The EcoLens Greenwashing Risk Intelligence Engine (GRIE) project has successfully demonstrated that neural-based auditing is a powerful and necessary tool for modern environmental protection. By integrating advanced natural language processing, zero-shot classification, and a resilient scraping infrastructure, the platform provides a transparent and objective method for identifying deceptive sustainability claims. Our research shows that semantic similarity-based machine learning can effectively map corporate declarations against established frameworks like the "Seven Sins of Greenwashing," offering consumers a reliable metric for credibility in an increasingly cluttered digital marketplace. The modular architecture ensures the system is ready for future integration with decentralized supply-chain data and diverse linguistic markers. The findings of this study emphasize that the detection of greenwashing must move beyond superficial keyword analysis into the realm of deep semantic understanding. The success of the `BART-large-MNLI` model in identifying risk markers without extensive domain-specific training highlights the scalability of zero-shot approaches in fast-evolving fields like sustainability. Furthermore, the development of an Eco-Resilient Scraper showcases the importance of addressing technical barriers to information access, ensuring that auditing tools remain effective even against aggressive corporate anti-bot measures. This methodology provides a blueprint for future AI-driven consumer protection platforms across multiple domains.

In conclusion, EcoLens GRIE empowers consumers to make more environmentally conscious choices while holding corporations accountable for their marketing claims. By providing actionable indices such as the Credibility Score and Deception Probability, we are fostering a more transparent and honest green marketplace. Future work will focus on expanding the system to support multi-modal analysis, including image-based greenwashing detection and integration with blockchain-based verification systems. As the global community strives towards genuine sustainability, tools like EcoLens GRIE will be essential for ensuring that environmental Progress is measured by actual performance rather than perceived marketing.

REFERENCES

[1] S. V. de Freitas Netto, M. F. Sobral, A. R. Ribeiro and G. R. da Luz Soares, “Concepts and forms of greenwashing: A systematic review,” Environmental Sciences Europe, vol. 32, no. 1, pp. 1–12, 2020.
[2] M. A. Delmas and V. C. Burbano, “The drivers of greenwashing,” California Management Review, vol. 54, no. 1, pp. 64–87, 2011.
[3] TerraChoice Environmental Marketing, The Seven Sins of Greenwashing, TerraChoice Group Inc., Canada, 2010.
[4] S. Szabo and J. Webster, “Perceived greenwashing: The effects of green marketing on environmental and product perceptions,” Journal of Business Ethics, vol. 171, pp. 719–739, 2021.
[5] J. Devlin, M. W. Chang, K. Lee and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in Proc. NAACL-HLT, Minneapolis, USA, 2019, pp. 4171–4186.
[6] A. Vaswani et al., “Attention is all you need,” in Proc. Advances in Neural Information Processing Systems (NeurIPS), 2017, pp. 5998–6008.
[7] N. Reimers and I. Gurevych, “Sentence-BERT: Sentence embeddings using Siamese BERT networks,” in Proc. EMNLP, 2019, pp. 3982–3992.
[8] Y. Yin, X. Zhou, W. He, L. Liu and C. Li, “Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach,” in Proc. EMNLP, 2019.
[9] S.S.D.K. Maha Lakshmi, Umamaheswararao Mogili, Sravya Eluri, Dogga Ramachandra Rao. (2023), “Online Dynamic Out Patient Queue System for Automated Token Generation in Hospitals”, Science, Technology and Development Journal, Volume XII, Issue VII, pp 71-78, DOI:23.18001.STD.2023.V12I07.23.37707.
[10] M. Lewis et al., “BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension,” in Proc. ACL, 2020.
[11] K. Ong, R. Mao, D. Varshney and E. Cambria, “Robust ESG analysis against greenwashing risks using artificial intelligence,” arXiv preprint arXiv:2502.15821, 2025.
[12] Mogili, U., & Mohamed, A. (2023, November). Artificial intelligence and machine learning in the fields of education, medical, and smart phones. In AIP conference proceedings (Vol. 2917, No. 1, p. 050012). AIP Publishing LLC.
[13] Mogili, U., Ampolu, K. V., Rajasekharam, B., & Timothy, M. J. AI-Driven Interaction in AR Environments, in Journal of Digital Economy, 2024, Volume 3, Issue 1, pp. 228-234.
[14] Timothy, M. J., Rajasekharam, B., Ampolu, K. V., & Mogili, U. Threat Detection Using AI in Cybersecurity Systems, in IJIS, 2023, Volume 7, Issue 1, pp. 1-7.
[15] Ampolu, K.V., Mogili, U., Timothy, M. J., & Rajasekharam, B. Machine Learning Models for Predictive Maintenance, in IJIS, 2022, Volume 6, Issue 4, pp. 1-7.
[16] Rajasekharam, B., Timothy, M. J., Mogili, U., Ampolu, K.V., Machine Learning Models for Predictive Maintenance, in JDE, 2023, Volume 2, Issue 2, pp. 95-101.
[17] Soujania, B., Ampolu, K. V., Timothy, M. J., & Mogili, U. (2025) Classifying Disease Information Forums through Semantic Similarity-Based Machine Learning, Science, Technology and Development Journal, Volume XIV, Issue II, pp 67-75.
[18] B Satish Kumar, Kavitha C., Mogili, U.R., S. Pallam Shetty (2022). “Application of Machine Learning To Enhance the Performance of The Prophet Routing Protocol For Delay Tolerant Networks”. Journal for Basic Sciences, Volume 23, Issue 5, 2107-2116, DOI:10.37896/JBSV23.5/2278.
[19] Sree Geeta, Umamaheswararao Mogili. (2022), “Use of Several Machine Learning Algorithms for Effective Prediction of Cyberbullying”, International Journal of Creative Research Thoughts, Volume 10, Issue 6, pp 17.

