

FAROOQ AYOUB DAR

Autonomous Pervasive Sensing
for Proactive Environmental
Sustainability



FAROOQ AYOUB DAR

Autonomous Pervasive Sensing
for Proactive Environmental
Sustainability



UNIVERSITY OF TARTU

Press

Institute of Computer Science, Faculty of Science and Technology, University of Tartu, Estonia.

Dissertation has been accepted for the commencement of the degree of Doctor of Philosophy (PhD) in Computer Science on March 31, 2026 by the Council of the Institute of Computer Science, University of Tartu.

Supervisor

Associate Prof. Huber Raul Flores Macario
Institute of Computer Science,
University of Tartu, Estonia

Opponents

Prof. PhD Mario Di Francesco
Department of Information and Communications Engineering,
Aalto University, Finland

Prof. PhD. Claudio Bettini
Department of Computer Science,
University of Milan, Italy

The public defense will take place on May 8, 2026 at 10:15 in Zoom.

The publication of this dissertation was financed by the Institute of Computer Science, University of Tartu.

ISSN 2613-5906 (print)

ISSN 2806-2345 (pdf)

ISBN 978-9908-57-187-4 (print)

ISBN 978-9908-57-188-1 (pdf)

Copyright © 2026 by Farooq Ayoub Dar

University of Tartu Press

<http://www.tyk.ee/>

*To my loving parents for their unconditional support,
to my siblings for always being there,
to my dearest wife for her endless love,
and to my son, who inspires me every day.*

ABSTRACT

Pervasive sensing and autonomous technologies present a transformative opportunity to address urgent challenges in environmental sustainability, including pollution, climate change, and large-scale ecosystem degradation. Central to this vision is the integration of sensor-rich systems and intelligent computation for continuous, scalable, and context-aware environmental monitoring. This thesis explores how affordable, adaptive, and resilient sensor networks—augmented by advanced autonomous processing—can support sustainable management of natural resources through real-time data acquisition and automated response.

Driven by the need for robust and cost-effective monitoring solutions in diverse and often inaccessible settings, this research investigates how pervasive sensing infrastructure can be paired with autonomous frameworks to overcome persistent obstacles of scalability, accuracy, and energy efficiency. The core research question guiding this work is: *How can pervasive, low-cost sensors and autonomous systems be effectively leveraged to enable dynamic, high-resolution monitoring and pollutant detection at scale?* Accordingly, the thesis investigates three primary avenues: innovative approaches to material identification, autonomous detection of plastic pollution, and underwater data analysis empowered by fog computing.

First, the thesis presents *MIDAS*, an innovative material identification technique based on thermal dissipation patterns, capable of robust classification across material types and application scenarios. *MIDAS* demonstrates up to 83% accuracy, outperforming conventional approaches and adapting reliably to variable user and environmental conditions.

Second, the *LIZARD* system is introduced – a multi-modal sensing pipeline that fuses thermal and optical data for the autonomous detection of plastic pollution. By targeting macro, meso, and microplastics, *LIZARD* bridges a critical monitoring gap and achieves up to 80% accuracy in real-world deployments, showcasing the feasibility of affordable, high-precision environmental sensing.

Third, a novel underwater fog computing framework is demonstrated, repurposing commercial off-the-shelf (COTS) devices as *micro-clouds* for distributed, real-time data processing in challenging aquatic and other remote environments. This solution enables cost-effective deployment and delivers resilient performance by adapting general-purpose technology for domain-specific needs.

Together, these contributions advance the state-of-the-art in environmental monitoring by demonstrating how pervasive sensors and autonomous technologies can deliver actionable insights through scalable, energy-efficient and application-tailored systems. By bridging the gaps in the practicality and analytical precision of deployment, this work lays a foundational framework for the next generation of sustainable and intelligent environmental management.

CONTENTS

List of original publications	18
1. Introduction	21
1.1. Background and Motivation	22
1.2. Research Challenges	23
1.3. Research Contributions	24
1.3.1. MIDAS	24
1.3.2. LIZARD	24
1.3.3. Micro-clouds	25
1.4. Thesis Structure	26
2. Background	27
2.1. Environmental Sustainability: A Multidimensional Concern	27
2.1.1. Water Pollution	28
2.1.2. Air Pollution	28
2.1.3. Noise Pollution	29
2.1.4. Biodiversity Loss	29
2.1.5. Ecosystem Threat Level Matrix	29
2.2. Data Collection and Analytics for Environmental Sustainability	31
2.2.1. Water Pollution Monitoring	31
2.2.2. Air Pollution Monitoring	32
2.2.3. Noise Pollution Monitoring	33
2.2.4. Biodiversity Monitoring	34
2.3. Challenges in Environmental Sustainability Monitoring	35
2.3.1. Technical Challenges	36
2.3.2. Scalability Challenges	36
2.3.3. Data Collection Challenges	37
2.3.4. Energy Challenges	37
2.3.5. Deployment Challenges	37
2.3.6. Connectivity Challenges	38
2.3.7. Data Processing and Interpretation Challenges	38
2.3.8. Regulatory and Compliance Challenges	38
2.3.9. Societal and Institutional Barriers	39
2.4. Defining Pervasive Sensing, Sensors, and Autonomous Technologies	39
2.4.1. Pervasive Sensing: The Foundational Philosophy	39
2.4.2. The <i>Senses</i> of the System: Types of Sensors	40
2.4.3. Scalable Deployment and Intelligent Processing: Autonomous Technologies	41
2.4.4. Integration of Computing Layers in Sensor-Based Environ- mental Monitoring	42
2.4.5. The Path Forward: An Integrated and Responsive Framework	44

2.5. Opportunities and Challenges of Sensor-Based Environmental Monitoring	44
2.5.1. Opportunities and Benefits	44
2.5.2. Limitations and Challenges	46
2.6. Case Studies: Real-World Applications and Impact Assessment . .	47
2.6.1. Air Quality Monitoring: Community-Based Indoor Air Quality System	47
2.6.2. Water Quality Monitoring: Precision Agriculture Water Management	48
2.6.3. Environmental Plastic Monitoring: From Traditional Assessment to Autonomous Marine Systems	49
2.6.4. Wildlife Monitoring: AI-Enhanced Acoustic Biodiversity Assessment Networks	49
2.6.5. Cross-Domain Analysis: Integrated Environmental Monitoring Framework	50
2.7. Summary	51
3. MIDAS - Thermal Dissipation for Object Characterization	53
3.1. Introduction	53
3.2. Related Work	54
3.3. Feasibility assessment	55
3.3.1. Testbed	55
3.3.2. Plastic thermal fingerprint dissipation	55
3.3.3. Other Thermal fingerprint dissipations	56
3.4. MIDAS pipeline	57
3.5. Robustness of Thermal Dissipation Fingerprints	59
3.5.1. Experimental setup	60
3.5.2. Baselines	62
3.6. Results	63
3.6.1. Differences in Thermal Fingerprints	63
3.6.2. Other Factors Influencing Thermal Dissipation	64
3.6.3. Dissipation Time Classification Performance	67
3.6.4. Comparison Against Other Approaches	67
3.6.5. Effect of thickness on thermal dissipation	69
3.7. Thermal Dissipation Fingerprint for Multiple Objects	72
3.7.1. From Single to Multiple Objects	72
3.7.2. Evaluation	72
3.8. Discussion	76
3.9. Applications for Environmental Sustainability: Outcomes and Implications	78
3.10. Summary	80

4. LIZARD - Pervasive Sensing for Plastic Litter Monitoring	81
4.1. Introduction	81
4.2. Related work	83
4.3. Motivation	84
4.4. LIZARD Sensing Pipeline	86
4.4.1. Phase 1 - Thermal imaging	88
4.4.2. Phase 2 - Light reflectivity	89
4.5. Experimental Setup	91
4.6. Results	94
4.7. Real-world practicability	101
4.8. Discussion	103
4.9. Applications for Environmental Sustainability: Outcomes and Impli- cations	105
4.10. Summary	107
5. Micro-Cloud: Fog Computing Framework for Underwater Pervasive Data Science	108
5.1. Introduction	108
5.2. Related Work	111
5.3. Feasibility Experiment	115
5.4. Underwater Micro-Cloud Design	118
5.4.1. Architecture	119
5.4.2. Components	119
5.4.3. Fog Provisioning Underwater	121
5.5. Computational Benchmarks	122
5.6. Results	124
5.6.1. Underwater Computing Performance	125
5.6.2. Underwater Capacity Performanace	126
5.6.3. Underwater Energy Consumption	127
5.6.4. Submersed Fog to Cloud Performance	128
5.6.5. Offloading from Underwater to the surface	128
5.7. Underwater Optical Communications	129
5.8. Ocean Deployment	131
5.9. Discussion	135
5.10. Applications for Environmental Sustainability: Outcomes and Impli- cations	139
5.10.1. Application Scenarios	139
5.10.2. Sustainability Impact	140
5.11. Beyond Aquatic Environments: Repurposing Micro-Clouds (E-Waste Smartphones) as Tiny Data Centers	140
5.11.1. Technical Implementation and Architecture	141
5.11.2. Applications and Deployment Scenarios	142
5.11.3. Proof-of-Concept Demonstration and Performance Analysis	143

5.11.4. Environmental Impact and Sustainability Benefits	143
5.12. Summary	144
6. Conclusions and Future Directions	146
6.1. Conclusions	146
6.2. Research Scope and Limitations	146
6.2.1. General Limitations	146
6.2.2. Specific Limitations of the Proposed Systems	147
6.3. Future Directions	148
Bibliography	150
Acknowledgements	181
Sisukokkuvõte (Summary in Estonian)	182
Curriculum Vitae	184
Elulookirjeldus (Curriculum Vitae in Estonian)	185

LIST OF FIGURES

1. Environmental Sustainability: A Multidimensional Concern	28
2. Computing layers in sensor-based environmental monitoring. . . .	43
3. Integration of pervasive sensing and autonomous solutions for environmental sustainability – highlighting underwater monitoring for reef, biodiversity monitoring; AGVs for pollution and litter detection, and aerial/terrestrial systems for waste management and deforestation monitoring.	45
4. Selected waste materials for preliminary experiments: A (Beer Can), B (Ceramic Cup), C (Takeaway Box), D (Plastic bottle), E (Glass Bottle), F (Coffee Cup), G (Plastic Cup), H(Cigarette Butt), I (Glass Jar), J (Milk pack), K (Aerosol Can), L (Rubber glove), M (Metal spoon) and N (Face mask).	56
5. Dissipation time of thermal fingerprints in different plastic materials and waste material using two different devices: (a) Grouped by RIC code, (b) Thermometer scanner FLIR TG267 (baseline), and (c) Smartphone CAT s60.	57
6. Processing pipeline of material classification based on dissipation time of thermal fingerprints.	58
7. Dissipation time of thermal fingerprint for two different objects: (a) Cardboard cup and (b) Cigarette butt.	58
8. Experimental conditions: (a) Rigid interaction, used in Fixed-hold and (b) Free interaction, used in Natural and Quick holds.	60
9. Experimental testbed and protocol steps: (a) Object obtains baseline temperature, (b) Object habituates to ambient temperature, and (c) Participant performs the experiment and puts the object in the marked target to measure its thermal fingerprint.	62
10. Thermal transferred conditions applied over three different objects (plastic bottle, cardboard cup and cigarette butt).	64
11. Influence of internal and external temperatures.	66
12. Baselines: (a) Trained model PlasticNet to identify and separate plastics from other object materials, and (b) Light reflectivity values of different materials measured with a photo-resistor.	69
13. Thickness of object materials analyzed with thermal dissipation fingerprint.	70
14. Effect on thermal dissipation with increase in the thickness of material.	71
15. Detection of multiple thermal dissipation fingerprints. a) Object arrangements, b) Extraction of thermal dissipation vectors.	73
16. Response time of multiple object identification when using MIDAS. (same video configuration). The bars within the figure correspond to different individual objects (A) or combination of objects (two objects: B, three objects: C, four objects: D) as given in Table 7.	75

17. Impact on response time when using videos of different length and frame rate (Disperse objects case). The bars within the figure correspond to different individual objects (A) or combination of objects (two objects: B, three objects: C, four objects: D) as given in Table 7.	76
18. Impact on response time when using videos of different length and frame rate (Agglomerated objects case). The bars within the figure correspond to different combinations of objects (two objects: B, three objects: C, four objects: D) as given in Table 7.	77
19. Combination of macro-plastics and micro-plastics (agglomerated and dispersed).	86
20. Thermal dissipation times for (a-d) 20 plastic samples, (e) Micro and meso plastics, and (f) Selected macro plastics.	87
21. Plastic fragment sensing and monitoring pipeline.	88
22. Wavelet transformation of light sensor data (top), change point scores (middle) and change point detection (bottom).	91
23. A. plastic types: (a) 1: Micro (<5 mm), (b) 2: Meso (5 mm to 15mm), and (c) 3: Macro plastics (> 15 mm); B. Plastic size comparison using coins as a reference; C. Arrangements of micro, meso, and macro plastic samples used in our analysis.	92
24. Lizard design alternatives.	93
25. Lizard prototype, (a) Main Components; (b) Underlying Sensor array; (c) Lizard in action	93
26. Light reflectivity values: (a),(b) plastic fragments of different sizes; (c),(d) 20 different plastic fragment samples.	95
27. Litter samples distribution. (a) Amount of plastic fragment samples identifiable in the wild (TACO dataset). (b) Litter distribution in extracted images	96
28. Light reflectivity and thermal dissipation values of plastic samples (a) thermal dissipation times for macro plastics, (b) Light reflectivity from different sampling points	97
29. Light reflectivity for different categories.	98
30. a) Comparison with ground truth b) Light values of same plastic fragments in different contexts	99
31. Baseline light reflectivity of different backgrounds	99
32. Light reflectivity values of plastic samples.	100
33. Clustering and classification results - (a) k-means clustering, (b) Confusion matrix for random forest classifier	100
34. Removal of plastic fragments via visual inspection. (a) Illustration of the method, (b) Removal rate.	102
35. LIZARD deployment. Sampling locations - backgrounds: (a) & (d) rocky (b) & (e) iron, (c) & (f) soil.	102
36. a)-c) Ambient luminosity during the day, d-f) Baseline background light values during the day.	103

37. Plastic fragment detection in the Wild	104
38. Improvement on classification accuracy.	104
39. Complexity of moving underwater collected data to computing infrastructure. (a) Surface-based deployments can only support underwater infrastructure that is close to coastal areas, (b) Major areas of marine litter are far from data centers and also in areas that are far from coastal regions, (c) Inland regions, such as lakes, rivers and streams, also can be poorly covered by communication infrastructure.	109
40. Capacity results of submerged micro-clouds when handling multiple users, a) Response time (Baseline - no water), b) Depth-1, c) Depth-2, d) Latency (Baseline - no water), e) Depth-1, f) Depth-2.	117
41. Different types of water motion influence connectivity to the submerged micro-cloud, a-b) Low, c-d) High (induced by water vehicles operating nearby).	118
42. Micro-cloud prototype: (a) Internal components, (b) Waterproof encasing, and (c) Deployment in the wild.	119
43. Initial micro-cloud deployment, (a) Micro-cloud architecture, (b) augmenting individual devices, (c) collaborating to improve fog processing performance, (d) dedicated micro-clouds in a location that can be used by underwater IoT devices in a contact-based manner.	120
44. Latencies of different applications, a) Computation Latency, b) Communication Latency	125
45. Results of computation and communication latency for concurrent users, (a) Baseline computation latency, (b) Underwater computation latency, (c) Baseline communication latency, (d) Underwater communication latency.	127
46. Results of energy consumption underwater, a) Baseline with multi-meter, b) Energy monitor (outside the water) c) Energy monitor deployed underwater.	127
47. Impact of concurrent users using iPokeman, a) Response time, b) Response latency, c) Response time underwater, c) Response latency underwater.	128
48. Response time and energy consumption for both apps (a-b) Stream app, (c-d) Code app.	129
49. Testbed for analysis of underwater communication with light.	130
50. Water motion (Calm water), a) Accumulated motion captured through all axes (variance), b) Experience motion (one axis only).	131
51. Light communication performance under different levels of induced water motion, a) Overall motion using AirPump (all axes) b) Motion using AirPump (one axis), c) Overall motion using a more potent motor (all axes) d) Motion using a more potent motor (one axis).	132

52. Underwater Micro-Cloud Computation Units (MCCUs) consisting of individual mobile devices, communicating using off-the-shelf integrated wireless interfaces while collaborating in the execution of a task, deployed in open sea at 8m depth. From left to right, top to bottom: (a) Master and worker within same unit, (b) Master to the left and two workers to the right, (c) Seabed deployment and (d) SCUBA diver survey.	133
53. Obtained sensor data during ocean deployment experiment time. From left to right: (a) Triaxial accelerometer, (b) CPU temperature, (c) RSSI, and (d) RAM Percentage. Red line depicts the time at the surface, while blue line is the time underwater. Vertical lines indicate the SCUBA diving time for underwater test.	135
54. (a) Vision, (b) tiny data centers approach overview and (c) use case application. Powering and assembling stage: (A) Basic tiny data center: A single phone powered using a 9V battery connected to an LM2596 DC-DC, (B) External power source, (C) LM2596 DC-DC power regulator, (D) Phone rear-view. Final stage: (E) Final prototype: 4 phones and common power regulator.	142

LIST OF TABLES

1. IUCN Red List Threat Classification System [242]	30
2. Ecosystem impacts by pollution type.	30
3. Summary of key challenges in environmental sustainability monitoring	36
4. Computing Layers in Sensor-Based Environmental Monitoring . .	42
5. Material classification accuracy (%) in different experimental conditions. Model data → Predicted. Classification Method: Random Forest (RF), Support Vector Machine (SVM) and Multi-layer Perceptron (MLPC).	68
6. Material thickness in each layer for Cardboard, Plastic and Mask respectively.	70
7. Table explaining grouping arrangement for prediction of Multiple Objects	74
8. State-of-the-art for identification of plastics and potential integration with autonomous vehicles (AV).	84
9. Summary of most relevant work for fog provisioning. Our work is the first to offer fog computing support consisting of multiple devices that is ready for underwater deployment and that uses low-cost COTS components.	112
10. Comparison of underwater computation times during the image classification for 10 minutes. Increasing the number of workers expectedly speeds up computations which implies that the connectivity between devices is sufficiently stable to operate the micro-cloud despite adverse water conditions.	136

LIST OF ABBREVIATIONS

Acronyms

- AGVs** Autonomous ground vehicles. 36, 44, 78, 85, 86, 105
- AGVs** Autonomous underwater vehicles. 41, 49, 52, 115, 119, 140
- AI** Artificial Intelligence. 21, 32, 34
- API** application programming interface. 43, 47
- ASVs** autonomous surface vehicles. 31, 49
- BOD** biological oxygen demand. 32
- CNNs** Convolutional Neural Networks. 34, 35, 50, 59, 62
- CO** carbon monoxide. 32
- COTS** commercial off-the-shelf. 6, 25, 53, 57, 108, 110, 139, 144
- FLIR** forward looking infrared. 53
- GDPR** General Data Protection Regulation. 38
- IoT** Internet of Things. 21, 27, 31, 47, 49
- IoUT** Internet of Underwater Things. 113, 114
- IUCN** International Union for Conservation of Nature. 29
- LDRs** light-dependent resistors. 93
- MEMS** micro-electro-mechanical systems. 44
- MFCCs** Mel–frequency cepstral coefficients. 34
- MLPC** Multi-Layer Perceptron Classifier. 59, 67
- PAM** Passive acoustic monitoring. 34, 121
- PM2.5** particulate matter. 28, 32, 48
- RF** Random Forest. 59, 67, 91
- RFID** Radio-Frequency IDentification. 53
- RIC** Resin Identification Code. 55
- ROIs** Region of Interest. 72, 90
- ROVs** remotely operated vehicles. 119, 120
- RPi4** Raspberry Pi 4. 115
- RSSI** Received Signal Strength Indicator. 134
- SDG** Sustainable Development Goals. 107
- SVMs** Support Vector Machines. 32, 59, 67, 91

TSS total suspended solids. 32

UAVs Unmanned Aerial Vehicles. 21, 27, 31, 35, 36, 41, 78

UWSNs Underwater Sensor Networks. 114

ViTs Vision Transformers. 59

VOCs volatile organic compounds. 32, 48

WSNs Wireless Sensor Networks. 21, 31, 34, 41, 46

LIST OF ORIGINAL PUBLICATIONS

Publications included in the thesis

1. Emenike, Hilary, **Farooq Dar**, Mohan Liyanage, Rajesh Sharma, Agustin Zuniga, Mohammad A. Hoque, Marko Radeta, Petteri Nurmi, and Huber Flores. "Characterizing everyday objects using human touch: Thermal dissipation as a sensing modality." In IEEE International Conference on Pervasive Computing and Communications, pp. 1-8. IEEE, 2021. DOI: 10.1109/PERCOM50583.2021.9439120
2. **Farooq Dar**, Hilary Emenike, Zhigang Yin, Mohan Liyanage, Rajesh Sharma, Agustin Zuniga, Mohammad A. Hoque, Marko Radeta, Petteri Nurmi, and Huber Flores. "The MIDAS touch: Thermal dissipation resulting from everyday interactions as a sensing modality." *Pervasive and Mobile Computing*, Vol. 84, 2022. DOI: 10.1016/j.pmcj.2022.101625 (*Extension for Publication 1*)
3. **Farooq Dar**, Mayowa Olapade, Abdul-Rasheed Ottun, Zhigang Yin, Mohan Liyanage, Agustin Zuniga, Monica Passanantti, Sasu Tarkoma, Petteri Nurmi, and Huber Flores. "LIZARD: Pervasive sensing for autonomous plastic litter monitoring." In IEEE/ACM Ninth International Conference on Internet of Things Design and Implementation, pp. 37-48. IEEE, 2024. DOI: 10.1109/IoTDI61053.2024.00008
4. **Farooq Dar**, Mohan Liyanage, Marko Radeta, Zhigang Yin, Agustin Zuniga, Sokol Kosta, Sasu Tarkoma, Petteri Nurmi, and Huber Flores. "Upscaling fog computing in oceans for underwater pervasive data science using low-cost micro-clouds." *ACM Transactions on Internet of Things*, Vol. 4, No. 2, pp. 1-29, 2023. DOI: 10.1145/3575801
5. Ngoy, Perseverance, **Farooq Dar**, Mohan Liyanage, Zhigang Yin, Ulrich Norbistrath, Agustin Zuniga, João Pestana, Marko Radeta, Petteri Nurmi, and Huber Flores. "Supporting Sustainable Computing by Repurposing E-Waste Smartphones as Tiny Data Centers." *IEEE Pervasive Computing*, Vol. 24, No. 1, pp. 53-60, 2025. DOI: 10.1109/MPRV.2025.3541558

Publications not included in the thesis

6. **Farooq Ayoub Dar**, Mayowa Olapade, Abdul-Rasheed Ottun, Zhigang Yin, Mohan Liyanage, Ulrich Norbistrath, Marko Radeta et al. "TOAD: Profiling and Evaluating 3D Printed IoT Rapid Prototype Designs." *ACM Transactions on Internet of Things*, Vol. 6, No. 3, pp. 1-31, 2025. DOI: 10.1145/3724128
7. **Farooq Dar**, Mohan Liyanage, Mayowa Olapade, Zhigang Yin, Abdul-Rasheed Ottun, Adeyinka Akintola, Francisco Airton Silva, and Huber

Flores. "Demo Abstract: PRINCE: Device Energy Estimation with a Single Photo." In IEEE/ACM Ninth International Conference on Internet-of-Things Design and Implementation, pp. 231-232. IEEE, 2024. DOI: 10.1109/IoTDI61053.2024.00031

Other published work of the author

8. Liyanage, Mohan, **Farooq Dar**, Rajesh Sharma, and Huber Flores. "GEESE: Edge computing enabled by UAVs." *Pervasive and Mobile Computing*, Vol. 72, 2021. DOI: 10.1016/j.pmcj.2021.101340
9. Yin, Zhigang, Mayowa Olapade, Mohan Liyanage, **Farooq Dar**, Agustin Zuniga, Naser Hossein Motlagh, Xiang Su et al. "Toward city-scale litter monitoring using autonomous ground vehicles." *IEEE Pervasive Computing*, Vol. 21, No. 3, pp. 74-83, 2022. DOI: 10.1109/MPRV.2022.3152926
10. Zuniga, Agustin, Mayowa Olapade, Naser Hossein Motlagh, Mohan Liyanage, Zhigang Yin, **Farooq Dar**, Ngoc Thi Nguyen et al. "Low-cost sensing for environmental sustainability." *IEEE Pervasive Computing*, Vol. 23, No. 4, pp. 76-86, 2024. DOI: 10.1109/MPRV.2024.3448199
11. Yin, Zhigang, Mohan Liyanage, Abdul-Rasheed Ottun, **Farooq Dar**, Mayowa Olapade, and Huber Flores. "Demo Abstract: A Smart Ring Monitoring Your Health using Hand-grip Strength." In *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems*, pp. 486-487, 2023. DOI: 10.1145/3625687.3628395

Author's Contribution to the Publications

- **MIDAS (Publications 1 and 2):** As the co-author in Publication 1 and lead author in the subsequent extension paper (Publication 2), the author played a central role in all stages of both studies. His responsibilities encompassed designing the conceptual pipeline, framing the theoretical foundations, coordinating the user study logistics, and leading experimentation, comprehensive data collection, analysis, and result interpretation. For Publication 1, he specifically developed the innovative, low-resource intensive ML pipeline designed to maintain high performance on resource-constrained devices. He also conducted the thermal dissipation experiments across distinct plastic samples, ensuring robust and replicable datasets. Building on this foundation, he assumed the lead author role for Publication 2, guiding the extension of the research. The central innovation of the extension paper involved expanding the experimentation to include multiple object scenarios—a pivotal advancement over the original study. He personally executed these experiments, managed the study content, and coordinated all aspects of the expanded analysis, which formed the main crux of the new publication.

- **LIZARD (Publications 3):** The author initiated the conceptual design and conducted preliminary experiments to establish the feasibility of object detection employing thermal imaging and light sensing modalities. All major experiments were independently conducted by the author. The principal experimental campaign comprised a comprehensive series of outdoor measurements carried out over a continuous ten-day period. These experiments were systematically performed at various intervals—early morning, afternoon, and late night—to ensure robust data collection across a wide range of environmental conditions. The author additionally led the experimental design and data acquisition, performed in-depth analyses, and contributed substantially to the writing and revision of the manuscripts.
- **Micro-clouds (Publications 4 and 5):** The author made significant contributions encompassing the conceptualization, experimental design, methodology, data collection, and analysis throughout the study. He independently conducted key experiments, ensuring rigorous and reproducible data acquisition. The author was responsible for developing the overarching theoretical framework and managing the experimental setup, including challenging field campaigns conducted under varied environmental conditions. In addition to leading data interpretation and validation, he played a central role in drafting, revising, and refining the manuscripts. Collaborative support was received for specific tasks, including apparatus setup and manuscript polishing, while the author maintained primary responsibility for the research progression and experimental outcomes. For the use-case scenario paper (Publication 5), the candidate led the data analysis efforts and contributed substantively as co-author.

1. INTRODUCTION

Environmental sustainability, as defined by the *Brundtland Commission* (World Commission on Environment and Development) in its report *Our Common Future*, refers to development that meets the needs of the present generation without compromising the ability of future generations to meet their own needs [288]. This foundational concept emphasizes the harmonious integration of economic development, environmental stewardship, and social equity, ensuring that progress today does not diminish the resources, ecosystem services, or opportunities available to future generations. The Brundtland Commission's definition has become the guiding principle for sustainable development policies worldwide, establishing environmental sustainability as a core consideration in addressing global challenges related to resource use, pollution, and ecosystem health.

The increasing severity of ecological challenges, such as climate change [167], air pollution [46], biodiversity loss [96], and resource depletion [30], necessitates innovative technological interventions. In recent years, advancements in pervasive sensing [233] and autonomous solutions have emerged as transformative tools for environmental monitoring and sustainability efforts. These technologies enable real-time, high-resolution data collection, fostering improved decision-making processes for ecosystem management and conservation [142]. Building upon these capabilities, the next generation of environmental monitoring systems increasingly leverages pervasive sensing and autonomous technologies to further enhance the scope, resolution, and efficiency of data collection.

Pervasive sensing refers to the deployment of sensor networks that continuously observe environmental parameters, offering valuable insights into pollution levels, natural resource consumption, and ecosystem health [272]. These systems rely on a diverse range of sensors, including Wireless Sensor Networks (WSNs) [12], low-power Internet of Things (IoT) devices [107], and edge-cloud computing frameworks. By integrating pervasive sensing with autonomous systems—such as drones, Unmanned Aerial Vehicles (UAVs) [194], and underwater sensor networks [133]—it becomes possible to enhance large-scale environmental data collection while minimizing human intervention. These technologies have already demonstrated their potential in applications such as climate monitoring [204], forest fire detection [194], urban sustainability [159], and precision agriculture [339].

Furthermore, emerging computing paradigms, including fog computing [261] and edge processing [250], are significantly enhancing the efficiency and scalability of environmental monitoring systems. By processing data closer to its source, these technologies reduce latency and energy consumption, making real-time monitoring feasible even in remote and resource-constrained environments. Additionally, machine learning and Artificial Intelligence (AI) [137] are playing a crucial role in extracting meaningful patterns from vast environmental datasets, enabling predictive analytics and informed decision-making [258].

Despite these advancements, challenges remain in terms of energy efficiency

[140], sensor coverage optimization, data accuracy, and system integration [228]. The reliability and long-term sustainability of these systems require further research in areas such as sensor calibration [196], power management, and interoperability between heterogeneous networks [109]. Moreover, ethical concerns related to data privacy, security, and regulatory compliance must be addressed to ensure responsible deployment of pervasive sensing technologies [158].

Building on these considerations, this thesis positions pervasive sensing and autonomous systems as core enablers of scalable environmental solutions. It explores the integration of pervasive sensing and autonomous systems in addressing key environmental sustainability challenges. It introduces solutions such as underwater fog computing, advanced sensing pipelines, and scalable autonomous sensing platforms as cost-effective tools for environmental monitoring and mitigation. By examining state-of-the-art developments, applications, and existing challenges, the thesis aims to advance sustainable environmental monitoring and management practices. The subsequent chapters detail the theoretical foundations, system designs, experimental validations, and implications of these contributions, highlighting their potential for practical, large-scale deployment.

1.1. Background and Motivation

The issue of environmental sustainability has emerged as one of the most pressing concerns of our time [273]. The rapid pace of industrialization, urbanization, and population growth has placed immense pressure on natural ecosystems [241]. Challenges such as pollution, habitat destruction, and climate change have significantly contributed to the degradation of the environment (Millennium Ecosystem Assessment, 2005) [312]. Among these challenges, the improper disposal and management of waste – especially plastic fragments – has become a significant threat, contaminating both land and aquatic ecosystems [246] [245] [244]. Microplastics, in particular, have drawn attention due to their persistence in the environment and their infiltration into food chains, posing risks to biodiversity and human health [168]. Furthermore, the disposal of daily waste into the environment and aquatic habitats pose a significant threat to the environment [130].

Technological innovations have created new possibilities to tackle these challenges [139]. Pervasive sensing, which involves deploying extensive sensor networks across large areas, allows for real-time monitoring of environmental parameters [155]. These insights are crucial for driving sustainability initiatives. Additionally, autonomous systems, including drones, underwater vehicles, and robotic platforms, enhance the efficiency and scope of monitoring and mitigation efforts [91]. By operating in diverse and often challenging environments, these technologies can play a pivotal role in addressing ecological issues.

However, despite these advancements, several obstacles remain that limit the broader application of pervasive sensing and autonomous systems [175]. High deployment costs, energy constraints of sensing devices [70], and the complexity

of real-time data processing are among the key challenges. Furthermore, many existing solutions are not scalable or robust enough to function effectively in remote or harsh conditions [232]. Overcoming these limitations is critical to unlocking the full potential of these technologies for promoting environmental sustainability.

Addressing these gaps not only holds promise for advancing technology but also aligns with global sustainability goals (UN SDGs) [31] [36], providing a pathway for innovative solutions to pressing ecological problems.

1.2. Research Challenges

The deployment of pervasive sensing and autonomous systems for environmental sustainability entails multifaceted challenges that must be addressed to realize robust, scalable, and sustainable monitoring solutions.

Environmental Complexity: Effective monitoring in remote terrestrial and aquatic environments demands resilient system architectures capable of operating reliably under extreme conditions—including high pressures, wide temperature fluctuations, corrosive elements, and limited or intermittent connectivity [65, 235]. The marine domain, in particular, imposes additional constraints such as biofouling, mechanical wear, and sensor degradation, which complicate long-term deployment and necessitate innovative protective and maintenance strategies.

Microplastic Detection and Material Characterization: Existing microplastic identification methodologies are largely constrained to laboratory settings and involve specialized, non-scalable instrumentation [23, 287]. The translation of these capabilities into deployable, autonomous, and pervasive sensing platforms for in-situ detection, particularly capable of resolving micro- and mesoplastics, remains an open technical challenge. Accurately characterizing heterogeneous materials through non-invasive sensing modalities such as thermal dissipation demands further advances to bridge sensitivity, specificity, and operational robustness.

Energy Efficiency and Power Autonomy: Sensor networks and autonomous monitoring platforms in inaccessible environments face stringent power constraints where traditional energy replenishment is impractical [54]. Achieving sustainable operation requires the development of ultra-low-power sensing and communication architectures complemented by energy harvesting techniques. Further, computational frameworks must balance energy consumption with processing demands to prolong operational lifetimes [140].

Real-Time Data Processing and Autonomous Decision-Making: The voluminous and heterogeneous data streams generated by pervasive environmental sensors necessitate efficient, real-time analytics and intelligent decision support [109, 262]. Architectures integrating edge and fog computing with advanced machine learning algorithms are critical to filter, aggregate, and interpret data locally, curtailing latency and bandwidth limitations while enabling adaptive and context-aware system behavior.

Sensor Reliability and Environmental Interference: Environmental variability—including fluctuating light conditions, water turbidity, thermal noise, and mechanical disturbances—frequently degrades sensor performance, challenging data accuracy and consistency [164]. Additionally, long-term deployments face challenges related to sensor drift, calibration loss, and component aging, necessitating robust calibration protocols and self-diagnostic capabilities to preserve data integrity over time [284].

Together, these challenges underscore the complexity inherent in advancing pervasive and autonomous environmental monitoring systems. This research endeavor addresses these gaps through the design and validation of innovative sensing frameworks, edge-enabled processing paradigms, and sustainable technology repurposing, thereby contributing toward scalable, resilient, and environmentally responsible monitoring solutions.

1.3. Research Contributions

This thesis makes significant contributions to the fields of pervasive sensing and autonomous systems. Rather than presenting isolated examples of environmental applications, the three core systems developed in this research – MIDAS, LIZARD, and Micro-clouds – were selected to address a continuous pipeline of the challenges outlined in Section 1.2. Specifically, they tackle the critical progression from fundamental sensor innovation, to scalable autonomous deployment, to resilient data processing at the edge.

1.3.1. MIDAS

Addressing the *Microplastic Detection and Material Characterization* challenges (Section 1.2) and broader *Technical Challenges* (Section 2.3), we introduce MIDAS, an innovative sensing solution that leverages thermal dissipation patterns resulting from human touch to characterize and identify everyday object materials. By capturing the transient heat transfer and modeling the decrease in thermal radiation, MIDAS achieves material classification with up to 83% accuracy. The system demonstrates robustness to variation across different users and shows the capability to detect thermal dissipation through objects up to 2 mm thick. Furthermore, MIDAS supports the analysis of multiple objects during interaction, representing a significant advancement in non-invasive material sensing technologies. This solution enhances material characterization by enabling the identification of waste materials based on their thermal signatures, thereby extending the functional scope of autonomous environmental monitoring systems.

1.3.2. LIZARD

Building on advanced sensing to tackle *Environmental Complexity* (Section 1.2) and the *Scalability and Deployment Challenges* (Section 2.3) of field monitoring,

this contribution addresses the issue of environmental littering. Littering poses a critical threat to ecosystems and human health, yet current litter monitoring methods are labor-intensive and reactive. The research presents LIZARD, a novel pervasive sensing system designed for autonomous vehicles to detect and monitor plastic pollution. LIZARD employs an innovative sensing pipeline that integrates thermal imaging to identify larger plastic debris through dissipation patterns and optical sensing to pinpoint areas with dense micro- and mesoplastic fragments. This represents the first pervasive sensing approach capable of microplastic detection in real-world environments and integration with autonomous platforms. Rigorous validation through controlled laboratory experiments and field deployments across three diverse locations demonstrates LIZARD's ability to detect plastics of varying sizes with up to 80% accuracy, with performance influenced by plastic size, surface type, and ambient lighting. Furthermore, LIZARD's seamless integration with ground drones facilitates (semi-)autonomous litter surveillance, offering a scalable and effective tool to address global plastic pollution challenges via pervasive sensing and autonomous monitoring.

1.3.3. Micro-clouds

Finally, to manage the data generated by pervasive sensors and address *Energy Efficiency and Power Autonomy, Real-Time Data Processing* (Section 1.2), and the severe *Energy, Connectivity, and Data Processing Challenges* (Section 2.3), we introduce the fog computing framework – Micro-clouds, using COTS devices. This system adopts a architecture, where multiple COTS devices collaborate as an autonomous cluster to deliver local fog computing capabilities. This research significantly advances the field of underwater data science and sustainable computing by developing a novel fog computing framework based on low-cost micro-cloud prototypes, rigorously validated through extensive benchmarks in both controlled laboratory and diverse field aquatic environments. The work addresses critical challenges such as optimizing communication interfaces and implementing failure detection mechanisms via accelerometers to enable resilient, scalable, and data-intensive underwater applications. Furthermore, building on this technological foundation, the study innovatively repurposes discarded smartphones as cost-effective, small-scale data centers, extending device lifecycles and mitigating the environmental impact of electronic waste. The prototype facilitates customizable sensor data collection and real-time analysis, demonstrating practical applicability across a range of contexts. Together, these contributions establish a versatile and eco-conscious computing infrastructure that not only enhances the capacity and complexity of underwater data processing but also promotes sustainability by fostering energy-efficient practices and reducing e-waste in environmental monitoring systems.

1.4. Thesis Structure

The thesis is organized to systematically develop the contributions outlined in previous section. Chapter 2 reviews relevant background and related work. Chapter 3 details the design, implementation, and evaluation of MIDAS, a thermal sensing system for material characterization. Chapter 4 presents LIZARD, a pervasive sensing solution for autonomous detection and monitoring of environmental plastic litter. Chapter 5 describes the Micro-clouds fog computing framework for underwater data science and sustainable computing. Chapter 6 concludes the thesis by summarizing key findings, discussing the research scope and limitations, and outlining directions for future work.

2. BACKGROUND

Environmental sustainability refers to the responsible interaction with the environment to avoid degradation of natural resources and ensure the long-term health of ecosystems. Traditionally the domain of environmental science, ecology, and policy, environmental sustainability has increasingly become a priority for computational research. As environmental challenges intensify due to industrialization, urbanization, and climate change, there is a growing demand for scalable, intelligent, and automated solutions to monitor, predict, and mitigate their impacts [105] [269].

Recent advances in pervasive sensing, embedded systems, machine learning, and autonomous computing solutions have opened new avenues for addressing environmental problems at scale and in real time. Pervasive sensing refers to the deployment of distributed sensor systems capable of collecting high-resolution spatiotemporal environmental data, often in previously inaccessible or unmapped regions [114]. When coupled with AI techniques, such systems allow for real-time analytics, early warning systems, and data-driven policymaking [343].

The rise of smart cities, citizen sensing, and low-cost IoT platforms has made it feasible to monitor phenomena such as air pollution [157], water contamination [101], and noise levels [248] using computational infrastructures. Moreover, autonomous systems such as UAVs, robotic platforms, and edge computing devices are increasingly used for in-situ environmental assessments and biodiversity monitoring [300] [292].

However, integrating these computational tools into environmental sustainability efforts presents a set of challenges. These include issues related to energy efficiency, data heterogeneity, sensor calibration, privacy, and lack of interoperability across systems and regions. Furthermore, the machine learning models used in such contexts often require domain adaptation, generalizability, and interpretability to be effective in varied and dynamic natural environments [136].

In the sections that follow, this chapter surveys four critical areas—water pollution, air pollution, noise pollution, and biodiversity loss—to identify how pervasive sensing and autonomous systems have been leveraged to address these problems. This review establishes the foundation for understanding the environmental context of this thesis and positions the proposed computational contributions within a growing interdisciplinary research landscape.

2.1. Environmental Sustainability: A Multidimensional Concern

Environmental sustainability addresses the responsible management of natural resources to prevent degradation and ensure their availability for future generations. It is a multidimensional issue encompassing water pollution, air pollution, noise pollution, and biodiversity loss as demonstrated in Figure 1. Each domain requires careful observation, mitigation, and technological intervention [138]. Pervasive

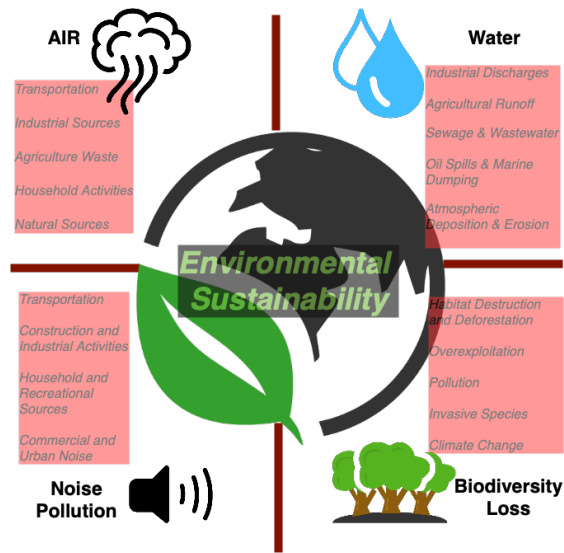


Figure 1: Environmental Sustainability: A Multidimensional Concern

sensing technologies are now being harnessed to provide the continuous, real-time data essential to making environmental sustainability achievable [147] [350], and they form the central focus of this review by linking environmental concerns with advanced computational responses.

2.1.1. Water Pollution

Water pollution stems from agricultural runoff [277], industrial waste, and urban discharge, leading to ecosystem disruption, public health crises, and economic losses [8]. Contaminants such as heavy metals, nitrates, and microplastics deteriorate water quality [256] [24]. Smart water management increasingly depends on sensor networks that provide real-time measurements of parameters like turbidity, pH, and conductivity [82]. These data streams are foundational for tracking contamination sources, enforcing environmental regulations, and ensuring water safety.

2.1.2. Air Pollution

Air pollution is a leading environmental and public health concern, responsible for approximately 7 million premature deaths annually according to the World Health Organization [155]. While industrial activities and general fossil fuel combustion are major contributors, road traffic is consistently the most dominant and immediate source of air degradation in urban environments. The high density of vehicles, combined with stop-and-go driving patterns and poor ventilation caused by urban "street canyons," leads to concentrated, localized spikes in harmful exhaust emissions. These common pollutants include particulate matter (PM_{2.5}), NO_x, and volatile organic compounds [216, 156]. Accurate, real-time air quality

data is vital for forming mitigation strategies, yet traditional monitoring stations are sparse, expensive, and offer limited spatial resolution—particularly in low- and middle-income regions. Effective air quality management hinges on high-resolution monitoring, enabled by both stationary networks and mobile sensors embedded in vehicles and drones.

2.1.3. Noise Pollution

Noise pollution, defined as unwanted or harmful sound that disrupts environmental quality and human well-being, has emerged as a significant urban sustainability challenge. Urbanization, transportation, and industrial machinery are the leading contributors, with road traffic acting as the most pervasive and chronic source of environmental noise in cities. The continuous combination of engine operations, exhaust systems, and tire-pavement friction creates a persistent baseline of acoustic disruption along densely populated urban corridors [19, 201]. Chronic exposure to these high noise levels is linked to stress, sleep disturbances, cardiovascular issues, and cognitive impairments, particularly in vulnerable populations [26]. Although there are environmental regulations in many countries, effective enforcement and mitigation require granular and dynamic data that traditional monitoring systems, typically sparse and expensive, struggle to provide. Smart acoustic sensors deployed in cities and natural reserves facilitate real-time monitoring and inform targeted traffic noise reduction policies [185].

2.1.4. Biodiversity Loss

Biodiversity loss is one of the most pressing environmental crises of our time, driven by habitat destruction, pollution, climate change, and invasive species. Protecting biodiversity is not only an ecological imperative but also a strategy to preserve ecosystem services that humanity depends on. Monitoring biodiversity is crucial for understanding ecosystem health, informing conservation strategies, and ensuring sustainable development [213]. However, traditional biodiversity assessment methods—such as manual species inventories, camera traps, and ecological surveys—are labor-intensive, episodic, and spatially limited [274]. In recent years, pervasive sensing and autonomous systems have emerged as transformative technologies to enable large-scale, continuous, and fine-grained biodiversity monitoring [28].

2.1.5. Ecosystem Threat Level Matrix

To provide a holistic view of pollution threats across ecosystems, the comprehensive ecosystem matrix categorizes various pollution types by impact severity. These severity levels are based on the International Union for Conservation of Nature (IUCN) Red List threat classification system [242], as defined in Table 1.

Threat Level	Impact Description	Decline Rate
Critical	Very rapid declines	> 30% over 10 years or three generations
High	Rapid declines	20% – 30%
Moderate	Slow but significant declines	< 20%
Low	Negligible declines	N/A
Minimal	No significant declines	N/A

Table 1: IUCN Red List Threat Classification System [242]

Building on these classifications, Table 2 synthesizes environmental data to detail ecosystem impacts by pollution type. It highlights the interconnected, synergistic nature of pollution threats that amplify ecosystem degradation. Ultimately, this matrix serves as a foundational reference for integrating computing methods into environmental monitoring, enabling data-driven mitigation strategies.

Pollution Type	Critical Impact	High Impact	Moderate Impact	Low Impact	Primary Ecosystem Effects
Ground-level Ozone (O₃)	Forest ecosystems in heavily polluted regions [10, 127]	Agricultural croplands (62% of EU forests exceeded critical levels in 2022) [10, 127]	Grasslands and natural vegetation [127, 78]	Remote/pristine areas [127]	Reduced growth rates, lower crop yields, biodiversity loss [127, 78, 77]
Nitrogen Deposition (NOX/NH₃)	Sensitive wetlands and heathlands [127, 78, 77]	73% of EU ecosystems above critical loads for eutrophication [127, 77]	Forest soils and grasslands [127, 77]	Marine ecosystems (indirect) [338]	Eutrophication, acidification, species composition changes [127, 77]
Sulfur Dioxide (SO₂)	Historically impacted forests near industrial sources [127, 78]	Freshwater systems in acid-sensitive regions [127, 78]	Coastal and marine environments [338]	Most terrestrial ecosystems (due to emission reductions) [127, 78]	Acidification, toxic metal release, fish mortality [78]
Heavy Metals	Arctic food webs (bioaccumulation) [207, 237]	Aquatic ecosystems near mining/industrial sites [207, 237, 215]	Urban soil ecosystems [207, 215]	Remote terrestrial areas [207]	Bioaccumulation, biomagnification, toxicity in food chains [207, 237, 215]
Particulate Matter (PM_{2.5})	Urban forest patches [305, 64]	Agricultural areas with high dust exposure [305]	Suburban ecosystems [305]	Rural/remote areas [305]	Reduced photosynthesis, leaf damage, altered plant physiology [305, 64] [215]
Water Pollution	Freshwater systems receiving industrial discharge [293, 207]	Coastal marine ecosystems [215]	River systems with agricultural runoff [293]	Groundwater systems [293]	Aquatic life mortality, habitat disruption, food web collapse [293, 215, 207]
Soil Contamination	Agricultural lands with heavy pesticide use [207, 215, 237]	Industrial brownfield sites [207, 215]	Urban soils [207]	Forest soils in protected areas [207]	Reduced soil fertility, crop contamination, biodiversity loss [207, 215, 237]
Plastic/Microplastic Pollution	Marine ecosystems, particularly gyres [338, 215, 293]	Coastal ecosystems and estuaries [338, 215]	Freshwater systems [215]	Terrestrial ecosystems [215]	Food web contamination, physical harm to wildlife, habitat degradation [338, 215, 293]
Thermal Pollution	Aquatic ecosystems near power plants [207]	River systems with industrial discharge [207]	Lake ecosystems [207]	Unimpacted water bodies [207]	Disrupted aquatic life cycles, reduced oxygen levels, species displacement [207]
Noise Pollution	Urban wildlife populations [106]	Transportation corridors [106]	Semi-urban areas [106]	Rural/remote areas [106]	Behavioral disruption, communication interference, habitat abandonment [106]
Light Pollution	Urban ecosystems [25]	Suburban areas [25]	Rural areas near development [25]	Dark sky preserves [25]	Disrupted circadian rhythms, altered migration patterns, reproductive impacts [25]

Table 2: Ecosystem impacts by pollution type.

2.2. Data Collection and Analytics for Environmental Sustainability

To complement the environmental goals outlined in the previous section, this section explores how advanced data collection and analytics play a critical role in monitoring and mitigating ecological threats. The foundation of modern environmental monitoring relies heavily on the deployment of Cyber-Physical Systems (CPS). By seamlessly integrating computational algorithms with physical sensing infrastructure, CPS enables the real-time monitoring, control, and modeling of complex environmental processes [317].

Within this overarching CPS framework, environmental data is primarily gathered through large-scale Wireless Sensor Networks (WSNs). The evolution of WSNs, particularly those utilizing mobile data collection elements such as autonomous vehicles or drones, has revolutionized environmental monitoring by extending coverage across remote and hazardous ecosystems, optimizing node energy consumption, and mitigating network transmission bottlenecks [72].

Furthermore, the ubiquitous presence of smart devices has given rise to Mobile Crowdsensing, a paradigm that leverages the participatory power of citizens. By utilizing the arrays of sensors embedded in commercial smartphones and wearables, crowdsensing platforms can gather vast amounts of spatiotemporal environmental data—such as urban air quality or noise levels—at an unprecedented scale and at a fraction of the cost of traditional static networks [44].

Together, these technologies transform raw environmental signals into actionable insights. By applying advanced analytics to the data harvested via CPS, WSNs, and crowdsensing platforms, these computing methods reinforce the broader sustainability agenda across the specific ecological domains detailed below.

2.2.1. Water Pollution Monitoring

Recent advances in IoT and embedded sensor networks have enabled the deployment of low-cost, real-time, and distributed water quality monitoring platforms [112]. These systems typically integrate multi-parameter probes capable of measuring physical (e.g., temperature, turbidity), chemical (e.g., pH, dissolved oxygen), and biological (e.g., bacterial content) characteristics of water. Data collected via these sensors can be transmitted wirelessly to cloud or edge infrastructure for aggregation, visualization, and analysis. Open-source technologies like Arduino and Raspberry Pi enable low-cost, flexible deployments in rural and underdeveloped areas.

In one implementation, WSNs have been used to monitor river systems for temporal changes in contaminant levels, enabling rapid detection of pollution events and adaptive management responses [266]. Furthermore, the use of autonomous surface vehicles (ASVs) and UAVs has expanded the scope of water monitoring to inaccessible or hazardous regions, such as floodplains, stormwater runoff zones, or industrial effluent sites [180].

A key computational contribution in this space lies in the application of machine learning algorithms for water quality prediction, anomaly detection, and decision support. For example, Support Vector Machines (SVMs), random forests, and deep learning architectures have been trained on sensor data to predict parameters like biological oxygen demand (BOD), total suspended solids (TSS), or heavy metal concentrations with high accuracy [329]. These models are increasingly deployed on edge devices, reducing latency and minimizing the need for high-bandwidth connectivity—an essential feature for remote field deployments.

Another critical area of research involves data fusion and calibration techniques. Due to the heterogeneous nature of sensors and environmental conditions, data normalization and quality assurance are essential for reliable decision-making. Some studies have proposed federated learning approaches to allow collaborative model training across multiple water monitoring sites without sharing raw data, thus ensuring data privacy while improving model generalizability [211] [295].

Despite these advances, challenges persist. Sensor calibration drift, energy constraints, data loss, and limited interoperability between platforms can undermine system reliability. Moreover, while AI models are powerful, their lack of interpretability in regulatory settings can hinder adoption by environmental agencies.

Nevertheless, the integration of pervasive sensing with autonomous systems and intelligent analytics has already begun to reshape water quality management, enabling not just reactive interventions but predictive and preventative strategies [195, 260]. As climate change and population growth continue to stress global water systems, these computational approaches are poised to become indispensable tools in sustainable water governance.

2.2.2. Air Pollution Monitoring

Traditional air-quality monitoring stations are sparse, expensive, and offer limited spatial resolution—particularly in low and middle-income regions. To address these limitations, sensing platforms and autonomous computing systems have emerged as key technologies in modern air quality monitoring. These systems leverage low-cost sensor nodes, mobile data acquisition platforms, and cloud-edge architectures to deliver scalable and fine-grained environmental data [157] [48]. Modern air pollution sensors can detect a range of parameters—such as PM_{2.5}, carbon monoxide (CO), nitrogen dioxide (NO₂), and volatile organic compounds (VOCs)—while being compact, power-efficient, and suitable for fixed or mobile deployment.

Mobile sensing platforms, including air quality sensors mounted on public vehicles, drones, or bicycles, provide flexible, high-resolution monitoring across urban spaces. For instance, a participatory sensing platform using mobile phones and sensor units attached to bicycles was deployed, enabling fine-scale air pollution mapping across city blocks [116]. Similarly, UAV-based monitoring has been used for vertical profiling of pollutants and identifying pollution sources in industrial

zones [160].

From a computational perspective, machine learning plays a central role in the processing and interpretation of air quality data. Data-driven models are used for spatial interpolation, anomaly detection, source attribution, and short-term forecasting. U-Air, a data fusion model combining meteorological data, traffic patterns, and urban topology with sparse sensor inputs to infer city-wide air quality, was introduced [343]. More recently, deep learning models, including convolutional and recurrent neural networks, have been used to model nonlinear pollution dynamics and temporal dependencies [302].

An additional innovation lies in the use of edge computing to perform local analytics, reducing latency and dependency on cloud infrastructure. This is particularly relevant in bandwidth-constrained environments or for real-time alerting in smart city scenarios [53]. Edge devices can perform on-site filtering, preliminary modeling, or federated learning, maintaining privacy while enabling collaborative air quality prediction [252].

Despite these advances, several challenges persist in ensuring data reliability. Low-cost sensors suffer from signal drift, cross-sensitivity, and calibration issues, necessitating robust correction techniques. Moreover, data heterogeneity and missing values complicate model training and validation. From a deployment standpoint, energy efficiency, network reliability, and interoperability standards remain open research issues.

Nonetheless, the integration of pervasive sensing with intelligent and autonomous computational frameworks has drastically enhanced our ability to monitor and respond to air pollution threats. These technologies not only provide visibility into hyperlocal environmental conditions but also offer predictive capabilities essential for proactive interventions, urban planning, and public health policy.

2.2.3. Noise Pollution Monitoring

Acoustic sensors integrated with AI can distinguish between traffic, industrial, and natural sound sources. These systems classify noise events, map spatial patterns, and provide real-time alerts. Such smart noise mapping supports zoning regulations, urban planning, and health impact assessments [210].

Recent developments in pervasive sensing technologies, crowdsensing platforms, and autonomous monitoring systems have introduced new paradigms for real-time, scalable noise pollution mapping. Low-cost acoustic sensors equipped with microcontrollers and edge computing capabilities can be deployed across urban environments to collect continuous data on sound levels and frequency characteristics [13]. These systems offer a scalable alternative to government-grade monitoring stations, with the added benefit of integration into smart city infrastructures.

One of the earliest efforts in this area was the NoiseTube project [183], which

demonstrated the feasibility of participatory noise mapping through smartphones. Users could record ambient sound levels and upload geotagged data to a central server, enabling the creation of crowdsourced urban noise maps. While this approach empowered citizen science, it also raised concerns about data quality, calibration, and privacy—challenges that persist in modern implementations.

To enhance data quality and automation, researchers have turned to dedicated embedded sensing systems and AI-based noise classification models. For example, WSNs have been employed to detect patterns in urban noise [248], while more recent studies use machine learning to classify noise sources such as traffic, construction, and industrial machinery [341, 56]. These models, often implemented using Convolutional Neural Networks (CNNs), can identify noise types from acoustic features like Mel–frequency cepstral coefficients (MFCCs) and spectrograms.

Edge computing also plays a pivotal role in on-device noise processing, enabling real-time classification and anomaly detection without relying on cloud resources. This is particularly useful in areas with limited connectivity or where latency-sensitive applications—such as real-time alerts near schools or hospitals—are needed [27]. Furthermore, energy-aware design in these systems ensures longer deployment periods and reduces maintenance overhead.

Another area of interest is spatiotemporal modeling of noise exposure. Using geostatistical techniques and AI-based inference models, researchers have created dynamic noise maps that adapt to diurnal and seasonal patterns [311]. These tools support urban planning, allowing policymakers to evaluate the effectiveness of traffic calming measures, zoning regulations, or green buffer zones in reducing noise exposure.

Despite these advances, there remain open challenges. These include developing robust calibration methods for heterogeneous sensors, ensuring long-term durability in outdoor conditions, handling privacy concerns around audio capture, and improving public trust in autonomous monitoring systems. Moreover, integrating noise data with multimodal urban sensing frameworks (e.g., combining air and noise data) is still in its early stages and offers a rich direction for interdisciplinary research.

In summary, the intersection of acoustic sensing, machine learning, and autonomous systems offers promising capabilities for high-resolution noise monitoring. These innovations are not only critical for regulatory compliance but also for promoting healthier, more livable urban environments through data-informed interventions.

2.2.4. Biodiversity Monitoring

Monitoring biodiversity is crucial for understanding ecosystem health, informing conservation strategies, and ensuring sustainable development. One of the most significant developments in this space is the use of acoustic sensing for wildlife monitoring, particularly in forest, wetland, and marine environments. Passive

acoustic monitoring (PAM) systems equipped with microphones or hydrophones are capable of detecting and recording vocalizations from birds, amphibians, bats, whales, and insects. These data streams, when processed with machine learning algorithms, can be used to detect species presence, behavior patterns, and even population dynamics [276] [9].

Advanced bioacoustic classification models, often built on convolutional or recurrent neural networks, are trained on spectrograms derived from recorded audio to identify species or events of interest. For instance, Stowell et al. demonstrated the use of CNNs to identify bird species in noisy, real-world acoustic environments [276]. Similarly, marine research has leveraged underwater acoustic sensors combined with unsupervised learning to detect cetacean activity over extended timeframes [173].

Beyond acoustics, camera-based systems—including motion-activated trail cameras and drones equipped with thermal or RGB cameras—have expanded observational capacity. UAVs are increasingly used for aerial biodiversity surveys, offering advantages in terms of spatial coverage and access to hard-to-reach habitats [151]. These platforms, when integrated with computer vision and object detection algorithms, can automatically identify species or detect changes in vegetation and land use [150].

From a systems perspective, edge computing and sensor fusion are critical for deploying biodiversity monitoring at scale. For example, autonomous sensor nodes deployed in rainforests or marine environments must operate under strict energy and connectivity constraints. Techniques such as adaptive sampling, energy-aware scheduling, and onboard inference (e.g., detecting only target species) allow for sustainable, long-term deployments [226].

Despite these innovations, key challenges remain in data quality, generalization of AI models across ecosystems, and integration of heterogeneous sensing modalities. Furthermore, ethical considerations related to wildlife disturbance, data privacy (especially with camera traps), and the digital divide in the deployment of conservation technology must be carefully addressed.

Nevertheless, the integration of pervasive computing, AI, and autonomous robotics is revolutionizing biodiversity monitoring. These tools offer unprecedented capabilities for real-time, non-invasive, and large-scale environmental assessment—contributing meaningfully to the global efforts for ecological sustainability.

2.3. Challenges in Environmental Sustainability Monitoring

Despite significant advances in sensor technologies, data analytics, and autonomous systems, there remain critical challenges to achieving reliable, large-scale, and sustainable environmental monitoring. These challenges are multifaceted—spanning technical, operational, regulatory, and social dimensions. To provide a structured overview of these interconnected issues, Table 3 synthesizes the key challenge

Challenge Category	Key Issues	Primary Impact
Technical	Sensor drift, environmental noise, lack of standardized hardware/protocols.	Unreliable data collection and difficult integration across platforms.
Scalability	Logistical and financial hurdles across large urban or remote ecosystems.	Restricts the geographic scope and density of sensor deployments.
Data Collection	Inconsistent data quality, gaps, and lack of adaptive sampling mechanisms.	Reduces the effectiveness of predictive models and data analysis.
Energy	Limited power sources in remote areas, weather-dependent energy harvesting.	Constrains sampling frequency, data fidelity, and operational lifespan.
Deployment	Physical degradation (UV, biofouling), complex terrains, need for rapid mobilization.	Increases maintenance costs, deployment risks, and limits response times.
Connectivity	Signal attenuation (foliage, topography), high cost of satellite IoT, heterogeneous routing.	Causes intermittent data transmission and delays real-time monitoring.
Data Processing	Cloud latency, edge/fog software complexity, "black-box" AI interpretability.	Hinders real-time decision-making and trust among domain experts.
Regulatory	Data privacy (GDPR), transboundary jurisdictional conflicts, lack of calibration standards.	Complicates legal compliance and renders data inadmissible for policy enforcement.
Societal	Misaligned institutional priorities, funding disparities, varying technical literacy.	Delays system adoption and complicates community engagement (e.g., citizen science).

Table 3: Summary of key challenges in environmental sustainability monitoring

categories, their primary constraints, and their overall impact on monitoring efforts. The remainder of this section categorizes and discusses each of these challenges in detail as they are encountered in current environmental sensing systems.

2.3.1. Technical Challenges

Several studies have underscored these limitations in sensor performance and integration. For example, Environmental noise and harsh conditions have been shown to degrade sensor accuracy over time [179, 169]. Efforts to standardize hardware interfaces have also been slow to gain global adoption [342].

Environmental sensors deployed in real-world settings often face harsh and variable conditions that can affect performance. Issues such as sensor calibration drift, environmental noise interference, and limited precision under extreme weather can lead to unreliable data collection. In addition, the lack of standardized hardware and communication protocols makes integration between platforms more difficult.

2.3.2. Scalability Challenges

The logistical complexity of scaling sensor networks in urban environments has been discussed [314], and the challenges faced by autonomous systems due to terrain and energy limits have also been highlighted [318].

Scaling environmental monitoring systems across cities, regions, or ecosystems presents logistical and financial hurdles. Deploying, maintaining, and upgrading large numbers of sensing units requires extensive planning and coordination. Geographic barriers, remote locations, and environmental fragility further complicate deployment. Autonomous platforms like UAVs and Autonomous ground vehicles

(AGVs) offer partial solutions but are still limited by battery life, payload capacity, and navigation constraints.

2.3.3. Data Collection Challenges

Federated learning techniques have been explored to mitigate issues arising from incomplete or noisy datasets, demonstrating improvements in distributed water quality monitoring [234, 200].

Environmental sensors generate vast amounts of data, which often contain gaps, noise, and inconsistencies due to power failures, sensor damage, or transmission issues. Ensuring consistent data quality across time and space is difficult. Moreover, many monitoring systems lack mechanisms for adaptive sampling, which would allow more efficient use of bandwidth and energy by collecting only significant or anomalous data.

2.3.4. Energy Challenges

Energy limitations remain a critical bottleneck, particularly in forested or underwater environments, as shown in the work of [308].

Many sensing systems are deployed in remote or hard-to-access environments where power sources are limited or unavailable. Although solar energy and other energy-harvesting methods are increasingly used, their effectiveness is constrained by weather, device size, and efficiency. Energy-aware scheduling and low-power communication protocols are required, but these often come at the expense of sampling frequency or data fidelity. These severe energy constraints are further compounded by the physical realities of placing and maintaining these sensors in the field.

2.3.5. Deployment Challenges

Environmental monitoring systems must be robust enough to survive variable climates, vandalism, or wildlife interference [35]. Beyond extreme weather, the physical degradation of sensor housings due to UV exposure, biofouling in aquatic environments, and corrosive coastal air significantly limits deployment lifespans [217, 65]. The logistical effort of deploying and retrieving devices in remote, underwater, or mountainous environments can be substantial, often requiring specialized equipment, vehicles, and trained personnel. Furthermore, access to complex topological terrains, such as dense forest canopies or deep marine trenches, adds both risk and cost to the deployment phase [76]. Additionally, real-time deployment in response to sudden events—like wildfires or floods—requires rapid mobilization capabilities and pre-established coordination among stakeholders, which is frequently hindered by a lack of accessible deployment routes or safe zones [268]. Even when physical deployment is successful, retrieving the gathered information continuously and reliably presents its own set of hurdles.

2.3.6. Connectivity Challenges

Reliable connectivity is essential for real-time data transmission and remote management of sensing systems. However, many target areas for environmental monitoring—such as rural communities, dense forest reserves, and open oceans—lack adequate cellular or broadband network infrastructure. Long-range, low-power protocols like LoRaWAN [117], Sigfox, and NB-IoT help bridge the gap for low-bandwidth applications, but they frequently struggle with signal attenuation caused by dense foliage, complex topography, and adverse weather conditions [67, 225]. While satellite communication systems provide global coverage for highly remote nodes, they remain prohibitively expensive and energy-intensive for continuous, high-frequency data streams [230]. Furthermore, integrating these heterogeneous communication protocols into a seamless, reliable network architecture presents significant routing and synchronization challenges, particularly when node connections are intermittent to conserve power [29]. Once the data successfully navigates these constrained communication channels, it must be ingested and analyzed, introducing significant computational hurdles.

2.3.7. Data Processing and Interpretation Challenges

The difficulties in interpreting black-box machine learning models in ecological applications, as well as the challenges of generalizing deep learning models across varied environments, have been outlined and emphasized in recent studies [143, 11].

The processing of high-volume, heterogeneous data from diverse sensors poses serious computational challenges. Traditional centralized cloud computing approaches often struggle with latency and scalability. While edge and fog computing offer localized processing, these approaches introduce complexity in software architecture, model deployment, and device coordination. Moreover, AI models trained on specific datasets may not generalize well across diverse ecosystems, and their black-box nature often hinders interpretability and trust among domain experts and policymakers. Beyond the technical difficulties of processing and interpreting this data, managing and storing it introduces complex legal considerations.

2.3.8. Regulatory and Compliance Challenges

Environmental monitoring systems must comply with a growing and often fragmented set of regulatory requirements [35]. These include stringent data privacy laws, such as the General Data Protection Regulation (GDPR), which complicate the deployment of camera or acoustic sensors in areas where human subjects might be inadvertently recorded or surveilled [203]. Additionally, cross-border environmental monitoring—such as tracking water quality in rivers that span multiple countries—introduces jurisdictional conflicts and requires harmonized legal frameworks that are often entirely absent [50]. In many regions, there is also a distinct lack of standardized certification for low-cost sensor calibration; consequently, the

collected data is frequently deemed legally inadmissible for enforcing environmental policies or industry compliance [189]. Finally, balancing open-data initiatives with the need to protect sensitive ecological information, such as the exact locations of endangered species or valuable natural resources, poses a significant ethical and regulatory dilemma for researchers and policymakers alike [38]. These legal and regulatory bottlenecks are often reflective of broader societal friction and a lack of unified institutional goals.

2.3.9. Societal and Institutional Barriers

Adoption of advanced environmental monitoring systems often requires alignment between governments, private entities, researchers, and local communities [108]. Differences in priorities, funding mechanisms, and technical literacy can delay or obstruct deployment. Citizen science initiatives have attempted to bridge these gaps, but ensuring data reliability and community engagement remains an ongoing challenge.

Together, these challenges highlight the need for a multipronged approach involving robust hardware design, adaptive software systems, interdisciplinary collaboration, and supportive policy frameworks. Overcoming them will be key to unlocking the full potential of pervasive sensing and autonomous systems in the service of environmental sustainability.

2.4. Defining Pervasive Sensing, Sensors, and Autonomous Technologies

Building upon the preceding discussion on the limitations of traditional environmental monitoring and the imperative to shift towards a proactive management paradigm, we now arrive at the core technologies that enable this transformation. This section delves into the components of the modern framework that moves us from sporadic, reactive analysis to continuous, automated insight.

Pervasive sensing forms the philosophical and technological foundation for autonomous environmental monitoring. This section introduces the key components of this emerging ecosystem, providing the practical tools to overcome the challenges outlined previously.

2.4.1. Pervasive Sensing: The Foundational Philosophy

Pervasive sensing is the ubiquitous deployment of interconnected sensors that continuously monitor environmental conditions. These networks form a cyber-physical interface between ecosystems and digital management tools, facilitating seamless feedback and control loops.

At the heart of this new approach is pervasive sensing (also called ubiquitous sensing). This concept directly addresses the shortcomings of sparse, manual data collection by envisioning a world where low-cost, interconnected sensors are

seamlessly integrated into the physical environment. The goal is to create a "digital skin for planet earth" [236], providing a fine-grained, live feed of environmental data.

Ubiquitous deployment is the solution to the problem of insufficient data coverage. Instead of select sampling points, sensors are distributed densely across forests, watersheds, or urban centers to achieve high-resolution spatial and temporal monitoring [98]. Novel approaches even propose using existing infrastructure, such as public transport or micromobility vehicles, as mobile sensing platforms to pervasively collect data in cities [90] [6].

Cyber-Physical Interface describes the critical link this creates. These systems merge the physical world (ecosystems) with computational algorithms (the cyber world). The sensors act as the bridge, translating physical phenomena (e.g., temperature) into digital data.

Seamless feedback and control loops represents the shift from a passive, data-gathering role to an active management one. Unlike traditional methods where data analysis can lag for weeks, this system enables real-time, automated action. For instance, soil moisture data can automatically trigger a targeted irrigation system. This moves beyond monitoring toward mitigation, connecting sensor data directly with civic engagement and remedial action [75].

Embedded networked sensing provides the foundation for using such networks in large-scale environmental monitoring [83]. More recent research explores how low-cost sensing improves awareness of pollutants and environmental sustainability challenges by increasing the coverage and resolution of available information [350].

2.4.2. The Senses of the System: Types of Sensors

The pervasive sensing network derives its power from a diverse array of specialized sensors, each acting as a digital sense organ. The history of remote sensing has seen the development of numerous sensors for mapping and data acquisition [191]. These are the tools that translate the physical world into actionable data:

Thermal Sensors: These detect infrared radiation to create heat maps. They are instrumental in overcoming the limitations of visual observation for early forest fire detection by identifying hotspots before a fire is visible. This technology enables a preemptive response that was previously impossible [193]. Combined with optical sensing, thermal data can also be used in novel ways, such as identifying plastic litter based on thermal dissipation patterns.

Light Sensors: By measuring solar radiation, these sensors provide the granular data needed for optimal renewable energy planning. In ecology, they quantify the impact of artificial light pollution, a modern environmental stressor that traditional observation struggles to measure effectively [99].

Gas Sensors: To move beyond infrequent, localized air quality checks, networks of gas sensors can provide a real-time city-wide map of pollutants. Recent advances include multispectral nondispersive infrared (NDIR) sensors that can

simultaneously detect methane (CH₄) and compensate for the presence of water vapor, significantly improving accuracy for greenhouse gas monitoring [100]. This continuous data stream empowers regulators to enforce air quality standards dynamically (Snyder et al., 2013).

Acoustic Sensors: In the field of bioacoustics, these sensors allow for non-invasive, continuous wildlife monitoring. By capturing and analyzing animal calls, they can track biodiversity and migration patterns on a scale and duration that human field researchers could never achieve [280].

Biosensors: These highly specific sensors use biological components to detect chemical signatures. They represent a significant leap forward for water quality monitoring, providing immediate alerts for pollutants, toxins, or pathogens [347]. For example, the SEAGULL project uses a light-based sensing solution to classify underwater plastic debris in-situ, removing the need for lab analysis [92] [95].

2.4.3. Scalable Deployment and Intelligent Processing: Autonomous Technologies

To overcome the logistical challenges and physical dangers associated with manual data collection in remote or hazardous areas, this framework relies on autonomous platforms equipped with the sensors described above. Comprehensive reviews highlight the cooperation between terrestrial WSNs, crowdsensing, and UAVs as the future of large-scale environmental monitoring [86].

Mobile Platforms (Drones, Autonomous underwater vehicles (AGVs), Robots): These platforms make monitoring scalable and dynamic. Unlike fixed stations, they can be deployed on demand to a specific event or programmed to cover vast, hard-to-access zones. For instance, drones with hyperspectral sensors are being developed to detect aquatic microplastic pollution over wide areas, a task impossible with fixed sensors [49]. Similarly, the integration of thermal sensors and computer vision onto ground drones for autonomous plastic litter monitoring has been demonstrated [186, 333].

Fog and Edge Computing: A key innovation that makes this high-volume, real-time data stream manageable is the integration of fog and edge computing. This represents a departure from traditional, centralized data processing models. *Edge Computing* places processing power directly on or near the autonomous device. A drone analyzing hyperspectral data can perform initial classification onboard, sending back only relevant threat data instead of terabytes of raw information [122]. *Fog Computing* acts as an intermediate layer, processing data from a group of sensors locally before sending a summary to the cloud [123]. This decentralized approach is crucial for reducing latency, saving bandwidth, and ensuring a truly responsive system, especially in areas with poor connectivity [262].

