Artificial Intelligence for Sustainable Cities | Volume 1 (2026) | ISSN 2979-1464

CARAVEL

Publish open access at Caravel Press

Artificial Intelligence for Sustainable Cities

Journal homepage: aiforsustainablecities.com

Technical article

Autonomous Coastal Waste Detection and Localization Using UAV

Eyüp Koyun [1]*, Nur Hazal Çufalcı [1], Berke Akdoğan [1] and Aytaç Uğur Yerden [2]

[1] Istanbul Gedik University, Faculty of Engineering, Department of Computer Engineering, Istanbul, Türkiye.
[2] Istanbul Gedik University, Faculty of Engineering, Department of Industrial Engineering, Istanbul, Türkiye.

ABSTRACT

Waste that accumulates along coastlines—particularly in areas that are difficult for people to access—and is predominantly composed of plastic derivatives constitutes an environmental problem that threatens marine ecosystems, especially marine life, and concerns all countries, whether they have a coastline or not. This study aims to develop a ground-station-based, efficient waste detection and localization project using low-cost, easily accessible commercial unmanned aerial vehicles (UAVs). As the methodological approach, the attention-based and up-to-date deep learning architecture YOLOv12, which has been proven to be successful in object detection, was employed for detecting coastal waste. A robust training dataset was constructed using a hybrid dataset formed from coastal waste images captured over different surface types along the coastlines of Istanbul, together with open datasets. The main novelty of the system is the "VideoLog Synchronization Module," which precisely matches the video stream with the flight data recorded by the drone. Through this module, each frame and the waste contained within it are labeled with geographic coordinates and made observable via an ASP.NET Core MVC–based interface. Results from tests conducted through drone flights indicate that the system detects small waste objects in complex backgrounds with high accuracy (mAP50: 0.783) and performs localization with a very high level of success. This study aims to enable municipalities, non-governmental organizations, as well as individuals and institutions involved in coastal waste detection and cleanup activities, to shift human effort from detection to collection, thereby saving labor and time.

ARTICLE INFO

Article History:
Received: 14 March 2026
Revised: 21 April 2026
Accepted: 07 May 2026

Keywords:
Coastal Waste Detection
Offboard Processing
YOLOv12 Object Detection
Geospatial Mapping
Artificial intelligence

Article Citation:
Koyun, E., Çufalcı, N. H., Akdoğan, B., & Yerden, A. U. (2026). Autonomous Coastal Waste Detection and Localization Using UAV. Artificial Intelligence for Sustainable Cities, 1(1), 148–162. https://doi.org/10.65582/aifsc.2026.010

1. INTRODUCTION

The primary objective of this study is to detect coastal waste—which is predominantly composed of solid plastics, negatively affects marine ecosystems and indirectly human health, and concerns all countries regardless of whether they have a coastline or not—using low-cost UAVs, and to serve as a resource for future cleanup activities. The aim is to present the matched waste-frame outputs, obtained by mapping the detected coastal waste according to their geographic locations, to the use of relevant individuals and institutions, and to transform this data into an effective resource for processes such as coastal waste management and cleanup

* Corresponding author. Email address: eyupkoyuncr@gmail.com (Eyup Koyun)

planning.

Traditional coastal waste detection and collection methods have not yet kept pace with the digitalization process and still largely rely on human labor. These methods remain insufficient in terms of time and efficiency compared to digital solutions when detection activities must be conducted over large areas, highlighting the need for faster and more efficient methods (Marin et al. 2024). The locations of waste accumulation along coasts are dispersed randomly by waves and wind; therefore, while manual field inspections are inadequate for monitoring pollution rates and waste, UAV-based systems provide systematic and scalable monitoring capabilities (Merlino et al. 2020). This study aims to present a digital solution by integrating UAVs and deep learning models such as YOLO in areas where manual detection methods are insufficient. Although satellite-based waste detection studies offer a strong alternative with their ability to provide rapid results over large areas, they have significant issues such as the inability to detect waste smaller than the pixel size. UAVs offer higher resolution capabilities compared to satellite imagery through low-altitude flights, making them highly effective for detecting macro-litter categories like plastic bottles and cans. However, detecting small fragments (< 4 cm) remains a challenge compared to manual surveys (Marin et al. 2024), though UAVs are increasingly becoming a primary priority for large-scale monitoring due to their efficiency (Pfeiffer et al. 2023).

In this study, instead of the onboard architecture commonly used in the literature for UAV-based waste detection—where waste is detected using powerful embedded processors mounted on the UAV—an offboard architecture was preferred, in which waste is detected by processing video and flight data recorded by the UAV at a ground station. It has been stated in the literature that the onboard waste processing architecture restricts flight durations due to power and weight constraints (Kraft et al. 2021); however, this limitation can be minimized through the use of an offboard architecture facilitated by high-speed communication networks (Duangsuwan & Prapruetdee 2024).

Another fundamental objective of the project is to process the waste detected by the model on the UAV-recorded video not only visually but also along with its geographic location, answering the question of where a waste item was detected rather than just detecting it, thereby creating an opportunity to take faster action in future cleanup plans since the location of the waste is known. By synchronizing the raw video data obtained from the UAV with flight logs, each coastal waste item detected by the model was precisely matched geospatially; if a frame contains waste, that frame and its corresponding location data were uploaded to Cloudinary as metadata, and the results were visually presented on a web-based dashboard. Through this approach, an easy-to-use web-based dashboard was developed for municipalities and non-governmental organizations to serve as a supportive resource for cleaning planning. This aligns with the urgent need for targeted management strategies and decision-support data emphasized in recent coastal monitoring studies (Terzi et al. 2025; Duangsuwan & Prapruetdee 2024).

In the literature, accurate UAV-based mapping is often associated with high-cost RTK-GPS modules (Leira et al. 2015) or involves time-consuming manual annotation processes conducted after flights (Marin et al. 2024). However, studies demonstrate that low-cost GPS sensors can also yield acceptable accuracy for surveillance applications (Leira et al. 2015). The "Video–Log Synchronization Module" developed within the scope of this study provides a cost-effective and software-based solution to this need.

This study is also fully aligned with Türkiye's National Artificial Intelligence Strategy (2021–2025) (Presidency of the Republic of Türkiye Digital Transformation Office 2021), 2030 Industry and Technology Strategy (Republic of Türkiye Ministry of Industry and Technology 2023), and 12th Development Plan (2024–2028) (Presidency of the Republic of Türkiye Strategy and Budget Directorate 2024), which emphasize the development of domestic and intelligent technologies.

The data processing and analysis workflow of the system is explained through the diagram shown in Figure 1.

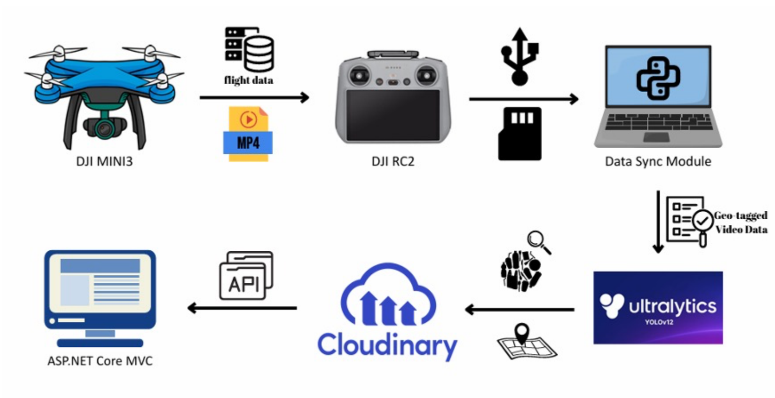

Figure 1. Architecture of the coastal waste detection and mapping system.

This study hypothesizes that integrating low-cost UAVs with YOLO-based detection and GPS synchronization can achieve operationally acceptable accuracy for coastal waste mapping. Instead of proposing a fundamentally new computer vision algorithm, this study aims to solve the severe operational and hardware constraints (e.g., rapid battery drain, heavy edge-computing payloads, and expensive RTK-GPS requirements) associated with UAV-based coastal monitoring. By utilizing state-of-the-art tools such as YOLOv12 and hybrid datasets as foundational components, the main methodological and architectural contributions of this paper are strictly defined as follows:

- Development of a Novel Video–Log Synchronization Module: A software-based georeferencing methodology was developed to mathematically synchronize 30 fps high-resolution optical video streams with 5 Hz UAV telemetry (SRT) logs. This module eliminates the need for expensive hardware modifications or onboard processing, accurately translating pixel coordinates of detected waste into real-world geographic coordinates (latitude, longitude, altitude).
- Design of a Sustainable Offboard Processing Architecture: A novel system-level architecture is proposed that deliberately decouples deep learning inference from the UAV payload. By processing the synchronized data at the ground station, the system successfully bypasses the rapid battery drain problem caused by onboard processors, enabling continuous coastal scanning (e.g., up to 38 minutes of flight time with Intelligent Flight Battery Plus) while still leveraging computationally heavy models like YOLOv12.
- Macro-Level Heatmap Generation for Human-Centric Cleanup: The synchronized detections are autonomously transformed into interactive geospatial heatmaps. Rather than targeting millimeter-level accuracy for robotic arms, this scalable engineering approach provides 1.5–2.5 m positional accuracy, which is statistically and practically proven to be perfectly adequate for guiding municipal human cleanup crews to coastal waste accumulation hotspots.

2. GENERAL INFORMATION (LITERATURE REVIEW)

This section examines existing studies in the literature on the detection of human-induced solid waste in coastal and marine ecosystems using digital methods; traditional approaches, namely satellite-based remote sensing systems and UAV-based waste detection systems, are comparatively reviewed in terms of technical, practical, and economic aspects. In particular, onboard architectures that use high-performance processors mounted on UAVs and offboard architectures that use UAVs solely as data collection tools are analyzed with respect to power consumption per flight, flight endurance, computational capacity, and scalability.

2.1. MANUAL AND TRADITIONAL DETECTION METHODS

The continuous degradation of marine and coastal ecosystems caused by human-induced waste is progressing toward an irreversible point on a global scale. Studies have shown that more than 80% of coastal pollution consists of land-based plastics (Terzi et al. 2025). Solid plastic waste, which constitutes the primary source of microplastic formation and makes up a large portion of coastal waste, persists in nature for many years and can no longer be considered a problem that can be left unattended. It is estimated that more than 8 million tons of plastic enter the oceans each year, and that in the near future the ratio of plastic to fish populations in the seas will reach 1:1 (Panwar et al. 2020). This situation clearly demonstrates how important the periodic detection and removal of coastal waste is.

2.2. SATELLITE-BASED REMOTE SENSING SYSTEMS

In traditional coastal waste cleanup methods, intervention generally occurs only after the amount of waste becomes visibly noticeable; however, by that time, plastics that have already mixed into seawater can no longer be collected. Manual cleanup activities carried out by municipal workers or volunteers from non-governmental organizations remain insufficient and inefficient in terms of both time and labor when scanning coastline lengths that extend for kilometers. In addition, methods based on human visual inspection are naturally inadequate for detecting waste in hard-to-reach areas such as rocky regions or locations with dense vegetation (Marin et al. 2024).

Technologies capable of remotely sensing coastlines at a macro scale are considered strategically important due to their advantage of covering large areas. These systems collect data by analyzing the electromagnetic spectrum through multispectral sensors integrated into orbiting satellites, and the acquired images are processed using deep learning architectures such as Convolutional Neural Networks (CNN) or Vision Transformers (ViT) to detect waste. Despite advantages such as providing ease of periodic monitoring at a global scale, the high cost of high-resolution datasets and the relatively large pixel size make the detection of small debris difficult. As sensor altitude increases, the detection of small waste becomes significantly more challenging (Pfeiffer et al. 2023; Duangsuwan & Prapruetdee 2024). Compared to satellite systems, waste detection systems using UAVs have been observed to be more successful in detecting small waste items, providing high-resolution data in a significantly shorter time (Marin et al. 2024).

2.3. UAV-BASED WASTE DETECTION: ONBOARD VS. OFFBOARD ARCHITECTURES

In the literature, many UAV-based solutions focus on performing waste detection onboard using embedded AI computing platforms such as the NVIDIA Jetson family (e.g., Xavier NX). However, although this "edge computing" approach enables real-time waste detection, it increases hardware weight and power consumption, which has been reported to shorten flight times and affect battery life efficiency (Kraft et al. 2021). Furthermore, the relatively small size of objects in UAV imagery requires high-resolution processing strategies (such as image tiling), which further increases computational cost and power consumption (Kraft et al. 2021).

Offboard waste detection using UAVs refers to an architecture in which UAVs act only as data collection platforms without requiring additional processing power beyond standard components such as a camera and a GPS module, while waste detection and location matching are carried out on high-performance computers at a ground station or in cloud-based environments. Kraft et al. (2021) acknowledge that high-accuracy object detection imposes significant computational and power demands, which challenges lightweight UAVs. However, they demonstrated that onboard detection is feasible by using specialized embedded platforms (e.g., NVIDIA Xavier NX) and optimized models like YOLOv4.

The main conclusion drawn from offboard waste detection studies is that removing high-performance processors and hardware load from the UAV allows single-charge flight durations to be extended from 15–20 minutes to nearly one hour—38 minutes in our case using the DJI Mini 3—and enables complex deep learning models, which are difficult to run onboard due to hardware limitations, to be executed with higher accuracy using the advanced computational capacity of the ground station. The most important advantage of this architecture is that high-performance, academic-level analysis can be conducted even with low-cost, standard commercial UAVs without requiring custom hardware.

On the other hand, a disadvantage is the inability to obtain real-time results. However, even if waste cannot be detected instantly, waste items and their locations can be accessed shortly after the UAV's video and flight data are transferred to the ground station. The literature on UAV-based waste mapping shows that precise positioning generally requires expensive RTK-GPS hardware or involves time-consuming manual analysis processes (Leira et al. 2015). The "Video–Log Synchronization Module" presented in this study overcomes this cost barrier by processing telemetry data provided by standard commercial UAVs through a software-based approach.

In addition, centralized data processing at the ground station facilitates the implementation of data privacy and cybersecurity protocols. This architecture aligns with recent secure system design methodologies (xT-

STRIDE) which emphasize that protecting sensitive sensor data against cyber threats is essential for the reliability of UAV operations (Yerden et al. 2025).