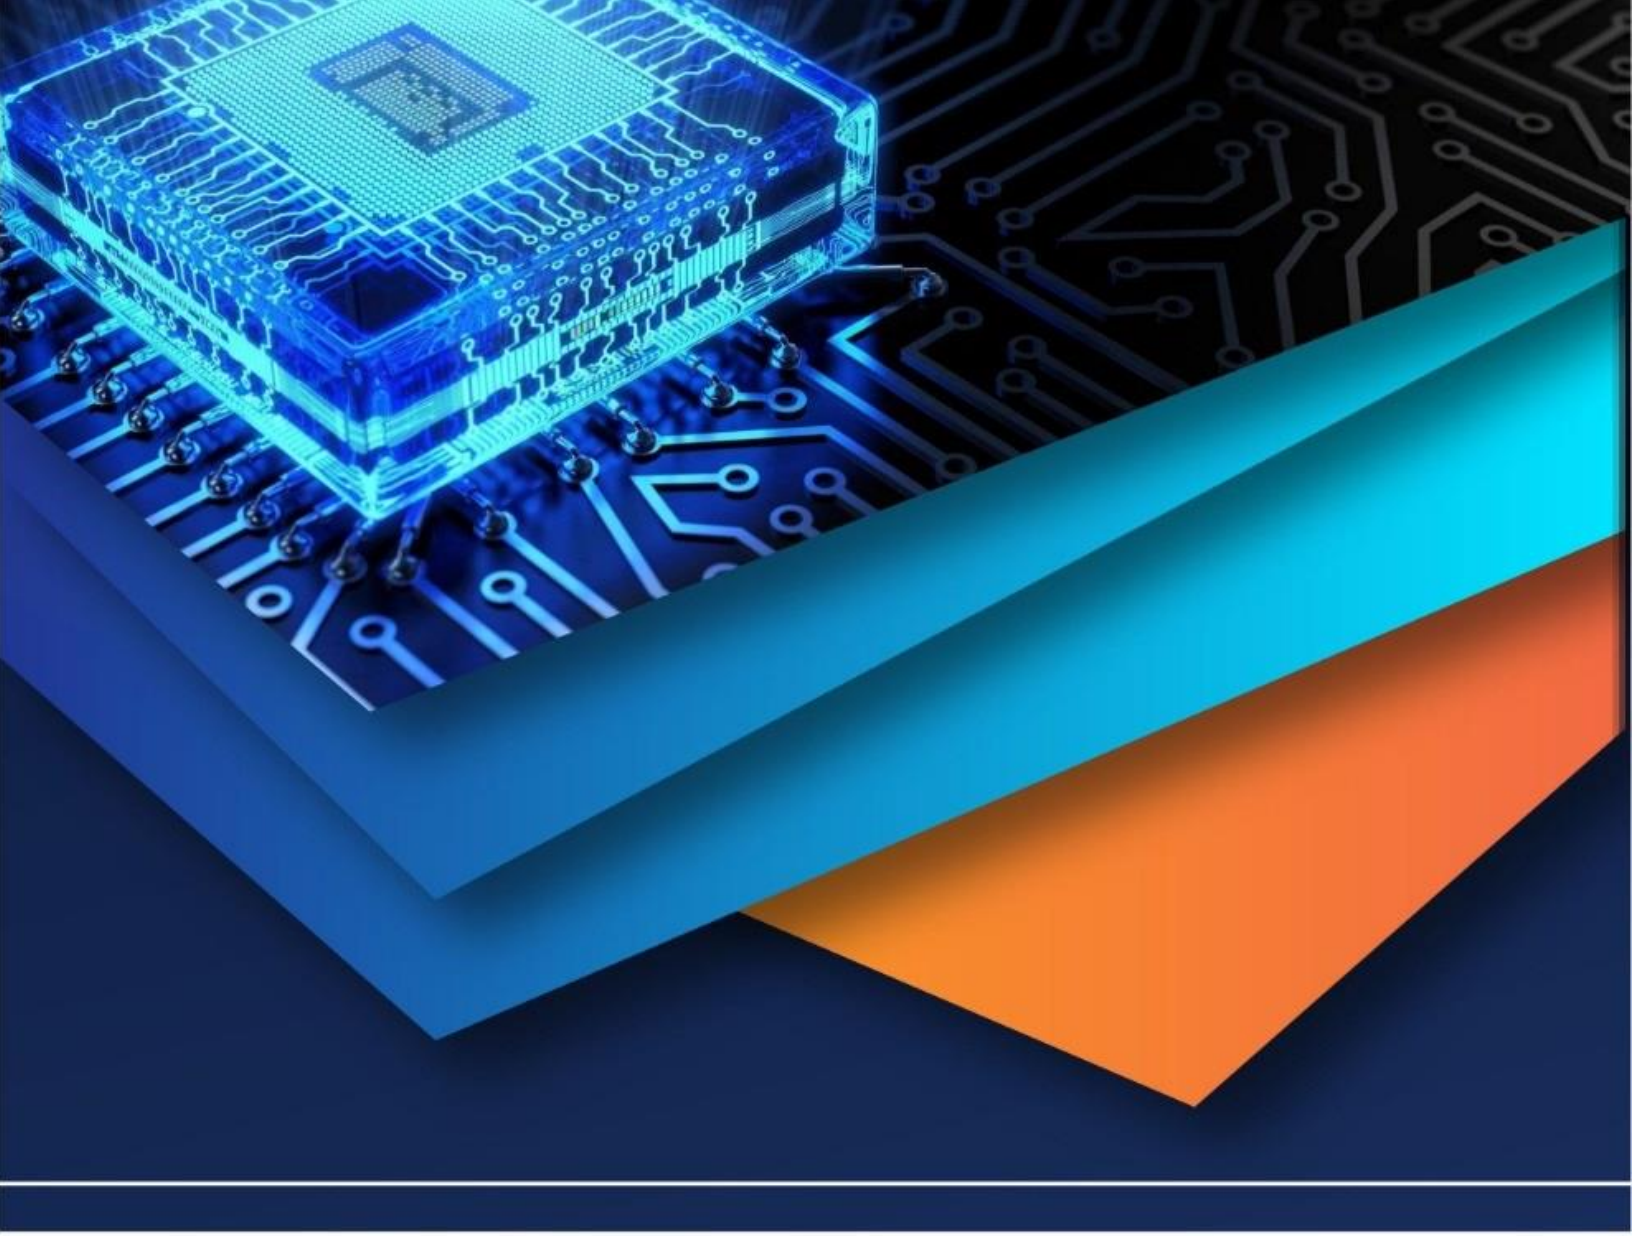

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY