Intelligent Environments 2026: Combined Workshop Proceedings
D. Micucci et al. (Eds.)

doi:10.3233/AISE260024

An Intelligent Sensor System for Adaptive Data Acquisition in Aquatic Environments

Andrej LEJMAN [a,1] Eike RODENBÄCK [a] Christoph MANSS [a]
Felix BECKER [a] and Frederic STAHL [a]

[a] *German Research Center for Artificial Intelligence (DFKI), Oldenburg, Germany*

ORCiD ID: Andrej Lejman https://orcid.org/0009-0002-2515-0059, Eike Rodenbäck https://orcid.org/0009-0000-9631-6226, Christoph Manss https://orcid.org/0000-0003-4851-2622, Felix Becker https://orcid.org/0009-0005-4305-3377, Frederic Stahl https://orcid.org/0000-0002-4860-0203

Abstract. This article introduces an intelligent sensor system installed on a mowing boat operating on the Masch Lake in Hanover, Germany, designed to control unwanted aquatic vegetation. The system adaptively regulates data acquisition based on real-time aquatic plant density estimation and integrates image streams from visible and near-infrared channels, along with 3D point cloud data. A lightweight neural network processes the data on an edge computing platform. By dynamically adjusting the data acquisition rate across four defined mowing stages, the system maximises the relevance of the collected information for subsequent data processing and analysis while minimising the acquisition and storage of irrelevant or redundant data. Field evaluations confirm that the system effectively captures relevant data during ecologically significant stages and suppresses acquisition during less relevant stages, resulting in a dataset of higher informational value. These findings highlight the system's potential to improve ecological monitoring and enhance operational efficiency in aquatic vegetation management.

Keywords. Aquatic plants, edge computing, artificial intelligence.

1. Introduction

Environmental monitoring and managing of aquatic ecosystems increasingly relies on adaptive data acquisition to capture dynamic and context-sensitive phenomena [1–3]. Traditional static sampling strategies, which typically collect data at fixed time intervals, often fail to reflect rapid ecological variations. These include seasonal fluctuations, local disturbances such as boat traffic or weather events,

[1]Corresponding Author: Andrej Lejman, andrej.lejman@dfki.de

and episodic occurrences like algal blooms or sediment resuspension [4]. Such phenomena require more flexible and responsive approaches to ensure relevant and timely data acquisition. Adaptive sampling algorithms, such as the Data-Driven Adaptive Sampling Algorithm (DDASA) [5] and Dynamically Sized Rate Adaptive – Parametric Machine Learning-based (DSRA-PMLO) [6], dynamically adjust data acquisition rates based on signal variability and predictive models. Although these methods primarily target low-dimensional sensor data, they emphasise the significance of context-aware acquisition strategies, which are designed to optimise data relevance.

Edge computing has gained relevance for mobile platforms with limited computational resources. Lou et al. [7] propose a framework that combines adaptive sampling with onboard preprocessing to reduce latency and filter data locally. Mobile robotic systems, such as Unmanned Surface Vehicles (USVs) and Autonomous Underwater Vehicles (AUVs), combine autonomous navigation with adaptive sampling and Artificial Intelligence (AI)-based decision-making [8–10].

AI techniques, such as pattern recognition and adaptive control, as demonstrated in recent robotic systems [11–14], inform not only autonomous navigation but also the visual detection and sampling logic implemented in this work. Despite recent advances, substantial research gaps remain. Many existing adaptive sampling methods are designed for low-dimensional environmental measurements and have not been extended to high-volume visual modalities, e.g., point cloud data.

This work addresses these gaps by presenting an intelligent sensor system for high-volume visual acquisition. We mounted the proposed system on a mowing boat[2] and tested it on the Masch Lake in Hanover, Germany, a shallow urban lake with dense vegetation. Using AI-based image analysis, the system monitors the boat's conveyor belt to dynamically adjust the data acquisition.

The paper is structured as follows: Section 2 details the methodology and AI logic, while Section 3 describes the hardware and datasets. Section 4 presents the field evaluation and performance analysis. Section 5 concludes the paper.

2. System Architecture

The system employs a modular, Docker-based [15] Robot Operating System (ROS) framework to ensure cross-platform portability. As illustrated in Figure 1, the architecture comprises three functional layers: sensing, decision-making, and control.