3. MATERIALS AND METHODS

In this section, the data collection, data processing, model training, synchronization, localization, and visualization stages of the developed coastal waste detection and mapping system are explained within the framework of experimental and applied methods. The process was designed as an end-to-end pipeline, and each step was structured to directly provide input to the next stage. The project flowchart showing the stages of the study is given in Figure 2.

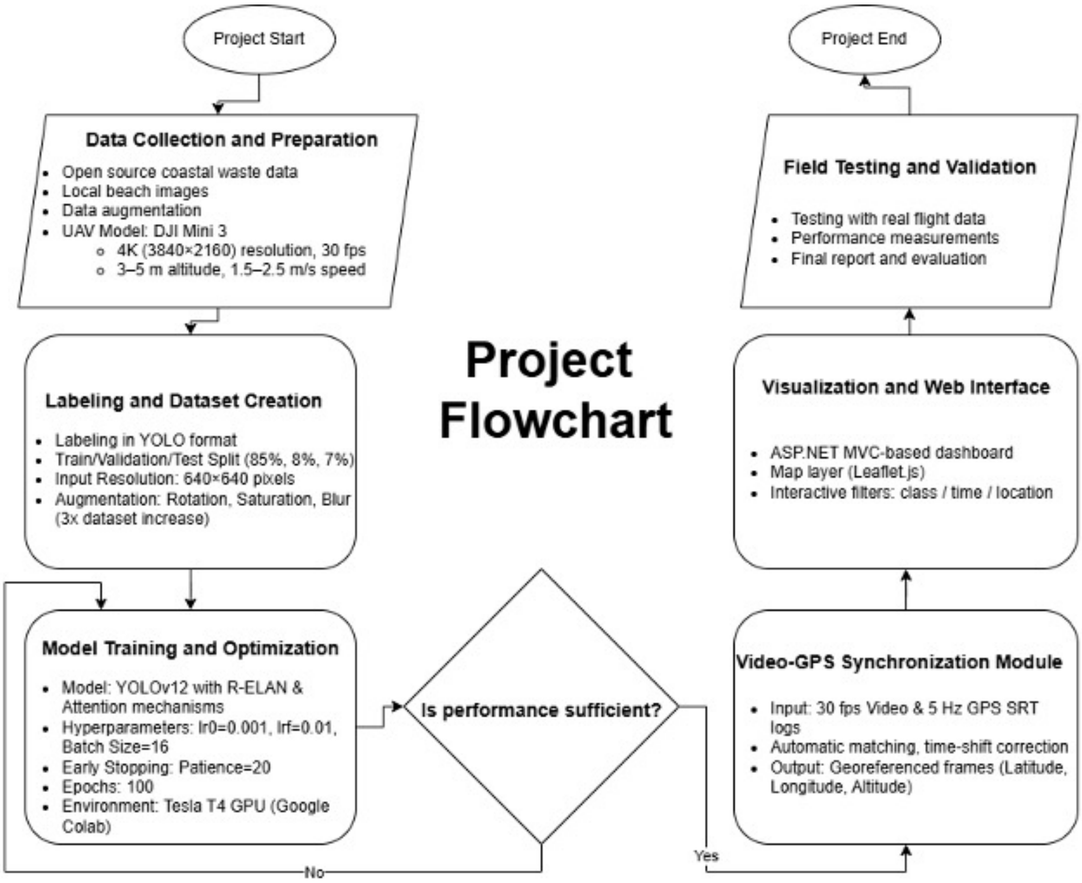

Figure 2. End-to-end system workflow diagram.

3.1. SYSTEM DESCRIPTION

To ensure the reproducibility of the study and address the limitations of treating the UAV as a “black box”, the hardware specifications and operational flight parameters were strictly documented. Data acquisition was performed using a DJI Mini 3 equipped with a 1/1.3-inch CMOS sensor (f/1.7 aperture, 82.1° FOV). To prevent spatial deformations at the image edges, the camera’s built-in lens distortion correction was applied. Video footage was captured at a high resolution of 4K (3840×2160 pixels) with an image capture rate of 30 fps.

During the field missions, flights were conducted at a constant low altitude of 3–5 meters above the coast surface. Flight log analysis reveals that while the UAV reached a maximum transit speed of 38 km/h, the average scanning speed during waste recording was maintained at 6.4 km/h (approx. 1.77 m/s). This operational speed involved a combination of slow cruising and stationary hovering to ensure image clarity and prevent motion blur under light wind conditions (average wind speed of 2.5 km/h). Furthermore, geospatial telemetry data (latitude, longitude, altitude) was extracted from the SRT logs with a GPS sampling rate of 5 Hz, which was subsequently synchronized with the 30 fps video stream.

3.2. DATASET CREATION

During the dataset creation stage, a hybrid dataset approach was adopted. The dataset used was created from two main sources:

1. UAVVaste Dataset: The open-source UAVVaste dataset, which includes different types of solid waste recorded by UAVs and produced by the authors of the study in the literature, was used as the main reference data source. This dataset was published on 14.11.2020 and includes a total of 772 images and 3,716 annotations belonging to these images. The annotations were labeled in bounding box format compatible with YOLO models.
2. Original Images (≈300 images): Approximately 300 original images were added to the dataset. These images were obtained along the Istanbul Maltepe coastline using a tripod, paying attention to an altitude suitable for waste detection by UAVs (3 – 5 m), and captured at 90-degree angles as if they were taken by a UAV. These images were collected under different weather conditions, locations, and lighting conditions in order to reflect real field conditions. Examples of the original images are shown in Figure 3.

(a) (b)

Figure 3. Original data samples. **(a)** Original data sample 1; **(b)** Original data sample 2.

To ensure high geographical diversity and prevent spatial overfitting, a hybrid dataset approach was utilized. The model was trained using the comprehensive UAVVaste dataset (Kraft et al. 2021), which includes diverse coastal topologies, alongside the 300 custom images captured in Istanbul. This combination guarantees that the YOLOv12 model can generalize across different beach compositions, sand colors, and waste distributions, ensuring its applicability beyond a single geographical location.

To provide full transparency on the dataset statistics, the combined hybrid dataset consists of 1,072 original images (772 from the UAVVaste dataset containing 3,716 annotations, and 300 custom images). The dataset was unified under a single class (‘rubbish’ or ‘coastal waste’) to maximize spatial detection capability rather than material-specific sorting. By combining these two sources, the generalization capability of the model was increased, and a more robust dataset was created, preventing dependence solely on open datasets (Zhu et al. 2023; Kraft et al. 2021).

3.3. DATA PREPROCESSING AND LABELING

During the data preprocessing and labeling stage, all collected images were labeled in accordance with the YOLO format. The labeling process followed the steps of normalizing litter image resolutions, eliminating frames containing unnecessary background or blur, and checking the compatibility of bounding box markings with waste size. Labeling was carried out using the Roboflow platform. The dataset was split into train/validation/test sets.

