

Bridging Monitoring and Mindset: A Retrieval-Augmented LLM Chatbot Framework for Real-Time Environmental Education

B.U.I Perera

Department of Electrical and Computer Engineering
The Open University of Sri Lanka
Nawala, Sri Lanka
uththaraisurindi@gmail.com

J.K. Epakanda

Department of Electrical and Computer Engineering
The Open University of Sri Lanka
Nawala, Sri Lanka
jepakanda@gmail.com

K.D.B. Aratthanage

Department of Electrical and Computer Engineering
The Open University of Sri Lanka
Nawala, Sri Lanka
kdara@ou.ac.lk

Abstract—Plastic pollution is a serious problem for our environment, and many public awareness campaigns are not able to change people's daily habits in a lasting way. At the same time, new Artificial Intelligence (AI) and smart sensor technologies can now monitor rivers, coasts, and ecosystems in real time, but this rich monitoring data rarely reaches ordinary citizens in a form they can understand and act on. This gap between “monitoring” (technical data collection) and “mindset” (public understanding and behavior) is the main motivation of this paper. We propose a conceptual framework named EcoLitLLM, a mobile-based chatbot that uses Retrieval-Augmented Generation (RAG) with a Large Language Model (LLM) to give simple, trustworthy, and location-based answers about plastic pollution and ecosystem health. The framework connects a verified environmental knowledge base with live AI/IoT ecosystem-monitoring feeds, so the chatbot's answers are grounded in current, local, and reliable evidence rather than only static content. The framework also supports offline use through an on-device small language model, multimodal litter recognition (multiple types of data) through the phone camera, and gamified citizen-science reporting that feeds new data back into the monitoring system. Since this is a conceptual paper, the methodology explains the framework-design process and gives logical pseudocode instead of an empirical evaluation. We also discuss expected outcomes, limitations, and directions for future testing of this framework in real communities.

Keywords—Retrieval-augmented generation(RAG), plastic pollution, ecosystem monitoring, environmental education, citizen science

I. INTRODUCTION

Plastic pollution is now one of the biggest environmental problems in the world. It affects rivers, oceans, soil, and even human health through microplastics. Governments, universities, and non-government organizations have tried many public awareness campaigns to reduce plastic use, using posters, television, and social media. However, most of these campaigns share information in one direction only. They give the same generic message to everyone, and they do not check whether people really understand the message or change their behavior afterward [1].

At the same time, a different research direction is growing very fast: Artificial Intelligence (AI) and smart technologies for monitoring ecosystem health. Sensors, drones, satellites,

and machine-learning models can now detect microplastics, classify waste, and track water quality almost in real time [13]–[21]. This monitoring data is very valuable, but it is mostly used by scientists, engineers, and policymakers. Ordinary citizens, students, and local communities rarely see this data, and even if they see it, the data is often too technical for them to understand.

This creates a clear disconnect. On one side, we have monitoring systems that know exactly how polluted a river is today. On the other side, we have education and awareness programs that speak in general terms and do not use this live data at all. Our motivation in this paper is to bridge these two sides using a conversational Artificial Intelligence system that ordinary people can use on their own mobile phone.

Large Language Models (LLMs) such as GPT-style models have shown that they can hold natural conversations and explain difficult topics in simple words. But LLMs by themselves can produce wrong information, called hallucination, especially about local or very recent facts [7], [8]. Retrieval-Augmented Generation (RAG) solves this problem partly by letting the LLM search a trusted knowledge source before answering [2]. Our idea is to extend RAG so that it retrieves not only static educational text, but also live monitoring data from AI/IoT ecosystem-monitoring systems. In this way, the chatbot becomes a bridge between the “monitoring” track and the “mindset” (education and awareness) track of environmental sustainability research.

This motivates the design question of this paper, How can a conversational AI framework ground its answers in both a verified educational knowledge base and live, hyperlocal ecosystem-monitoring data, while remaining usable offline in low-connectivity communities? We address this question with EcoLitLLM, a conceptual framework described in the contribution section below.

The main contributions of this paper are

- We identify and describe the research gap between digital awareness campaigns and AI-based ecosystem-monitoring systems.
- We propose EcoLitLLM, a conceptual framework for a mobile chatbot that grounds an LLM's answers in both a verified knowledge base and real-time monitoring data.

- We design a layered system architecture that supports offline use, multimodal litter recognition, hallucination control, and citizen-science feedback.
- We present logical pseudocode for the main algorithms of the framework, since this is a conceptual, design-stage contribution.

The rest of this paper is organized as follows. Section II reviews related work. Section III explains the research gap in more detail. Section IV presents the proposed conceptual framework. Section V describes the system architecture. Section VI explains the methodology and gives the pseudocode. Section VII discusses expected outcomes, Section VIII lists limitations, and Section IX concludes the paper.

II. RELATED WORK

A. LLM Chatbots and RAG in Education

Retrieval-Augmented Generation has become a popular way to reduce hallucination and to connect Large Language Models with domain-specific knowledge [1], [2]. Several studies have applied RAG chatbots in higher education, for example, to answer student questions using course material [3], to support database and information-system teaching [4], and to solve technical problems that appear when deploying LLM-based educational software at scale [5]. RAG chatbots have also been used to answer questions from medical textbooks by retrieving passages before generation [6]. These works show that RAG chatbots are practical and are already trusted in education, but almost none of them connect the retrieval layer to live, real-world sensor data; they retrieve from fixed documents only.

B. Hallucination, Trust and Explainability

Hallucination remains one of the biggest problems for any LLM-based system that must be trusted by the public [7], [8]. Adaptive retrieval methods have been proposed to detect when a model is uncertain and to trigger extra retrieval or verification steps before giving a final answer [9]. Beyond hallucination, researchers have also studied explainable AI (XAI) as a way of building user trust in AI systems that make important decisions [10], [11]. Miller [12] explains that good explanations must be built the way humans naturally explain things to each other, not only as technical outputs. These ideas are important for our framework, because an environmental chatbot must be both accurate and understandable to non-expert users.

C. AI for Plastic and Waste Detection

A large number of studies apply deep learning to detect and classify waste, including plastic, from images, mostly using Convolutional Neural Networks (CNNs) or YOLO-style object detectors [13], [16], [17]. Some works fuse features from several pre-trained deep models to improve waste-classification accuracy [14], while others use airborne hyperspectral imagery to detect garbage over wide areas [15]. Large multi-class datasets have also been proposed, such as TrashBox with 17,785 images in seven classes, on which transfer-learning models achieved up to 98.47% test accuracy [18]. Domain-adaptation techniques have been used to improve the detection of small floating objects on water surfaces in new environments [19]. These studies indicate that image-based waste recognition is technically feasible.

D. AI for Ecosystem Monitoring

Deep learning has been applied to a variety of environmental remote-sensing tasks, including land cover mapping, parameter retrieval, and data fusion, with applications in the atmosphere, vegetation, hydrology, and ocean color, and is reported to outperform traditional models [20]. In marine and coastal settings, deep-learning models have been applied to water-quality estimation and forecasting using satellite and in situ sensor or IoT data [21]. Ensemble learning, which combines multiple base models to improve accuracy and robustness, is widely used across machine-learning applications [22]. The reliability of these models depends on their data: quality problems such as missing, incorrect, or duplicate data reduce accuracy and can lead to unreliable predictions [23], and environmental deep-learning studies note limited labeled training data, noisy and incomplete multi-sensor observations, and low interpretability [20], [21]. Machine learning has also been used to predict microplastic transport in open channels from hydraulic, vegetation, and microplastic-density factors in laboratory experiments, using explainable AI to interpret the results, and the authors describe it as a promising tool for real-world predictions [24]. A review of deep learning for marine water quality reports that the vast majority of the works it surveyed focus on forecasting water parameters and that few include decision-making components, and it proposes reinforcement learning for decision-making as a future research direction [21].

E. Citizen Science and Gamification

Gamification elements such as badges and competition appear to motivate participants in online-gaming contexts, and gaming has been proposed as a way to engage younger and more diverse audiences [25]; an exploratory study of a gamified citizen-science app found that likely users were more motivated by badges and peer competition, although sustained engagement was not tested [26].

F. Awareness Campaigns and Behavior Change

Public awareness and education campaigns, from school programs to mass-media initiatives, are widely used to reduce plastic pollution [27], [28]. However, simply providing factual information is often insufficient to change behavior, and increased awareness has not been found to result in action [27], [28]. People also need to see personal relevance [27]. A narrative review of microplastic public-perception research reports mixed success for school programs, media campaigns, and eco-labels, and describes a gap between scientific knowledge, public behavior, and policy action [27]. The effect of campaigns on behavior is also often not evaluated rigorously [27]. Awareness differs across regions, which highlights the need to tailor communication to local contexts [27], and targeted interventions for different audiences are recommended by most of the studies reviewed in [28]. Simply educating people about the plastic crisis is not enough to change how they behave in real [28]. These findings suggest that a generic message is not enough, which motivates the personalized, location-aware, conversational framework proposed in this paper.

G. Positioning Against Recent Environmental RAG and IoT-LLM Systems

Recent LLM-environment systems each cover part of what EcoLitLLM proposes, but none combines all of it. MeteoChat [29] uses fine-tuning plus RAG to generate audience-adapted reports, but from archived (not live) data,

with no offline mode or citizen reporting. IoT-LLM [30] improves LLM reasoning over real-world sensor data through retrieval, but as a general benchmark framework, not a public chatbot. K  k et al. [31] propose a three-tier physical/edge-fog/cloud architecture for industrial predictive maintenance relevant to our edge-layer design (Section V, L3) but built for a different domain. ChatENV [32] jointly reasons over satellite imagery and sensor data for scenario simulation, but works on aerial imagery rather than citizen phone photos and isn't aimed at public education. EcoLitLLM is designed to be the first to bring live hyperlocal retrieval, offline inference, phone-camera litter recognition, and a consented citizen-reporting loop into one RAG-grounded chatbot, a claim not yet tested empirically, since this remains a conceptual paper.

III. RESEARCH GAP

Based on the related work above, we identify four specific gaps that motivate this paper

- Digital awareness campaigns are mostly one-way and generic. They broadcast the same message to everyone and do not personalize information to the user's location, language, or existing knowledge level, and they rarely measure real behavior change over time.
- Monitoring data stays inside technical silos. AI/IoT ecosystem-monitoring systems produce valuable, real-time data about pollution and ecosystem health, but this data is written for scientists and regulators, not for the general public.
- Most educational LLM chatbots are grounded only in static text. They do not retrieve from live monitoring feeds, so they cannot tell a user, for example, how polluted their own nearby river is today.
- Offline and low-connectivity access is not addressed. Many rural and coastal communities, including in Sri Lanka, do not have stable internet access, but almost all LLM chatbot designs assume constant cloud connectivity.

The proposed EcoLitLLM framework is designed to directly address these four gaps, as explained in the next sections.

IV. PROPOSED CONCEPTUAL FRAMEWORK

Fig. 1 shows the overall conceptual framework, which we call EcoLitLLM. On the left side (“Mindset”), community users such as students, fisherfolk, and coastal or rural residents interact with a mobile application through chat, camera, and location services. On the right side (“Monitoring”), AI/smart ecosystem-monitoring systems such as sensors, drones, and satellite feeds continuously produce data about water quality, litter density, and ecosystem indicators.

The two sides are connected through a shared Retrieval-Augmented Grounding Layer in the middle. This layer performs a vector search over two sources at the same time: a verified environmental knowledge base (curated facts, local waste-disposal rules, and curriculum-aligned content) and a live monitoring data feed (real-time readings from the monitoring systems). Before the final answer is generated, a hallucination filter checks whether the draft answer is actually supported by the retrieved evidence. If it is not well

supported, the system regenerates the answer in a stricter mode or reduces confidence and asks the user a clarifying question.