(I) The *sensing layer* manages hardware-level data acquisition through an RGB Depth (RGB-D) sensor and a Global Navigation Satellite System (GNSS) unit. It captures pixel-aligned Red Green Blue (RGB) and Near Infrared (NIR) imagery alongside 3D point clouds, covering the conveyor belt area. Simultaneously, the GNSS sensor provides precise positioning data to enable accurate georeferencing of all recorded information.

[2] https://www.berky.de/en/aquatic-weed-harvester/, accessed on 17-12-2024

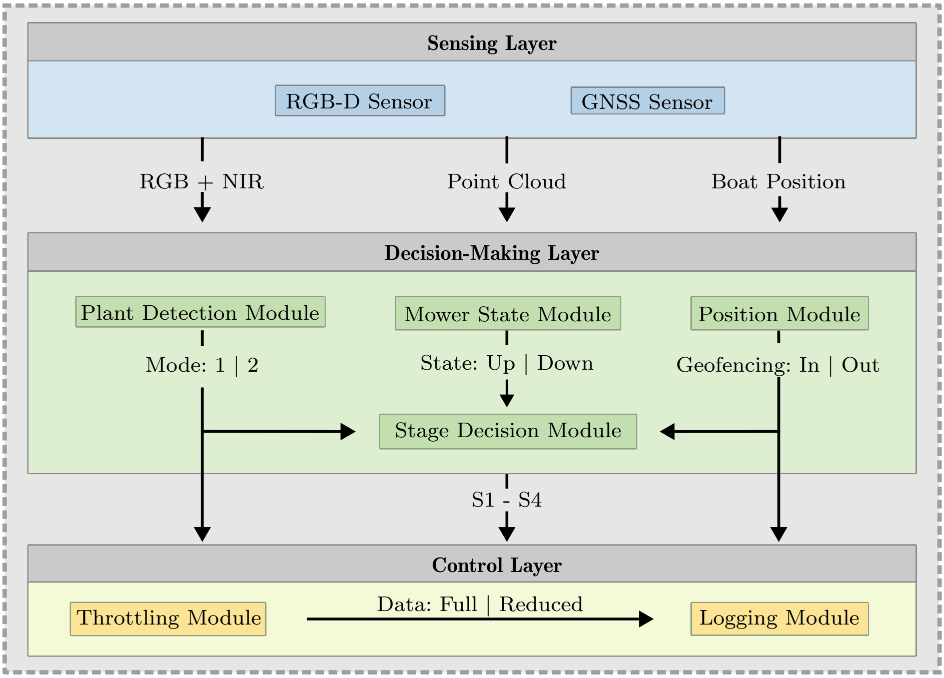

Figure 1. The software architecture comprises three layers: Sensing, Decision-Making, and Control. RGB-D and GNSS sensors provide RGB+NIR, point cloud, and positioning data. Within the Decision-Making Layer, modules evaluate plant detection, mower state, and geofencing to determine the operational stage (S1—S4). Finally, the Control Layer manages data throttling (Full | Reduced) and logging for transmission to a remote server.

(II) The *decision-making layer* interprets sensor data to assess the operational context and classify the system into four stages (S1–S4). A central plant detection module employs a lightweight Neural Network (NN) for pixel-wise multispectral segmentation, building on established approaches [16] and their adaptation for aquatic environments [17]. By processing aligned RGB and NIR channels, the NN exploits spectral cues for plant discrimination. The architecture features a single hidden layer with 64 neurons and ReLU activation. To ensure efficient deployment on edge devices, we optimised the architectural parameters (number of hidden layers and neurons) and training hyperparameters (learning rate, batch size, and dropout rate) using the Optuna framework [18]. Training utilised the Adam optimiser and a binary cross-entropy loss function on a dataset of 950 manually annotated images, partitioned into training (80 %), validation (10 %), and testing (10 %) subsets. On the held-out test set, the model achieves an accuracy of 84.9 % and an F1 score of 0.836 while maintaining a minimal footprint of 4,545 parameters.

Based on the network's segmentation output, the system calculates the percentage of aquatic plants within the observed area. Comparing this value against a predefined threshold ($0 \leq t \leq 4$) determines the acquisition mode: Mode M1 for high data acquisition and Mode M2 for reduced data acquisition, as depicted in Figure 2.