To ensure the dataset volume is sufficient for training a deep architecture like YOLOv12 and to prevent overfitting, rigorous data augmentation techniques were applied during the preprocessing stage. All images were auto-oriented and resized to fit a 640×640 pixel resolution (with black edges to preserve the aspect ratio). The dataset size was artificially tripled (producing 3 outputs per training example) by applying the following augmentations: horizontal flipping, random rotation between -7° and +7°, saturation variation between -7% and +7%, brightness variation between -7% and +7%, exposure adjustments between -1% and +1%, and artificial Gaussian blur of up to 0.1 pixels. The use of advanced data augmentation techniques (e.g., Mosaic, Mixup) increases the model’s robustness against diverse environmental conditions and provides noticeable improvements in mAP values (Bochkovskiy et al. 2020; Tian et al. 2025).

The 300 custom images were acquired using a tripod-mounted camera to establish a highly controlled baseline for the dataset. To ensure these images closely represent true UAV flight conditions, strict operational similarities were maintained. Both the tripod and the UAV utilized 4K resolution sensors, and the tripod height was specifically adjusted to 3–5 meters to perfectly match the drone's operational scanning altitude. Furthermore, because the drone is operated at a very low cruising speed (1.5–2.5 m/s) with a 3-axis mechanical gimbal, real-world motion blur and mechanical vibrations are inherently kept to a minimum. To account for any residual motion blur that might occur during actual flight, an artificial Gaussian blur of up to 0.1 pixels was deliberately introduced during the data augmentation phase. While we acknowledge as a minor limitation that tripod-captured images cannot 100% replicate the dynamic micro-vibrations and exact aerodynamic conditions of an airborne UAV, they serve as a highly accurate proxy. Combining these simulated images with the actual airborne footage from the UAVVaste dataset successfully bridged any potential domain gap.

3.4. MODEL TRAINING

During the model training stage, the YOLOv12 model was selected for waste detection. In this study, the YOLOv12 architecture was specifically selected over its predecessors (such as YOLOv8 and YOLOv11) based on recent comparative studies in UAV-based object detection. According to Tian et al. (2025), YOLOv12 integrates an attention-centric framework with Residual Efficient Layer Aggregation Networks (R-ELAN), addressing the inefficiencies of traditional CNN-based architectures. This allows YOLOv12 to outperform YOLOv8 and YOLOv11 in terms of mAP and inference latency, particularly for small object detection tasks. Recent implementations have highlighted the model's capabilities; for example, Buleu et al. (2025) reported that YOLOv12 outperforms its predecessors in precision-demanding tasks. Since the primary contribution of this study is the novel end-to-end offboard processing system and the Video–Log Synchronization Module, we relied on these established state-of-the-art baseline validations rather than conducting redundant architectural ablation studies.

Regarding the training configuration, the model was fine-tuned using weights pretrained on the COCO dataset. The hyperparameters were rigorously configured to ensure stable convergence: the model was trained for a maximum of 100 epochs with an early stopping patience of 20 epochs to prevent overfitting (Goodfellow et al., 2016). The optimization was performed using Stochastic Gradient Descent (SGD) with a momentum of 0.937 and a weight decay of 0.0005. The initial learning rate was set to 0.001 (lr0=0.001) and linearly decayed to a final learning rate factor of 0.01 (lrf=0.01). A batch size of 16 was utilized. Model training was conducted in the Google Colab environment using Tesla T4 GPU support. To measure the model's performance, the mAP@0.5 and mAP@0.5–0.95 metrics—standardized in the literature (Lin et al. 2014)—were used.

3.5. VIDEO–LOG SYNCHRONIZATION MODULE

One of the other core components of the system is the Data Synchronization Module, which matches the frames in the video recorded by the UAV with the UAV telemetry data on a location–frame basis. At this stage, SRT log files recorded simultaneously with the videos produced by the DJI Mini 3 are processed, enabling the waste detections to be presented in a location-based manner.

The synchronization process is carried out through the following steps: each frame in the video file recorded by the UAV is extracted according to its timestamp, and the GPS (latitude, longitude, altitude) information of the corresponding frame is parsed from the SRT file. To ensure precise mapping, navigation data from the UAV autopilot, including GPS and IMU measurements, are time-synchronized with the camera's image stream. Object detection and location matching are carried out by processing the video stream frame by frame using the YOLOv12-trained model. For every bounding box produced by the model, the corresponding video frame is extracted, the GPS data with the same timestamp is retrieved, and the object class (waste), confidence score, and location information are stored together (Leira et al. 2015; Leira et al. 2021).

3.6. CLOUD STORAGE AND DATA MANAGEMENT

Following the detection and localization of coastal waste via the Video-Log Synchronization Module, the resulting georeferenced data must be efficiently stored and managed for end-users. For each detected waste item, the cropped image, class information, date-time, and GPS coordinates are stored as separate fields. In this system, the Cloudinary infrastructure was preferred due to its secure storage of images, automatic URL generation, seamless metadata access, and the availability of a free starter plan to support the project's low-cost objective.

This architectural choice perfectly aligns with recent theoretical advancements in environmental data analytics. As comprehensively analyzed by Verma et al. (2024), the integration of Cloud Computing with distributed sensing devices (such as UAVs and IoT nodes) fundamentally transforms environmental monitoring by overcoming traditional limitations in data granularity, geographic coverage, and real-time analysis. By leveraging scalable cloud infrastructures, our system ensures that the massive datasets generated by 4K UAV imagery are efficiently managed without relying on heavy local storage. Ultimately, this cloud-integrated approach provides a robust foundation for the proposed ASP.NET Core MVC-based web dashboard, enabling rapid, data-driven decision-making and facilitating collaborative, real-time interventions by municipal cleanup crews (Verma et al. 2024).

3.7. WEB-BASED DASHBOARD AND VISUALIZATION

In the final stage of the study, the frames in which waste detections were matched with their corresponding locations were visualized using a free web-based dashboard hosted on the Render platform. The main components of the dashboard are: (1) Map-based visualization – Waste detected in the images is displayed on a map using markers that indicate location and heatmaps that represent density, based on geographic coordinate data; and (2) Visual gallery and detailed inspection – Images recorded separately for each waste detection can be viewed in a gallery.

Figure 5 and Figure 6 present two different visualizations generated by the web-based dashboard, showing the detected waste as pinned markers on a satellite map and as a heatmap for analysis purposes.