Two generation engines are available: a Cloud LLM for complex or nuanced questions, and an on-device Small Language Model (SLM) that can answer common questions and classify litter images even without an internet connection. This dual design directly answers the offline-access gap identified in Section III.

The final output combines a localized answer, a short micro-lesson, a gamified nudge (points or a small challenge), and, when relevant, a citizen-science prompt asking the user to report litter they see nearby. When the user submits such a report with their consent, it is fed back into the monitoring-data layer, closing the loop between public engagement and scientific monitoring. This feedback loop is the key novel element that lets “mindset” actions strengthen “monitoring” coverage, and “monitoring” evidence strengthen “mindset” education, at the same time.




Fig. 1. Conceptual framework showing how the mobile chatbot bridges the “mindset” side (public users, education) with the “monitoring” side (AI/smart ecosystem-monitoring systems), through a shared retrieval-augmented grounding layer.

V. SYSTEM ARCHITECTURE

To make the conceptual framework implementable, Fig. 2 breaks it down into six layers.

A. L1 – Presentation Layer

The chat interface, camera capture for litter photos, and a gamification dashboard that shows the user's points and local leaderboard position.

B. L2 – Application/Orchestration Layer

A query handler and session manager, a geolocation and localization service that finds the user's municipality, and a module that adapts the reading difficulty and language of the response for the user.

C. L3 – Edge Layer

An on-device Small Language Model that answers frequently-asked questions without internet access, and an on-device litter-image classifier (a lightweight CNN) that can run offline on a normal smartphone.

D. L4 – RAG / Grounding Layer

Embedding and vector-search components, a hallucination/fact-check filter, and a citation builder that

attaches the evidence source to every answer so the user can see where the information came from.

E. L5 – Cloud LLM Layer

A full cloud-hosted LLM for complex reasoning and a personalized learning-path generator that slowly increases the difficulty of content as the user learns more.

F. L6 – Data Sources Layer

The verified environmental knowledge base, the live AI/IoT ecosystem-monitoring feeds, and a citizen-submitted reports database. As shown by the red arrow in Fig. 2, citizen reports can enrich the monitoring feed for future queries, which is the same feedback loop described in Section IV.




Fig. 2. Layered system architecture of EcoLitLLM, from the mobile presentation layer (L1) down to the underlying data sources (L6).

VI. METHODOLOGY

Because this paper is at the conceptual and design stage, we do not present an empirical evaluation with real users or a working prototype. Instead, we follow a design-oriented research method with four steps: (1) identifying the research gap from the literature reviewed in Section II; (2) defining clear design objectives for a solution, namely grounding, locality, offline support, and citizen feedback; (3) constructing the conceptual framework and layered architecture presented in Sections IV and V; and (4) expressing the internal logic of the framework as structured pseudocode, so that other researchers or developers can later implement and test it. Below, we present this last step as four algorithms that describe EcoLitLLM's core reasoning.

Algorithm 1 describes how a user query is handled end-to-end, including the offline fallback and the hallucination check.

ALGORITHM 1: RAG-Grounded Query Handling

Input: user_query, user_location, connectivity_status
Output: verified_response

```

1: receive user_query from chat interface
2: detect_language(user_query)
3: if connectivity_status == OFFLINE then
4:   response <- EdgeSLM.answer(user_query)
5:   if confidence(response) < THRESHOLD then
6:     queue(user_query) for sync when online
7:     mark response as "may be outdated"
8:   end if
9: else
10:  embed <- Embed(user_query)
11:  docs <- VectorSearch(embed, KnowledgeBase,
    MonitoringFeed, user_location)
12:  draft <- CloudLLM.generate(user_query, docs)
13:  grounded <- HallucinationFilter(draft, docs)
14:  if grounded == TRUE then
15:    response <- AttachCitations(draft, docs)
16:  else
17:    response <- CloudLLM.regenerate(user_query,
    docs, strict_mode = TRUE)
18:  end if
19: end if
20: log_interaction(user_query, response, user_location)
21: return response

```

Algorithm 2 shows how the framework finds location-specific rules and monitoring data, so the same question can get a different, correct answer in different towns.