Figure 2. Visualisation of the four mowing stages (S1—S4) and their corresponding data acquisition modes. Each stage is defined by mower position and plant presence on the conveyor belt, illustrated through schematics and RGB images. Adaptive acquisition is split into Mode M1 (full data) for stages S1—S2 and Mode M2 (reduced data) for S3—S4.

Simultaneously, the system evaluates the 3D point cloud to distinguish between active mowing (S1, S3) and inactive transit (S4). It detects whether the mower is lowered or raised by calculating the average distance between the RGB-D sensor and the conveyor belt. Additionally, a point-in-polygon algorithm [19] processes GNSS data to identify whether the boat is within a geofenced unloading area (S2). Finally, the system fuses these spectral, geometric, and spatial inputs to classify the current operational stage (S1–S4).

(III) The *control layer* executes acquisition logic by regulating data flow based on the identified stage. A throttling module dynamically adjusts ROS publication rates; in mode M2, it reduces image streams to 1 Hz and disables the volumetric point cloud stream entirely to conserve bandwidth and storage. Simultaneously, a logging module manages data persistence, implementing a timeout-based pause when the boat enters unloading zones to prevent the accumulation of redundant idle data.

3. Hardware Setup and Dataset

The sensor system, installed on a mowing boat, features an NVIDIA Jetson Xavier NX as its central unit for local AI processing and decision-making. Mounted on the roof and aligned with the conveyor belt, a FRAMOS D455 RGB-D sensor captures multi-modal data, including RGB, NIR, and depth point clouds, while an integrated GNSS sensor and LTE link provide precise georeferencing and remote server connectivity. The mowing boat's front-mounted cutters sever the plants, while its conveyor belt carries the material directly into a central holding area.

For the recording frequencies, we set all sensor streams in Mode M1 to approximately 2 Hz. This frequency ensures a sufficient spatial overlap between consecutive frames at typical operating speeds. In Mode M2, we set the recording frequency of the RGB and NIR image streams to 1 Hz, while point cloud acquisition is disabled entirely.

The evaluation relies on two complementary datasets, referred to as adaptive and static dataset. The adaptive dataset (recorded May/Aug 2025), using the adaptive system installed on the mowing boat. The static dataset (July/Aug 2024), using the previous static system. This dataset also served as the basis for training the neural network for stage recognition (see Section 2) and is used

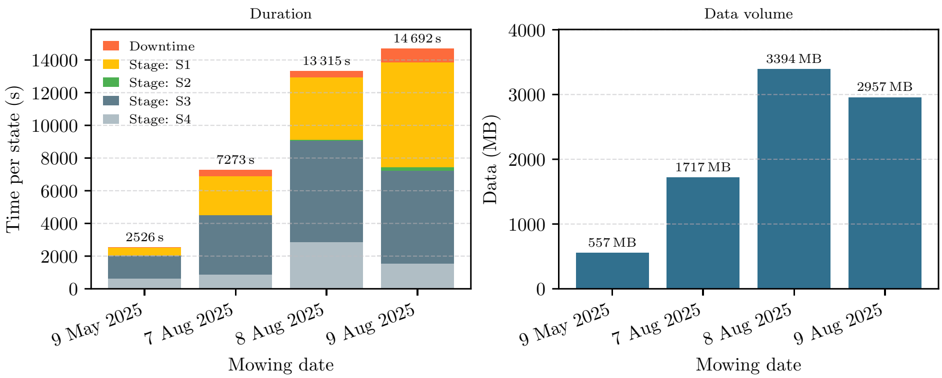

Figure 3. System performance across four mowing days. The left chart shows the duration of stages S1–S4 and RGB-D downtime; the right chart displays the total recorded data volume. Notably, the 9th of August shows lower data volume than the 8th of August, despite a longer total duration, due to increased sensor downtime.

here to simulate the adaptive system's behaviour under identical environmental conditions.

4. Evaluation

This section evaluates the performance and efficiency of the adaptive system under real-world conditions at Masch Lake. The analysis focuses on the system's dynamic response to plant density and provides a comparative assessment against a static baseline.

4.1. Adaptive System Performance Analysis

Figure 3 presents a comprehensive overview of the system's performance across four mowing days in the adaptive dataset. On the 9th August 2025, the system records the longest total mowing time. Stage S1 dominates the active runtime with a duration of 6416 s, indicating high aquatic plant density that causes the system to remain in Mode M1 for extended periods. Despite this prolonged operation, the system records a higher data volume on 8th August, capturing 3394 MB compared to 2957 MB captured on the 9th August. This discrepancy results from extended and unintended RGB-D sensor downtime: on 9th August, the sensor remains inactive for more than 13 minutes, whereas on 8th August, downtime amounts to just over 6 minutes.