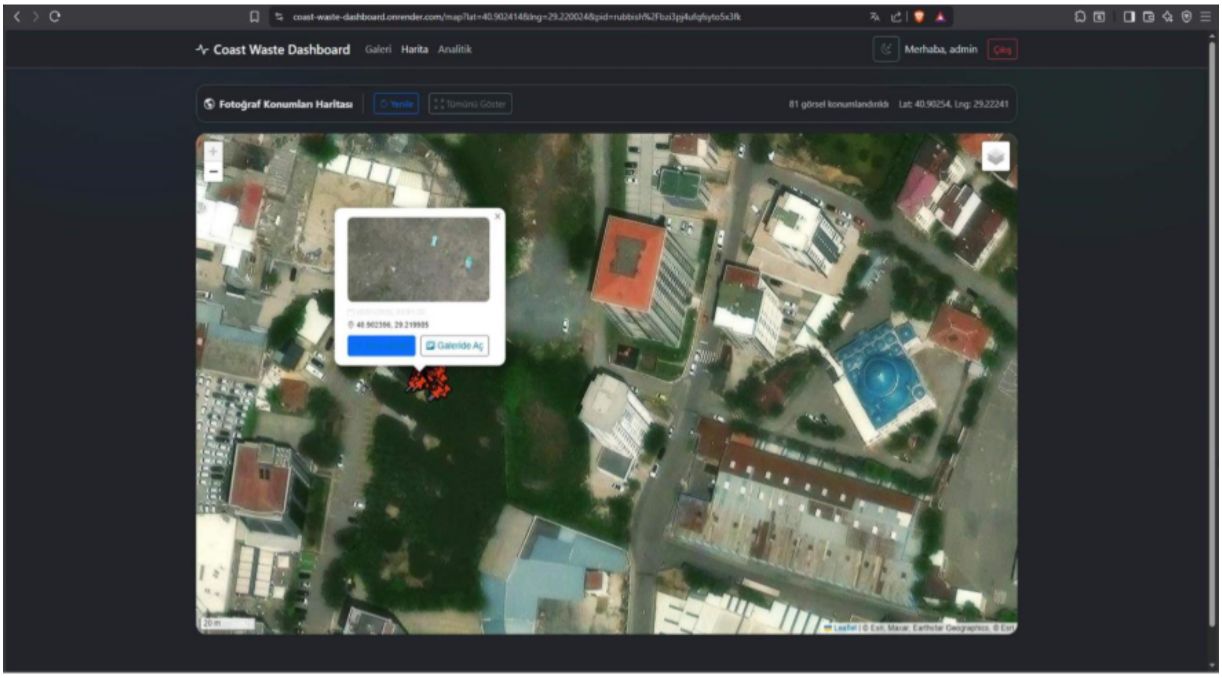

Figure 5. Pinned representation of waste on the satellite map.

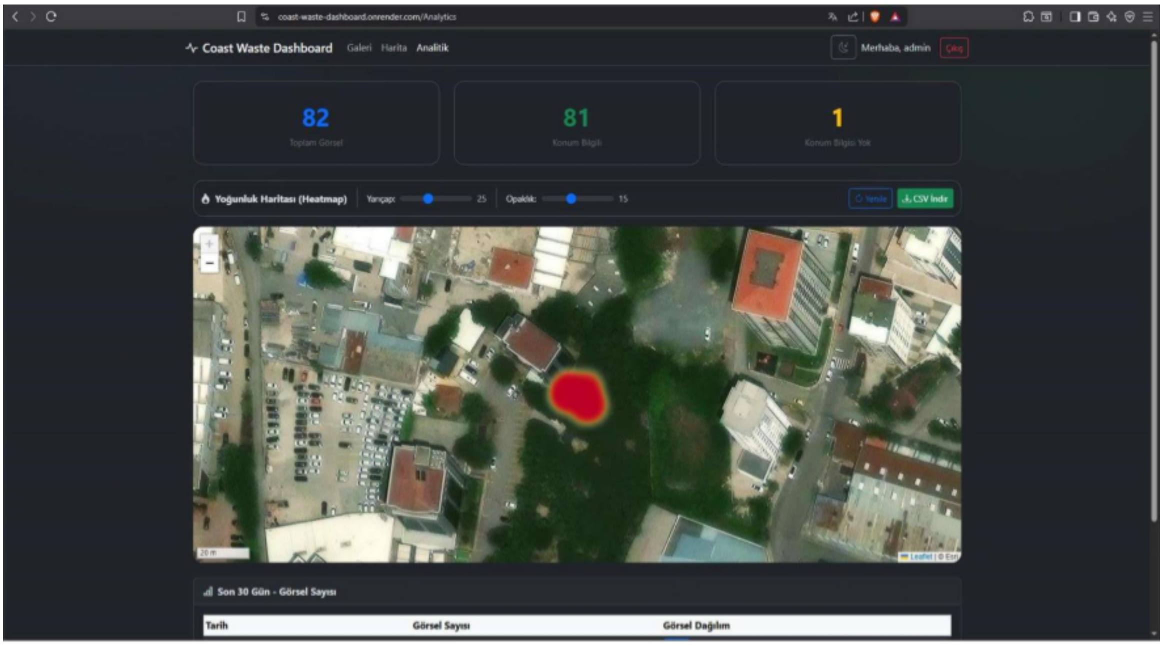

Figure 6. Heatmap visualization showing waste density of the area on the analysis page.

In conclusion, the developed system produces outputs that are easily interpretable by different user profiles. In this context, a coastal waste monitoring infrastructure suitable for field conditions has been established through the use of low-cost UAVs and a ground-station-based AI processing approach.

3.8. EXPERIMENTAL DESIGN

In terms of operational parameters, the primary objective of this study was to establish a reliable performance baseline for the proposed offboard detection architecture rather than conducting a multi-variable sensitivity analysis. Therefore, critical flight parameters were deliberately standardized based on established UAV monitoring protocols in the literature. The camera angle was fixed at a 90-degree nadir position to minimize perspective distortions and occlusions caused by coastal vegetation. The flight speed was kept constant at 1.5–2.5 m/s to prevent motion blur, and flights were executed under optimal, sunny conditions with low wind speeds (< 2.5 km/h). While environmental variables such as extreme sunlight glare, wave movement, and high winds undeniably affect the detection accuracy, analyzing the isolated impact of each of these variables requires extensive multi-season testing and is considered a scope for future work. The constraints introduced by these environmental factors are further discussed in Section 5.3.

Furthermore, it is important to note the geographical limitations of the system's validation. While the training phase utilized a hybrid dataset incorporating the publicly available UAVVaste dataset to introduce geographical and topological diversity, the real-world field deployment and quantitative validation of the proposed system were strictly limited to a specific coastline in Istanbul. Consequently, although the model demonstrated robust performance in this specific operational environment, its detection efficacy may vary across different coastal topographies, sand colors, or under different regional environmental conditions. Extensive cross-regional field validation remains a necessary subject for future large-scale deployments.

4. RESULTS AND DISCUSSION

In this section, experimental results of coastal waste detection and mapping using developed UAVs are presented, and the success of the system is discussed by comparing the results with similar studies in the literature. While evaluating the success of the system; model accuracy, success in detecting small wastes, location matching accuracy, and overall system efficiency were addressed.

4.1. MODEL DETECTION PERFORMANCE

The system, built on the YOLOv12 model, demonstrated successful performance in waste detection demo studies conducted on coasts with complex backgrounds. To provide a comprehensive evaluation of the YOLOv12 model, multiple performance metrics beyond mAP@0.5 were analyzed, including the strict mAP@0.5:0.95, Recall, and F1-score. According to the final validation results, the model achieved a Precision of 0.850, a Recall of 0.694, and an F1-score of 0.764. The mAP@0.5:0.95 metric, which averages performance across multiple Intersection over Union (IoU) thresholds and heavily penalizes poor bounding box localization, was recorded at 0.491. The mAP@0.5 value obtained was 0.783.