ALGORITHM 2: Hyperlocal Grounding

Input: user_query, user_location
Output: localized_context

```

1: municipality <- ReverseGeocode(user_location)
2: rules <- KnowledgeBase.lookup(municipality,
    "waste_disposal_rules")
3: monitoring <- MonitoringFeed.lookup(municipality,
    radius = 10km)
4: if rules is empty then
5:   rules <- KnowledgeBase.lookup(
    nearest_region(municipality))
6:   mark_as_approximate(rules)
7: end if
8: localized_context <- Merge(rules, monitoring)
9: return localized_context

```

Algorithm 3 decides whether the query should be sent to the cloud LLM or handled on the device, balancing accuracy, cost, and offline availability.

ALGORITHM 3: Offline / Edge Routing Decision

Input: connectivity_status, query_complexity_score
Output: routing_decision

```

1: if connectivity_status == ONLINE and
    query_complexity_score > LOW then
2:   routing_decision <- CLOUD_LLM

```

```

3: else if connectivity_status == ONLINE and
  query_complexity_score <= LOW then
4:   routing_decision <- EDGE_SLM // save cost
5: else
6:   routing_decision <- EDGE_SLM_OFFLINE_MODE
7: end if
8: return routing_decision

```

Algorithm 4 shows the multimodal path, where a photo of litter becomes both an educational response and a new citizen-science data point that can enrich the monitoring feed.

ALGORITHM 4: Multimodal Litter Recognition

Input: photo, user_location
Output: educational_response, citizen_report

```

1: features <- EdgeCNN.extract(photo)
2: plastic_type <- Classify(features)
3: recyclability <- KnowledgeBase.lookup(
  plastic_type, municipality)
4: educational_response <- Compose(plastic_type,
  recyclability, impact_fact)
5: citizen_report <- { photo_metadata, plastic_type,
  user_location, timestamp }
6: submit(citizen_report) to ReportsDatabase
7: if consent_given(user) then
8:   MonitoringFeed.update(citizen_report)
9: end if
10: return educational_response, citizen_report

```

Table I maps each objective to the architectural layer(s) in Section V that implement it and the algorithm(s) in this section that operationalize it. We treat a design objective as satisfied at the design-stage level defined below.

TABLE I. TRACEABILITY FROM RESEARCH GAP TO DESIGN OBJECTIVE, ARCHITECTURE AND ALGORITHM

Research Gap	Design objective	Architecture layers	Algorithm
Awareness campaigns are one-way/generic and don't personalise to the user's location	Locality	Application layer	Algorithm 2
Monitoring data stays inside technical silos, not reaching the public	Citizen feedback	Presentation Layer, Data Sources Layer	Algorithm 4

Educational chatbots are grounded only in static text, not live monitoring feeds	Grounding	RAG Layer	Algorithm 1
Offline and low-connectivity access is not addressed	Offline support	Edge Layer	Algorithms 1, 3, 4

VII. DISCUSSION

If EcoLitLLM is implemented and piloted, we expect several outcomes, based on patterns already reported in the literature reviewed in Section II. First, grounding answers in both a knowledge base and live monitoring data should reduce hallucination compared with a plain LLM chatbot, because every claim can be traced to a retrieved source [2], [9]. Second, the hyperlocal grounding logic (Algorithm 2) should make the chatbot more useful than generic awareness campaigns, because a user in one town can get different, correct disposal instructions from a user in another town, directly addressing research gap 1 from Section III. Third, the offline/edge routing (Algorithm 3) should make the tool usable in low-connectivity rural and coastal areas, which is important for a country like Sri Lanka where internet access is not always stable. Fourth, the citizen-report feedback loop (Algorithm 4) should slowly increase the coverage and freshness of the monitoring feed itself, turning ordinary users into informal citizen scientists, similar to the engagement effects reported for gamified citizen-science apps [26].