While the exact cause of sensor downtime remains unconfirmed, system logs indicate temperature-related warnings and device disconnects. This suggests that thermal instability or internal sensor errors contribute to the failure. These issues will be further investigated in future deployments to improve system reliability.

In addition to the sensor downtime, the system exhibits notably lower average acquisition frequencies on 9th August compared to the previous day. Specifically, frequencies for RGB dropped from 1.01 Hz to 0.75 Hz, NIR from 0.97 Hz to 0.76 Hz, and the point cloud from 0.34 Hz to 0.16 Hz. These technical issues,

Table 1. Average system metrics per stage (9th Aug 2025). Values are shown as: RGB / NIR / Point Cloud (PC). Reduced activity in S3/S4 illustrates adaptive resource efficiency and system sensitivity

Metric	S1	S2	S3	S4
Duration (s)	6415.5	218.7	5666.9	1554.3
Freq. (Hz)	1.0/1.1/0.2	1.0/1.0/0.1	0.5/0.5/0.0	0.5/0.5/0.0
Vol. (MB)	292/246/1826	14/10/32	138/107/**154***	44/32/**62***

likely stemming from the aforementioned thermal instability, directly cause this frequency reduction. As the exact origin remains under investigation, we defer a detailed analysis of these correlations to future work. This reduction leads to fewer data packets generated per unit time, directly impacting the overall data volume. The combination of prolonged sensor inactivity and decreased acquisition frequency explains why the system captures less data on 9th August, despite spending more time in Stage S1.

For the quantitative evaluation of correct stage detection, we first assess the mower state (S1/S3 vs. S4). Over the full four days of recording, the system reaches an accuracy of 0.9062. The *down* state (S1/S3) achieves high precision (0.9983) and good recall (0.8874), resulting in an F1 score of 0.9396. However, the *up* state (S4) shows good recall (0.9931) but low precision (0.6569).

Next, we evaluate the detection of aquatic plants, which are relevant for S1 and S3. On a pixel-wise scale for the whole dataset, the detection reaches an accuracy of 0.7611. S1 achieves high precision (0.9850) but moderate recall (0.7285), resulting in an F1 score of 0.8376. This indicates a bias towards high-confidence detection, favouring precision over recall.

To illustrate the practical consequences of these classification weaknesses, we focus now on 9th August 2025, the day with the longest mowing duration. Table 1 summarises these results, listing the acquisition frequency, duration, and total data volume for each data type and mowing stage. These metrics reflect how the system dynamically adjusts its data acquisition strategy based on contextual relevance.

The observed frequencies align with the system's throttling configuration, as defined in the Throttling Module. In Stage S1, the system prioritises data acquisition by recording at the highest average frequencies: RGB and NIR channels operate at over 1 Hz, while point cloud data is recorded at 0.21 Hz, resulting in the largest data volume for this stage. During Stage S2 (unloading stage), RGB and NIR frequencies remain close to 1 Hz. We observe the point cloud frequency dropping to 0.11 Hz, although the throttling configuration does not mandate this reduction. In contrast, the system conserves resources during Stages S3 and S4 by lowering all frequencies below 0.52 Hz, with point cloud data dropping to just 0.02 Hz in Stage S3.

The presence of point cloud data in Stages S3 and S4 deviates from the intended design, as the system should disable this data type during non-mowing stages. Specifically, the system records 113 point cloud messages in Stage S3 and 45 in Stage S4. This anomaly suggests that the adaptive logic briefly switches into high-acquisition Mode M1 in response to transient aquatic plant detections before reverting to Mode M2. During these rapid transitions, the Throttling Module may