It should be noted that the detection task in this study was intentionally formulated as a single-class problem (unified as "rubbish" or "coastal waste") rather than material-specific sorting (e.g., plastic, metal, glass). This approach was chosen to maximize the overall detection capability and localization efficiency for guiding human cleanup crews to waste accumulation zones, regardless of the material type. Therefore, the reported overall metrics directly represent the per-class performance of the targeted waste category. In the literature, detection performances for UAV-based waste monitoring vary significantly; for instance, Pfeiffer et al. (2023) reported mAP values around 0.31 for general litter classes, while recent optimized models like QSB-YOLO achieved up to 0.84 (Zhu et al., 2023). Considering these findings, the mAP@0.5 of 0.783 obtained in this study positions the proposed system at the upper tier of existing solutions.

As shown in Figure 7, the obtained performance is at the upper band of existing studies and therefore can be considered successful.

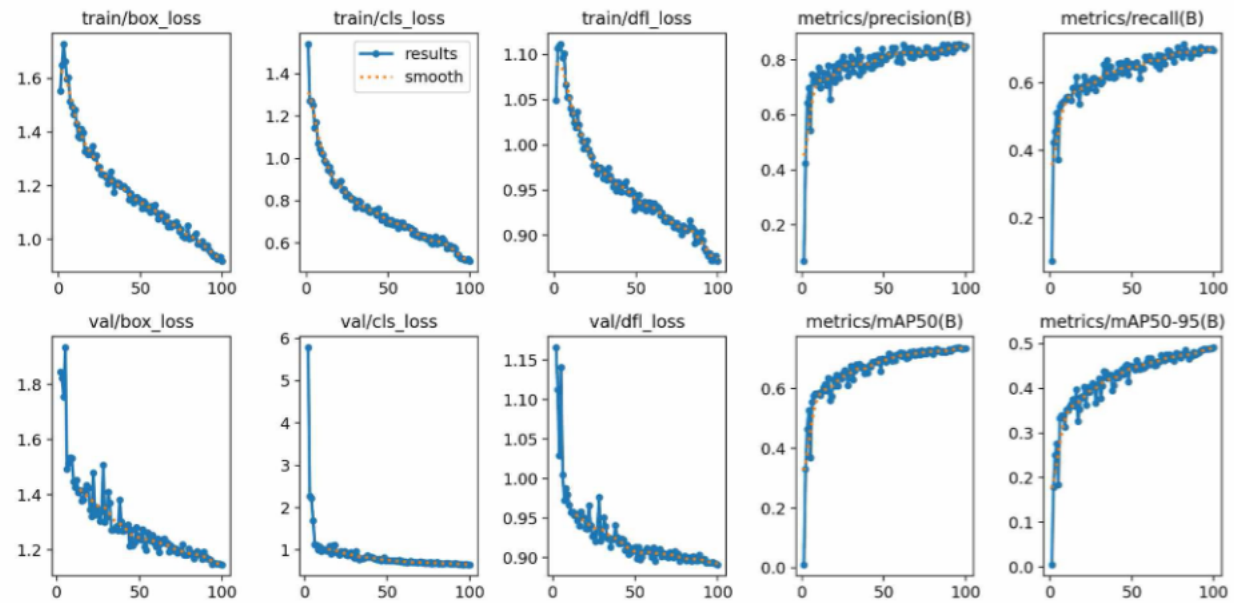

Figure 7. Change in training and validation performance of the model according to epochs.

Regarding failure case analysis, missed detections were occasionally observed for waste objects partially buried under sand, under heavy shadows, or obscured by vegetation. False positives were primarily triggered by natural coastal elements such as white sea shells mimicking styrofoam or sun glare on wet rocks mimicking metallic cans. These results are similar to previous UAV coastal waste detection studies, which noted the sensitivity of RGB-based approaches to image limitations and lighting variability caused by the angle of sunlight (Merlino et al., 2020). As seen in Figure 8, the output of the YOLOv12-based study on a video frame captured by the UAV is shown.

Figure 8. Waste detections of the model on the image.

4.2. POSITIONING AND SYNCHRONIZATION ACCURACY

Thanks to the video-log synchronization module developed with Python, the frames in the video and the flight data in the SRT file were processed, and frame-location matching was automatically performed at the offboard station. In field experiments, the positional error was measured at an average of 1.5–2.5 meters, and this value is directly dependent on the performance of the GPS module on the UAV used.

To validate this accuracy without relying on expensive RTK-GPS hardware, the UAV's flight logs were extensively analyzed. The drone's Global Navigation Satellite System (GPS + GLONASS + Galileo) successfully connected to an average of 24.26 satellites during the flight, maintaining an excellent GPS signal quality (gpsLevel) of 4.91 out of 5. According to the manufacturer's specifications, this multi-constellation GNSS setup provides a horizontal hovering accuracy of ±1.5 meters (DJI 2025). The alignment of our measured 1.5–2.5 meter error range with the hardware's theoretical limits and the high satellite connection rate proves that the 5 Hz GPS sampling rate and software-based synchronization provide highly reliable georeferencing for coastal waste mapping.

Furthermore, when compared to other localization methods, while Real-Time Kinematic (RTK) and Post-Processed Kinematic (PPK) techniques provide centimeter-level positioning accuracy (typically 1–3 cm), they require continuous, low-latency communication links to reference networks or the deployment of local base stations, which significantly increases operational complexity and hardware costs (Eker et al. 2021). In contrast, standard Single Point Positioning (SPP) via multi-constellation GNSS, as adopted in this study, typically provides horizontal accuracies in the range of 2–5 meters in open-sky baseline conditions (Jiang et al. 2026). Our obtained 1.5–2.5 meter accuracy is well within this expected range and proves highly efficient. Although RTK/PPK methods and ground-truth validation are indispensable for precise engineering surveys, the standard multi-constellation GNSS approach offers a scalable, cost-effective balance that is perfectly adequate for macro-level coastal waste density mapping and for successfully guiding human cleanup crews.

Regarding real-world applications, this 1.5–2.5 meter localization error is highly acceptable and has no adverse effect on practical cleanup planning. Since the primary objective of the proposed system is to generate macro-level waste density heatmaps and guide human cleanup crews to accumulation hotspots, millimeter-level precision—which is mandatory for autonomous robotic retrieval—is not required. When a municipal cleanup crew is directed to a specific coordinate provided by the dashboard, a spatial deviation of up to 2.5 meters falls easily within the human visual field, allowing workers to naturally spot and retrieve the macro-litter without any operational delay.

4.3. OPERATIONAL EFFICIENCY

When compared with traditional manual coastal inspection methods frequently used by municipalities and non-governmental organizations, significant gains in operational efficiency were achieved. To scientifically support the claim that a one-kilometer coastline can be scanned in under ten minutes, detailed flight log analyses were conducted.

During a continuous field test, the UAV covered a total distance of 1,134 meters in exactly 13 minutes and 38 seconds, executing a combination of variable-speed sweeps and stationary hovering. The average scanning speed was calculated at 6.4 km/h. At this average operational speed, scanning a linear 1-kilometer coastline takes precisely 9 minutes and 44 seconds, proving that the sub-10-minute target is not only realistic but completely field-tested.

Moreover, regarding energy efficiency, this 13.5-minute flight consumed 44% of the battery capacity. According to the manufacturer's specifications, the maximum flight time of 38 minutes is measured in a controlled, windless environment at a continuous speed of 21.6 km/h. In contrast, real-world coastal scanning requires frequent stationary hovering and acceleration against coastal winds, which naturally increases the power draw. Even under these demanding real-world field conditions, consuming only 44% of the battery for a complete 1.1 km detailed scan demonstrates that the proposed offboard system provides robust sustainability for large-scale monitoring. It successfully avoids the rapid battery drain and drastically reduced flight times typically caused by the heavy power requirements of onboard edge-computing processors.