For future evaluation, we suggest three simple studies once a working prototype exists (i) an A/B test comparing knowledge retention between a static awareness poster/app and the EcoLitLLM chatbot; (ii) a hallucination audit where domain experts check a sample of chatbot answers against the retrieved evidence; and (iii) a longitudinal engagement study measuring how many users keep using the app and reporting litter after 30, 60, and 90 days. Answer quality would be assessed based on factual accuracy, citation precision, citation recall, and retrieval relevance, while system performance would be evaluated using response latency, offline model accuracy, battery and memory consumption on low-end smartphones. In addition, user impact would be assessed using task completion, learning gain, behavior-change indicators, user trust, and reporting retention.

VIII. LIMITATIONS

This paper presents a conceptual framework and a pseudocode algorithm. The expected outcomes presented in Section IV are not yet proven, as the system has not been implemented or tested with real-world users. The framework depends on user participation and the quality of local AI/IoT monitoring feeds, which may not be available in every area, such as small or rural towns. Due to the memory and battery limitations of low-end phones, running a small language model and an image classifier fully offline may also be difficult. In addition, this framework does not address multilingual support for Sinhala and Tamil. It also does not address the privacy concerns related to geolocation data collection, which must be handled carefully through clear user consent.

IX. CONCLUSION AND FUTURE WORK

The paper proposed a conceptual framework for a mobile, Retrieval-Augmented LLM chatbot that links public environmental education with real-time AI/IoT ecosystem-monitoring data. The framework grounds answers in a verified knowledge base and live local monitoring feeds. It also supports offline use, litter recognition through the phone camera, and a consent-based citizen-science loop that feeds validated reports back into monitoring. The main contribution of the framework is the bidirectional loop. Live monitoring data informs and grounds the educational responses, while validated citizen observations are incorporated into future monitoring. The framework is designed to address the gaps identified between awareness campaigns and monitoring systems, but it has not yet been implemented or tested, so its expected benefits remain hypotheses. Future work will build a prototype with multilingual support, connect it to a real local monitoring source, and run a pilot study in a coastal or rural community in Sri Lanka to evaluate knowledge retention, answer accuracy, and long-term engagement.