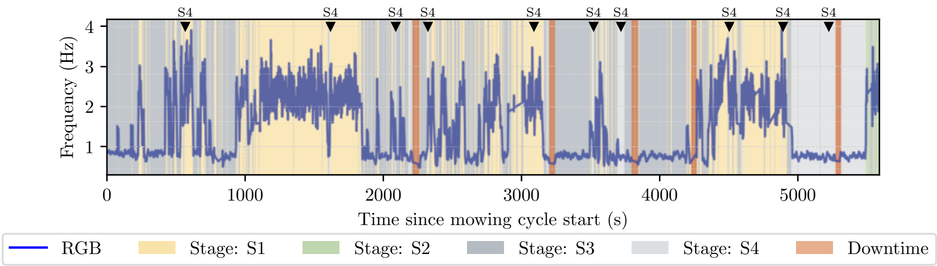

Figure 4. Stage-dependent frequency adaptation over a single mowing cycle (∼5500 s). The plot tracks RGB acquisition frequency, with background shading indicating stages S1–S4 and sensor downtime. Arrows highlight instances where Stage S4 was incorrectly triggered during active mowing.

not immediately suppress point cloud publication, leading to a small number of unintended or residual messages.

These findings confirm the system's adaptive capabilities while also revealing areas for improvement. In particular, the presence of point cloud data during non-mowing stages highlights the need for stabilisation mechanisms that prevent premature mode transitions in response to short-lived changes in aquatic plant density.

To analyse the system's adaptive behaviour in detail, this evaluation focuses on a representative mowing cycle of approximately 5500 s, selected from a complete mowing process on 9th August 2025 consisting of three cycles as shown in Figure 4. This cycle includes all relevant stages (S1–S4) and follows the same sequence as the others, making it suitable for demonstrating the system's dynamic response.

Figure 4 demonstrates that the system increases its frequency during Stages S1 (yellow) and S2 (green) and reduces it during Stages S3 (dark gray) and S4 (light gray). However, two key issues emerge: (i) Frequency instability within each stage: The acquisition frequency fluctuates rather than stabilising at a consistent rate, indicating that the system reacts to minor input variations instead of maintaining a fixed frequency per stage. (ii) Incorrect stage transitions: Stage S4 occasionally appears between Stages S1 and S3, despite ongoing mowing activity. These misclassifications likely stem from limitations in the mower state detection algorithm, which estimates the mower's state using the average point cloud distance within a predefined Region of Interest (ROI). A static threshold, empirically calibrated to represent the mower being parallel to the water surface, determines the state: values above the threshold indicate *down*, while lower values indicate *up*. However, dense plants or nearby objects can distort the point cloud and lower the average distance, leading to false detections of the mower being raised.

These findings highlight the need for more robust stage detection and frequency stabilisation. The current system deliberately relies on existing onboard sensors to identify mowing stages to minimise hardware complexity. However, the results show that this approach does not achieve sufficient accuracy. Future work should therefore explore the integration of additional sensors, such as an echo sounder, to improve the reliability of mower state classification.

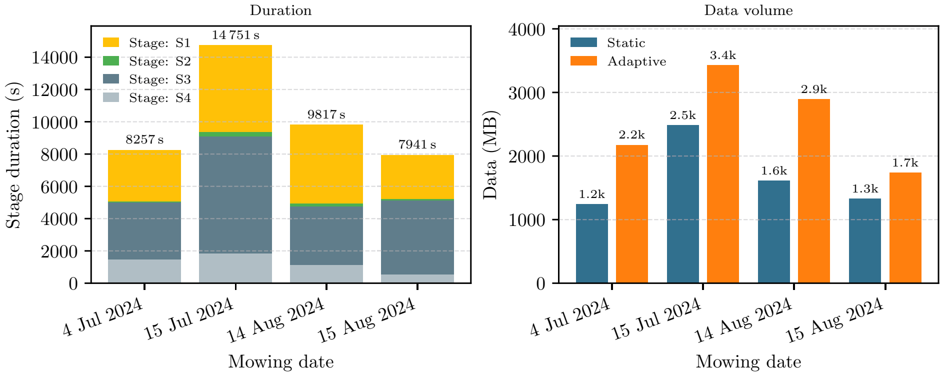

Figure 5. Comparison of static and adaptive systems over four mowing days. The left chart shows stage durations (S1–S4); the right chart displays the resulting data volumes. The adaptive system's responsiveness to plant density leads to extended S1 phases and increased data acquisition.

Table 2. Static vs. adaptive system comparison (15th July 2024). Format: RGB / NIR / PC.