4.3. OPERATIONAL EFFICIENCY

When compared with traditional manual coastal inspection methods frequently used by municipalities and non-governmental organizations, significant gains in operational efficiency were achieved. To scientifically support the claim that a one-kilometer coastline can be scanned in under ten minutes, detailed flight log analyses were conducted.

During a continuous field test, the UAV covered a total distance of 1,134 meters in exactly 13 minutes and 38 seconds, executing a combination of variable-speed sweeps and stationary hovering. The average scanning speed was calculated at 6.4 km/h. At this average operational speed, scanning a linear 1-kilometer coastline takes precisely 9 minutes and 44 seconds, proving that the sub-10-minute target is not only realistic but completely field-tested.

Moreover, regarding energy efficiency, this 13.5-minute flight consumed 44% of the battery capacity. According to the manufacturer's specifications, the maximum flight time of 38 minutes is measured in a controlled, windless environment at a continuous speed of 21.6 km/h. In contrast, real-world coastal scanning requires frequent stationary hovering and acceleration against coastal winds, which naturally increases the power draw. Even under these demanding real-world field conditions, consuming only 44% of the battery for a complete 1.1 km detailed scan demonstrates that the proposed offboard system provides robust sustainability for large-scale monitoring. It successfully avoids the rapid battery drain and drastically reduced flight times typically caused by the heavy power requirements of onboard edge-computing processors.

4.4. OFFBOARD ARCHITECTURE EVALUATION

Overall, the findings show that the proposed system offers effective solutions to constraints such as high operational costs, limited area scanning, and optimum flight time of manual and onboard UAV waste detection systems frequently emphasized in the literature, thanks to its offboard architecture. By adopting the offboard processing architecture, the short flight time problem caused by the high energy consumption and carrying capacity constraints of onboard systems has been eliminated. These constraints are among the issues that cause onboard UAV-based object detection applications to become studies that are successful on paper but limited in operational use, far from efficiency in field trials (Duangsuwan & Prapruetdee 2024). These results are consistent with other recent research emphasizing the advantages of UAV-based digital solutions compared to labor-intensive manual methods for monitoring marine ecosystems (Marin et al. 2024; Duangsuwan & Prapruetdee 2024).

5. CONCLUSION

5.1. SUMMARY OF FINDINGS

This study has put forward a successful system in field trials, including steps for detection, location matching, and visual display on a dashboard, by using ordinary UAVs that only have camera and GPS features—which are easily accessible—together with YOLO-based object detection approaches for the detection and spatial mapping of human-sourced solid plastic waste in coastal and marine ecosystems.

As emphasized in the literature, a large portion of the waste causing marine pollution and damaging the ecosystem consists of land-based plastic waste, and this situation affects not only the marine ecosystem but also humans indirectly through the food chain by exposure to microplastics (Terzi et al. 2025). In this context, the developed system aimed to be an auxiliary resource providing regular, periodic, modular, and scalable monitoring of coastal waste.

The YOLOv12 architecture exhibited high performance in detecting small, irregularly shaped, and low-contrast coastal waste in complex backgrounds thanks to its advanced attention-focused mechanisms (Tian et al. 2025). The results observed in this study support that attention and transformer-based architectures are one of the correct approaches to overcome this difficulty, in light of the results obtained.

5.2. OFFBOARD ARCHITECTURE AND OPERATIONAL ADVANTAGES

The construction of the system on an offboard architecture proved its field suitability by providing significant gains in operational efficiency. Focusing only on the data collection task without creating an additional processor load on the UAV thanks to the offboard architecture allowed for the maximization of flight time and made it possible to scan wide coastal areas in a single flight. The hybrid data approach adopted in this study—combining custom UAV footage with public datasets—produced stable results and improved generalization capabilities. The findings obtained support the fact that data diversity should be considered a priority in model training and that it directly determines success.

5.3. LIMITATIONS OF THE STUDY

While the proposed system demonstrates high accuracy and operational efficiency, it is important to explicitly acknowledge certain limitations inherent to UAV-based monitoring systems.

First, the detection performance is highly dependent on favorable lighting and weather conditions. The field flights and dataset acquisitions in this study were conducted under optimal conditions in June, characterized by bright sunlight and low wind speeds (approximately 2.5 km/h). Under suboptimal environmental conditions—such as overcast skies, varying sunlight angles causing glare, or high winds causing motion blur and drone instability—the effectiveness of RGB-based visual detection is expected to decrease (Merlino et al. 2020).

Secondly, detecting objects that are under heavy shadows, obscured by vegetation, or partially submerged in water remains a fundamental challenge for RGB optical sensors (Marin et al. 2024). To mitigate this, the custom 300-image dataset created for this study deliberately included waste items that were partially buried in sand or covered by coastal vegetation. Training with this heterogeneous data, combined with the advanced attention-centric architecture of YOLOv12 (Tian et al. 2025), significantly improved the model’s ability to extract features from partially occluded objects compared to older YOLO generations. However, in cases of severe occlusion or deep submersion, false negatives are still unavoidable.

Furthermore, the spatial mapping relies on consumer-grade GNSS sensors equipped on commercial UAVs. Although the observed positional error of 1.5–2.5 meters is highly acceptable for guiding human cleanup crews to waste accumulation zones, it lacks the centimeter-level precision of costly RTK-GPS systems (Leira et al. 2015). Thus, the current system is perfectly suited for macro-level density mapping, but it may not be precise enough for tasks requiring millimeter accuracy, such as autonomous robotic arm retrieval.

Finally, the deliberate choice of an offboard processing architecture inherently results in a lack of real-time detection capabilities. Since video streams and telemetry logs are processed at the ground station post-flight, the system cannot provide instantaneous feedback or dynamic obstacle-avoidance alerts based on detected waste. However, this limitation is a necessary trade-off; offloading the computational burden from the UAV avoids the rapid battery drain associated with edge-computing, extending the flight time to nearly 38 minutes and allowing the use of computationally heavy, state-of-the-art models like YOLOv12 without requiring expensive payload modifications.

5.4. FUTURE DIRECTIONS AND PUBLIC IMPACT

Finally, by presenting the waste-location matching visually through cloud-based storage and a web-based dashboard, it is aimed to make the results accessible and understandable even by non-technical users, and to lead to the formation of public awareness. In the literature, it is stated that in environmental waste detection and observation projects, the lowering of model outputs to a level that normal users can understand through digital platforms, as in this study, accelerates the process and operational efficiency (Merlino et al. 2020).

Evaluated generally, this study has put forward a coastal waste detection and monitoring system that is scalable, periodically applicable, and adaptable to real field conditions by bringing together low-cost commercial UAVs with the most up-to-date deep learning architectures and cloud-based storage technology.