REFERENCES

- [1] W. Fan, Y. Ding, L. Ning, S. Wang, H. Li, D. Yin, T.-S. Chua, and Q. Li, "A Survey on RAG Meeting LLMs: Towards Retrieval-Augmented Large Language Models," in *Proc. 30th ACM SIGKDD Conf. Knowledge Discovery and Data Mining (KDD)*, 2024, pp. 6491–6501, doi: 10.1145/3637528.3671470.
- [2] S. Barnett, S. Kurniawan, S. Thudumu, Z. Brannelly, and M. Abdelrazek, "Seven Failure Points When Engineering a Retrieval Augmented Generation System," in *Proc. IEEE/ACM 3rd Int. Conf. AI Engineering – Software Engineering for AI (CAIN)*, 2024, pp. 194–199, doi: 10.1145/3644815.3644945.
- [3] M. Maryamah, M. M. Irfani, E. B. Tri Raharjo, N. A. Rahmi, M. Ghani, and I. K. Raharjana, "Chatbots in Academia: A Retrieval-Augmented Generation Approach for Improved Efficient Information Access," in *Proc. 16th Int. Conf. Knowledge and Smart Technology (KST)*, IEEE, 2024, pp. 259–264, doi: 10.1109/KST61284.2024.10499652.
- [4] A. Neumann, Y. Yin, S. Sowe, S. Decker, and M. Jarke, "An LLM-Driven Chatbot in Higher Education for Databases and Information Systems," *IEEE Transactions on Education*, vol. 68, pp. 103–116, 2025, doi: 10.1109/TE.2024.3467912.
- [5] N. Chondamrongkul, G. Hristov, and P. Temdee, "Addressing Technical Challenges in Large Language Model-Driven Educational Software System," *IEEE Access*, vol. 13, pp. 12846–12858, 2025, doi: 10.1109/ACCESS.2025.3531380.
- [6] S. Bhavani Peri, S. Santhanalakshmi, and R. Radha, "Chatbot to Chat with Medical Books Using Retrieval-Augmented Generation Model," in *Proc. IEEE NKSS Int. Conf. Digital Transformation (NKCon)*, 2024, doi: 10.1109/NKCon62728.2024.10774900.
- [7] L. Huang, W. Yu, W. Ma, W. Zhong, Z. Feng, H. Wang, Q. Chen, W. Peng, X. Feng, B. Qin, and T. Liu, "A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions," *ACM Transactions on Information Systems*, vol. 43, no. 2, pp. 1–58, 2024, doi: 10.1145/3703155.
- [8] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y. Bang, D. Chung, and P. Fung, "Survey of Hallucination in Natural Language Generation," *ACM Computing Surveys*, vol. 55, no. 12, pp. 1–38, 2023, doi: 10.1145/3571730.
- [9] H. Ding, L. Pang, Z. Wei, H. Shen, and X. Cheng, "Rowen: Adaptive Retrieval-Augmented Generation for Hallucination Mitigation in LLMs," in *Proc. ACM SIGIR Conf. Research and Development in Information Retrieval, Asia Pacific Region (SIGIR-AP)*, 2025.
- [10] A. Adadi and M. Berrada, "Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)," *IEEE Access*, vol. 6, pp. 52138–52160, 2018, doi: 10.1109/ACCESS.2018.2870052.
- [11] V. Chamola, V. Hassija, A. R. Sulthana, D. Ghosh, D. Dhingra, and B. Sikdar, "A Review of Trustworthy and Explainable Artificial Intelligence (XAI)," *IEEE Access*, vol. 11, pp. 78994–79015, 2023, doi: 10.1109/ACCESS.2023.3294569.
- [12] T. Miller, "Explanation in Artificial Intelligence: Insights from the Social Sciences," *Artificial Intelligence*, vol. 267, pp. 1–38, 2019, doi: 10.1016/j.artint.2018.07.007.
- [13] H. Abdu and M. H. Mohd Noor, "A survey on waste detection and classification using deep learning," *IEEE Access*, vol. 10, pp. 128151–128165, 2022, doi: 10.1109/ACCESS.2022.3226682.
- [14] K. Ahmad, K. Khan, and A. Al-Fuqaha, "Intelligent fusion of deep features for improved waste classification," *IEEE Access*, vol. 8, pp. 96495–96504, 2020, doi: 10.1109/ACCESS.2020.2995681.
- [15] D. Zeng, S. Zhang, F. Chen, and Y. Wang, "Multi-scale CNN based garbage detection of airborne hyperspectral data," *IEEE Access*, vol. 7, pp. 104514–104527, 2019, doi: 10.1109/ACCESS.2019.2932117.
- [16] A. S. Kumar, A. Mathai, R. Dinesh, H. Mohan, and D. S., "Waste detection and classification using YOLO algorithm and sensors," in *Proc. 16th Int. Conf. Sens. Technol. (ICST)*, 2023, pp. 1–6, doi: 10.1109/ICST59744.2023.10460825.