Metric	**Static**	**Adaptive**	**Change (%)**
Messages	13.2k / 13.1k / 792	12.0k / 11.9k / 1.7k	-8.8 / -9.6 / +110.7
Rate (kB/s)	51.0 / 45.8 / 72.1	50.2 / 39.8 / 415.2	-1.5 / -13.2 / +476.2
Freq. (Hz)	0.87 / 0.87 / 0.05	0.94 / 0.93 / 0.29	+7.6 / +6.7 / +453.3

4.2. Comparison of the Adaptive and Static System

In this section, we apply the adaptive system by simulating its behavior on the previously defined static dataset. As this baseline was recorded at fixed frequencies (1 Hz for RGB and NIR; 0.1 Hz for point cloud), it allows for a direct, quantitative comparison with the adaptive approach under identical environmental conditions.

Figure 5 illustrates the differences between the simulated adaptive system and recorded static system. The left plot in Figure 5 shows the daily duration of each mowing stage (S1–S4) for the simulated adaptive system. For example, on the 15th of July 2024, the prolonged Stage S1 indicates high vegetation densities. The right plot in Figure 5 compares the total data volume generated by both systems. The adaptive system consistently produces more data by increasing acquisition frequency during relevant stages, thereby capturing more informative content.

Table 2 summarises key metrics for RGB, NIR, and point cloud data, including the number of messages, acquisition frequency, and data rate for the 15th July 2024.

The adaptive system records fewer RGB and NIR messages but substantially more point cloud data compared to the static system. This is a result of the targeted activation of point cloud acquisition during the relevant mowing stages (S1 and S2) while suppressing it during transitions (S3 and S4). Overall, the adaptive system produces 37.75 % more data.

Figure 5 and Table 2 reflect a behaviour reminiscent of the Jevons paradox [20]. Although the adaptive system is designed to optimise data acquisition by prioritising relevant information and reducing redundancy, its responsiveness to

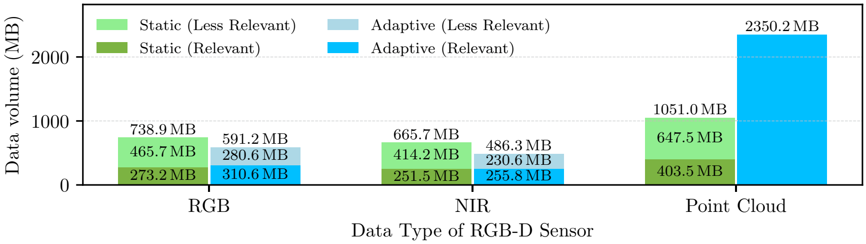

Figure 6. Data volume comparison by relevance for static and adaptive systems. The adaptive approach reduces less relevant data while prioritising acquisition during critical stages, most notably for the point cloud.

environmental conditions leads to increased recording activity during high-density stages. As a result, the total data volume exceeds that of the static system. This illustrates how efficiency improvements can, under certain conditions, result in higher overall resource consumption.

However, if we define data recorded during the Stages S1 and S2 as relevant, as they involve active aquatic plant detection and transport, the adaptive system records less data. Thus, Figure 6 compares the data volume across relevant and less relevant stages for both systems. The static system records substantial point cloud data during low-relevance stages, resulting in low information density. In contrast, the adaptive system suppresses point cloud acquisition entirely in these stages, ensuring that all recorded 3D data is contextually meaningful.

5. Conclusion

This study presents an intelligent sensor system for adaptive data acquisition for aquatic environmental management use cases. Field evaluations at the Masch Lake confirm that the system effectively captures critical data during ecologically relevant stages, while suppressing data acquisition during less relevant stages. Compared to a static system, the adaptive approach collects a significantly larger volume of data, with an increase of up to 37.75 %.

While this study focuses on artificial lakes, the underlying adaptive concept applies to various domains requiring contextual sensor adjustment. To address the increased data volume, reminiscent of the Jevons paradox, future optimisation must balance acquisition frequency to maximise information content without resulting in excessive resource consumption. Further efficiency gains may include adaptive image resolution or advanced compression. To enhance reliability, we aim to improve stage recognition by evaluating alternative algorithms and broader datasets. Additionally, integrating further sensor modalities, such as an echo sounder, will ensure more accurate state classification and prevent false detections, ultimately enhancing system stability in challenging environments.

Acknowledgment

This work is done within the HAI-x project, which is funded by the BMFTR (Funding number: 16IW23003)

References

[1] Alippi C, Camplani R, Galperti C, Roveri M. A robust, adaptive, solar-powered WSN framework for aquatic environmental monitoring. IEEE Sensors Journal. 2010;11(1):45-55.