AUTHOR CONTRIBUTIONS

Eyup Koyun: Conceptualisation, methodology, data curation, writing of the original draft, and review and editing of the final manuscript, read and approved the final version of the manuscript. **Nur Hazal Cufalci**: Conceptualisation, methodology, data curation, writing of the original draft, and review and editing of the final manuscript, read and approved the final version of the manuscript. **Berke Akdogan**: Conceptualisation, methodology, data curation, writing of the original draft, and review and editing of the final manuscript, read and approved the final version of the manuscript. **Aytac Ugur Yerden**: Conceptualisation, methodology, data curation, writing of the original draft, and review and editing of the final manuscript, read and approved the final version of the manuscript.

COMPETING INTERESTS

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

DATA ACCESSIBILITY

The dataset used for training the model in this study is shared under the CC BY 4.0 license and is publicly accessible via the Roboflow Universe platform. Repository link: universe.roboflow.com/berkea8/iha-goruntulerinden-kiyi-atigi-tespiti.

ETHICAL APPROVAL

This study did not require ethical approval as it did not involve human participants, animal subjects, or sensitive personal data.

FUNDING

This study received no external financial funding or grants.

REFERENCES

Bochkovskiy, A., Wang, C. and Liao, H.M., 2020. YOLOv4: Optimal speed and accuracy of object detection. arXiv. DOI: https://doi.org/10.48550/arXiv.2004.10934.

Buleu, B., Robu, R. and Filip, I., 2025. A deep learning-based system for automatic license plate recognition using YOLOv12 and PaddleOCR. Applied Sciences, 15(14), 7833. DOI: https://doi.org/10.3390/app15147833.

DJI, 2025. DJI Mini 3 User Manual v1.2. SZ DJI Technology Co., Ltd., Shenzhen, China.

Duangsuwan, S. and Prapruetdee, P., 2024. Drone-enabled AI edge computing and 5G communication network for real-time coastal litter detection. Drones, 8(12), 750. DOI: https://doi.org/10.3390/drones8120750.

Eker, R., Alkan, E. and Aydın, A., 2021. A comparative analysis of UAV-RTK and UAV-PPK methods in mapping different surface types. European Journal of Forest Engineering, 7(1), 12–25. DOI: https://doi.org/10.33904/ejfe.938067.

Goodfellow, I., Bengio, Y. and Courville, A., 2016. Deep learning. Cambridge, MA: MIT Press.

Jiang, C., Zhou, X., Chen, H. and Liu, T., 2026. UAV positioning using GNSS: A review of the current status. Drones, 10(2), 91. DOI: https://doi.org/10.3390/drones10020091.
Kraft, M., Piechocki, M., Ptak, B. and Walas, K., 2021. Autonomous, onboard vision-based trash and litter detection in low altitude aerial images collected by an unmanned aerial vehicle. Remote Sensing, 13(5), 965. DOI: https://doi.org/10.3390/rs13050965.
Leira, F.S., Trnka, K., Fossen, T.I. and Johansen, T.A., 2015. A lightweight thermal camera payload with georeferencing capabilities for small fixed-wing UAVs. In: Proceedings of the 2015 International Conference on Unmanned Aircraft Systems (ICUAS), Denver, CO, USA, 9–12 June 2015, pp. 485–494. DOI: https://doi.org/10.1109/ICUAS.2015.7152327.
Leira, F.S., Helgesen, H.H., Johansen, T.A. and Fossen, T.I., 2021. Object detection, recognition, and tracking from UAVs using a thermal camera. Journal of Field Robotics, 38, 242–267. DOI: https://doi.org/10.1002/rob.21985.
Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P. and Zitnick, C.L., 2014. Microsoft COCO: Common objects in context. In: Computer Vision – ECCV 2014, Zurich, Switzerland, 6–12 September 2 014, pp. 740–755. DOI: https://doi.org/10.1007/978-3-319-10602-1_48.
Marin, D., Ciucă, A.-M., Filimon, A. and Stoica, E., 2024. Marine litter and plastic detection on Black Sea beaches by using unmanned aerial vehicle (UAV). Ovidius University Annals of Chemistry, 35(1), 43–50. DOI: https://doi.org/10.2478/auoc-2024-0006.
Merlino, S., Paterni, M., Berton, A. and Massetti, L., 2020. Unmanned aerial vehicles for debris survey in coastal areas: Long-term monitoring programme to study spatial and temporal accumulation of the dynamics of beached marine litter. Remote Sensing, 12(8), 1260. DOI: https://doi.org/10.3390/rs12081260.
Panwar, H., Gupta, P.K., Siddiqui, M.K., Morales-Menendez, R., Bhardwaj, P., Sharma, S. and Sarker, I.H., 2020. AquaVision: Automating the detection of waste in water bodies using deep transfer learning. Case Studies in Chemical and Environmental Engineering, 2, 100026. DOI: https://doi.org/10.1016/j.cscee.2020.100026.
Pfeiffer, R., Valentino, G., D'Amico, S., Piroddi, L., Galone, L., Calleja, S., Farrugia, R.A. and Colica, E., 2023. Use of UAVs and deep learning for beach litter monitoring. Electronics, 12(1), 198. DOI: https://doi.org/10.3390/electronics12010198.
Presidency of the Republic of Türkiye Digital Transformation Office, 2021. National artificial intelligence strategy (2021–2025). Ankara: Presidency of the Republic of Türkiye. Available at https://wp.oecd.ai/app/uploads/2021/12/Turkey_National_Artificial_Intelligence_Strategy_2021-2025.pdf [Accessed 11 May 2026].
Presidency of the Republic of Türkiye Strategy and Budget Directorate, 2024. 12th development plan (2024–2028). Ankara: Presidency of the Republic of Türkiye. Available at https://www.sbb.gov.tr/wp-content/uploads/2025/03/Twelfth-Development-Plan_2024-2028.pdf [Accessed 11 May 2026].
Republic of Türkiye Ministry of Industry and Technology, 2023. 2030 industry and technology strategy. Ankara: Republic of Türkiye Ministry of Industry and Technology. Available at https://www.sanayi.gov.tr/assets/pdf/plan-program/2030IndustryAndTechnologyStrategyENG.pdf [Accessed 11 May 2026].
Terzi, Y., Gündoğdu, S., Yandi, İ., Altinpinar, İ., Öztürk, R.Ç. and Gedik, K., 2025. Marine litter on the Turkish Black Sea shoreline: Abundance, composition, and sources. Waste Management, 205, 115027. DOI: https://doi.org/10.1016/j.wasman.2025.115027.
Tian, Y., Ye, Q. and Doermann, D.S., 2025. YOLOv12: Attention-centric real-time object detectors. arXiv. DOI: https://doi.org/10.48550/arXiv.2502.12524.
Verma, R., Taneja, H., Singh, K.D. and Singh, P.D., 2024. Enhancing data analytics in environmental sensing through cloud IoT integration. Journal of Climate Change, 10(2), 6. DOI: https://doi.org/10.3233/JCC240013.
Yerden, A.U., Şenol, S., Kara, M. and Dilibal, S., 2025. xT-STRIDE threat model for unmanned air vehicle security. International Journal of Information Security, 24, 169. DOI: https://doi.org/10.1007/s10207-025-01082-4.
Zhu, H., Zhu, D., Qin, X. and Guo, F., 2023. Efficient and accurate beach litter detection method based on QSB-YOLO. International Journal of Advanced Computer Science and Applications, 14(7). DOI: https://doi.org/10.14569/IJACSA.2023.0140735.