- [17] M. G. Vimal Kumar, M. Kumar, K. Narayana Rao, P. Syamala Rao, A. Tirumala, and E. Patnala, "Advanced YOLO-based trash classification and recycling assistant for enhanced waste management and sustainability," in *Proc. 2nd Int. Conf. Intell. Cyber Phys. Syst. Internet Things (ICoCI)*, Sep. 2024, pp. 726–733.
- [18] N. V. Kumsetty, A. B. Nekkare, S. Kamath S., and A. Kumar M., "TrashBox: Trash detection and classification using quantum transfer learning," in *Proc. 31st Conf. Open Innov. Assoc. (FRUCT)*, Apr. 2022, pp. 125–130, doi: 10.23919/FRUCT54823.2022.9770922.
- [19] R. Chen, J. Wu, Y. Peng, Z. Li, and H. Shang, "Solving floating pollution with deep learning: A novel SSD for floating objects based on continual unsupervised domain adaptation," *Eng. Appl. Artif. Intell.*, vol. 120, Art. no. 105857, 2023, doi: 10.1016/j.engappai.2023.105857.
- [20] Q. Yuan, H. Shen, T. Li, Z. Li, S. Li, Y. Jiang, H. Xu, W. Tan, Q. Yang, J. Wang, J. Gao, and L. Zhang, "Deep learning in environmental remote sensing: Achievements and challenges," *Remote Sens. Environ.*, vol. 241, art. 111716, 2020, doi: 10.1016/j.rse.2020.111716.
- [21] A. F. Gambin, E. Angelats, J. S. González, M. Miozzo, and P. Dini, "Sustainable Marine Ecosystems: Deep Learning for Water Quality Assessment and Forecasting," *IEEE Access*, vol. 9, pp. 121344–121365, 2021, doi: 10.1109/ACCESS.2021.3109216.
- [22] I. D. Mienye and Y. Sun, "A survey of ensemble learning: Concepts, algorithms, applications, and prospects," *IEEE Access*, vol. 10, pp. 99129–99149, 2022, doi: 10.1109/ACCESS.2022.3207287.
- [23] Y. Gong, G. Liu, Y. Xue, R. Li, and L. Meng, "A survey on dataset quality in machine learning," *Inf. Softw. Technol.*, vol. 162, art. 107268, 2023, doi: 10.1016/j.infsof.2023.107268.
- [24] A. Z. Fazil, P. I. A. Gomes, and R. M. K. Sandamal, "Applicability of machine learning techniques to analyze microplastic transportation in open channels with different hydro-environmental factors," *Environ. Pollut.*, vol. 357, art. 124389, 2024, doi: 10.1016/j.envpol.2024.124389.
- [25] G. Newman, A. Wiggins, A. Crall, E. Graham, S. Newman, and K. Crowston, "The future of citizen science: Emerging technologies and shifting paradigms," *Front. Ecol. Environ.*, vol. 10, no. 6, pp. 298–304, 2012, doi: 10.1890/110294.
- [26] A. Bowser, D. Hansen, Y. He, C. Boston, M. Reid, L. Gunnell, and J. Preece, "Using Gamification to Inspire New Citizen Science Volunteers," in *Proc. 1st Int. Conf. Gameful Design, Research, and Applications*, ACM, 2013.
- [27] J. Khorsandi, L. Blank, K. Momchilov, M. Dagovetz, and K. Batra, "From ocean to table: How public awareness shapes the fight against microplastic pollution," *Urban Sci.*, vol. 9, no. 10, art. no. 418, Oct. 2025, doi: 10.3390/urbansci9100418.
- [28] A. MacDonald, D. Allen, L. Williams, P. Flowers, and T. R. Walker, "People, plastic, and behaviour change: A comment on drivers of plastic pollution, barriers to change and targeted behaviour change interventions," *Environ. Sci.: Adv.*, vol. 2, no. 4, pp. 551–557, 2023, doi: 10.1039/d2va00248e.
- [29] A. Lo Duca, R. Lo Duca, A. Marinelli, D. Occhiuto, and A. Scariot, "Semi-automated reporting from environmental monitoring data using a large language model-based chatbot," *ISPRS Int. J. Geo-Inf.*, vol. 15, no. 2, Art. no. 80, Feb. 2026, doi: 10.3390/ijgi15020080.
- [30] T. An, Y. Zhou, H. Zou, and J. Yang, "IoT-LLM: A framework for enhancing large language model reasoning from real-world sensor data," *Patterns*, to be published. [Online]. Available: <https://arxiv.org/abs/2410.02429>.
- [31] İ. Kök, O. Demirci, and S. Özdemir, "When IoT meet LLMs: Applications and challenges," in *Proc. IEEE Int. Conf. Big Data (BigData)*, Washington, DC, USA, 2024, pp. 7075–7084, doi: 10.1109/BigData62323.2024.10825187.
- [32] H. Elgendy, A. Sharshar, A. Aboeitta, and M. Guizani, "ChatENV: An interactive vision-language model for sensor-guided environmental monitoring and scenario simulation," *IEEE Trans. Geosci. Remote Sens.*, vol. 64, Art. no. 4703710, 2026, doi: 10.1109/TGRS.2026.3685864.