[2] Khatei A, Ganie PA, Posti R, Swamy JM. Aquatic Environmental Monitoring with Sensor Network. In: Information Technology in Fisheries and Aquaculture. Springer; 2025. p. 151-66.

[3] Sudriani Y, Sebestyén V, Abonyi J. Surface water monitoring systems—the importance of integrating information sources for sustainable watershed management. IEEE Access. 2023;11:36421-51.

[4] Adebayo AI, Olubanjo KT, Fadeke AM, Uyanah JJ, Zirra AT, Igbaoreto WA, et al. From static sampling to dynamic insights: The future of water quality monitoring with sensors, IOT, and drones. Science World Journal. 2024;20(1):454-66.

[5] Shu T, Xia M, Chen J, De Silva C. An energy efficient adaptive sampling algorithm in a sensor network for automated water quality monitoring. Sensors. 2017;17(11):2551.

[6] Algabroun H. Dynamic sampling rate algorithm (DSRA) implemented in self-adaptive software architecture: a way to reduce the energy consumption of wireless sensors through event-based sampling. Microsystem Technologies. 2020;26(4):1067-74.

[7] Lou P, Shi L, Zhang X, Xiao Z, Yan J. A data-driven adaptive sampling method based on edge computing. Sensors. 2020;20(8):2174.

[8] Li T, Xia M, Chen J, Gao S, de Silva C. A hexagonal grid-based sampling planner for aquatic environmental monitoring using unmanned surface vehicles. In: 2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC). IEEE Press; 2017. p. 3683–3688. Available from: https://doi.org/10.1109/SMC.2017.8123205.

[9] Tan YT, Kunapareddy A, Kobilarov M. Gaussian Process Adaptive Sampling Using the Cross-Entropy Method for Environmental Sensing and Monitoring. In: 2018 IEEE International Conference on Robotics and Automation (ICRA). IEEE Press; 2018. p. 6220–6227. Available from: https://doi.org/10.1109/ICRA.2018.8460821.

[10] Koparan C, Koc AB, Privette CV, Sawyer CB. Adaptive Water Sampling Device for Aerial Robots. Drones. 2020;4(1). Available from: https://www.mdpi.com/2504-446X/4/1/5.

[11] Rahmati M, Nadeem M, Sadhu V, Pompili D. UW-MARL: Multi-agent reinforcement learning for underwater adaptive sampling using autonomous vehicles. In: Proceedings of the 14th International Conference on Underwater Networks & Systems; 2019. p. 1-5.

[12] Dang T, Liu J. Adaptive Water Environment Optimization Strategy Based on Reinforcement Learning. Comput-Aided Des Appl. 2024;21:1-18.

[13] Huang W, Zhan J, Sun Y, Han X, An T, Jiang N. Context-aware adaptive sampling for intelligent data acquisition systems using DQN. arXiv preprint arXiv:250409344. 2025.

[14] Manss C, Shutin D, Wiedemann T, Viseras A, Mueller J. Decentralized multi-agent entropy-driven exploration under sparsity constraints. In: 2016 4th International Workshop on Compressed Sensing Theory and its Applications to Radar, Sonar and Remote Sensing (CoSeRa). IEEE; 2016. p. 143-7.

[15] Merkel D, et al. Docker: lightweight linux containers for consistent development and deployment. Linux j. 2014;239(2):2.

[16] Tholen C, Rodenbäck E, Nolle L, Rettig R, Stahl F. On the Development of a Pixel-Wise Plastic Waste Identification System for Multispectral Remote Sensing Applications. In: Bramer M, Stahl F, editors. Artificial Intelligence XLI. Cham: Springer Nature Switzerland; 2025. p. 47-60.

[17] Rodenbäck E, Manss C, Schneider J, Stahl F. Evaluation of Aquatic Plant Detection with Remote Sensing Using Onboard Camera. In: 2025 13th European Workshop on Visual Information Processing (EUVIP); 2025. p. 1-6.

[18] Akiba T, Sano S, Yanase T, Ohta T, Koyama M. Optuna: A next-generation hyperparameter optimization framework. In: Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining; 2019. p. 2623-31.

[19] Hormann K, Agathos A. The point in polygon problem for arbitrary polygons. Computational geometry. 2001;20(3):131-44.

[20] Alcott B. Jevons' paradox. Ecological economics. 2005;54(1):9-21.