Layer	Primary Roles	Enhanced Details
<i>Sensing Layer</i>	Physical acquisition of environmental variables (e.g., temperature, pollutants, water quality, noise).	Networks encompass fixed, mobile, and wearable sensors; employ time sync, self-diagnosis, modular design for durability.
<i>Edge Computing Layer</i>	On-device data validation, preprocessing, event detection, and rapid autonomous response.	Performs anomaly detection, thresholding, compression and privacy filtering prior to network transmission.
<i>Network/Communication Layer</i>	Bidirectional, secure, and adaptive data transfer across the entire topology.	Multi-modal (LoRaWAN, LTE, Wi-Fi, mesh, satellite), robust against failures, uses encryption and self-healing protocols.
<i>Fog Computing Layer</i>	Intermediate data aggregation, local network processing, and localized load balancing.	Clusters coordinate multiple edge sensors, manage short-term data buffering, and execute complex localized decision-making.
<i>Cloud Computing Layer</i>	Centralized integration, scalable analytics, historic data storage, and cross-system orchestration.	Advanced AI/ML for pattern recognition, cross-domain environmental fusion, and open APIs for third-party integration.
<i>Application Layer</i>	User-facing decision support, visualization dashboards, alerts, and system control interfaces.	Custom dashboards per user group, real-time mobile/app notifications, AI-driven recommendations and collaborative tools.

Table 4: Computing Layers in Sensor-Based Environmental Monitoring

2.4.4. Integration of Computing Layers in Sensor-Based Environmental Monitoring

The transformation from conventional to next-generation environmental monitoring hinges on the integration of multiple computing layers within a unified ecosystem. This layered architecture is fundamental for collecting, processing, transmitting, analyzing, and operationalizing environmental data, supporting high-resolution, real-time, and actionable insights required for proactive management.

The computing layers architecture interconnects physical sensors, localized on-device computation, intermediate local networks, centralized cloud analytics, and end-user interfaces. By distributing intelligence from field-deployed sensor nodes all the way to the cloud, this approach maximizes data quality, responsiveness, scalability, and decision support [126]. Each layer plays a distinct yet interdependent role, allowing the entire system to function smoothly even as data volumes and heterogeneity grow, as shown in Table 4. . To resolve ambiguities between localized and localized-network computing, the architecture explicitly distinguishes between edge and fog paradigms. The detailed layered description, aligned with Figure 2, is outlined below:

Sensing Layer: Ensures robust, continuous acquisition of high-granularity environmental data. Modularity allows the addition or replacement of sensors. Time synchronization enables precise correlation across sites. Health-status self-reports aid maintenance and ensure data reliability under harsh environmental conditions.

Edge Computing Layer: Resides directly on or immediately adjacent to the sensing hardware. Many modern environmental sensing devices now possess sufficient microprocessing capabilities to act as edge devices themselves. This layer reduces latency by performing immediate, on-device data validation (e.g., removing outliers, checking integrity), compressing data, and applying privacy filters before any network transmission occurs. Near-sensor processing enables the rapid detection of acute emergencies (such as a sudden pollution spike) and can

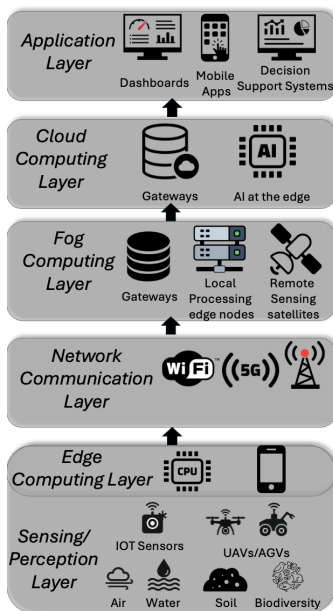


Figure 2: Computing layers in sensor-based environmental monitoring.

autonomously trigger immediate local interventions.

Network Communication Layer: Connects the distributed edge devices to local gateways and the broader internet. It employs hybrid communication protocols (combining LoRaWAN, mesh networking, cellular, and satellite) to ensure reliable, redundant, and secure transmissions. This is crucial for resilience in both urban and remote deployments. Self-healing protocols reroute data when nodes fail, while encryption ensures data privacy.

Fog Computing Layer: Operates at the local or regional network edge (e.g., smart gateways, local routers, or micro-clouds), serving as an intermediate bridge between the device-level edge and the centralized cloud. While edge computing is inherently device-specific, fog computing aggregates, filters, and processes data from multiple edge nodes within a specific geographic area. It manages localized load balancing, short-term data buffering during wider network outages, and complex localized decision-making that requires context from several nearby sensors.

Cloud Computing Layer: Provides the compute resources to archive vast multi-modal datasets. It runs advanced analytics (including deep learning for forecasting and anomaly detection), and deliver insights through application programming interface (API)s. Centralization ensures cross-domain system integration (air, water, biodiversity) and historical comparison, while also enabling rapid scaling.

Application Layer: Offers tailored dashboards, real-time notifications, and advanced AI-driven recommendations. End-users range from policy makers (who need high-level summaries and trend analysis), to scientists (requiring granular

data and query tools), and the public (who receive alerts or participate in citizen science efforts). Application interfaces can also send commands back down the stack for remote system configuration or event-driven responses.

This multilayered computing integration is the digital backbone of modern environmental intelligence. It ensures that data flows seamlessly from the physical environment through advanced distributed analytics to actionable information – empowering rapid detection, comprehensive insight, and robust decision-making for proactive stewardship.

2.4.5. The Path Forward: An Integrated and Responsive Framework

Together, these above technologies and sensors provide a modular, responsive, and decentralized framework for continuous environmental insight. This framework is not merely a technological upgrade; it is the practical implementation of the proactive and predictive paradigm called for in the preceding sections. By overcoming the core limitations of traditional methods, this integrated system of pervasive sensors and autonomous technologies provides the continuous, high-resolution insight necessary for effective environmental sustainability as shown in Figure 3. The figure illustrates how pervasive sensing and autonomous solutions operate collaboratively across terrestrial and aquatic ecosystems. Underwater micro-clouds monitor coral reefs, fish biodiversity, and oil pipelines in real time, while AGVs equipped with thermal and light sensors detect litter and pollution. Drones and IoT-enabled waste management systems enhance monitoring above water, ensuring early detection, rapid response, and predictive management of environmental challenges such as plastic pollution, reef degradation, and industrial spillage.

2.5. Opportunities and Challenges of Sensor-Based Environmental Monitoring

The pervasive sensor systems detailed in the previous section hold enormous potential to revolutionize environmental monitoring. By shifting from a reactive to a proactive model, this technological framework offers transformative benefits. However, realizing this potential requires a clear-eyed acknowledgment of the practical challenges that must be addressed for widespread, reliable deployment.

2.5.1. Opportunities and Benefits

Affordability and Democratization: Low-cost sensors democratize access to data collection, especially in low-income and rural areas. The proliferation of micro-electro-mechanical systems (MEMS) and advances in manufacturing have drastically reduced the cost of environmental sensors. This affordability democratizes data collection, moving it from the exclusive domain of well-funded government agencies and research institutions into the hands of community groups, small municipalities, farmers, and even individuals. This is particularly impactful

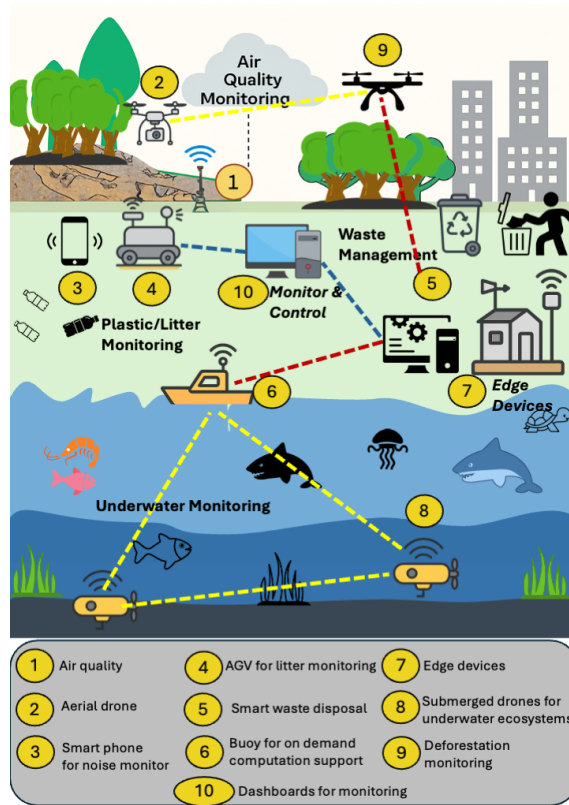


Figure 3: Integration of pervasive sensing and autonomous solutions for environmental sustainability – highlighting underwater monitoring for reef, biodiversity monitoring; AGVs for pollution and litter detection, and aerial/terrestrial systems for waste management and deforestation monitoring.

in low-income and rural regions that have historically been *data deserts*. For example, community-led initiatives can deploy networks of affordable air quality sensors to gather hyperlocal evidence of pollution, empowering them to advocate for policy change with data-driven arguments. The use of low-cost sensors for air quality monitoring has been reviewed, highlighting their role in citizen science and supplementing official monitoring networks, especially in developing countries [142].

Real-Time Insight and Proactive Response: Continuous data streams enable early warning systems and proactive responses. Unlike traditional monitoring, which often involves manual sample collection followed by lab analysis with significant time delays, pervasive sensor networks provide a continuous, high-frequency stream of data. This enables the creation of sophisticated early warning systems. For instance, a network of acoustic and seismic sensors near a volcano can detect subtle changes in gas emissions or ground tremors that precede an eruption, providing critical hours or days for evacuation [190]. This moves environmental

management from a post-event forensic analysis to a proactive, preventative posture, allowing for interventions before a situation escalates into a full-blown disaster. The concept of real-time monitoring for disaster mitigation is well-established. The development of real-time streamflow and water-quality monitoring networks for early flood warnings and contaminant detection has been detailed, showcasing the shift from reactive to proactive water resource management [205].

Scalability and Modularity: Modular sensor networks can be expanded incrementally based on geographic or budgetary constraints. Pervasive sensing frameworks are inherently modular. A monitoring project does not require a massive, one-time investment to cover an entire region. Instead, organizations can deploy a small pilot network in a critical area and then scale it incrementally. For example, a conservation group could begin by monitoring a single river for agricultural runoff and, as funding or needs grow, expand the network to cover the entire watershed, adding different sensor types as required. This architectural flexibility allows for adaptive deployment that can evolve with changing environmental priorities and budgetary realities. The architectural principles of WSNs inherently support scalability. A foundational survey on WSNs discusses how their multi-hop communication protocols and distributed nature allow for the flexible and scalable addition of new sensor nodes [307].

2.5.2. Limitations and Challenges

However, to deploy these systems effectively, their current limitations must be acknowledged and managed.

Sensor Accuracy and Calibration: Low-cost sensors may suffer from calibration issues. There is often a trade-off between cost and accuracy. While low-cost sensors are excellent for indicating trends and identifying hotspots, they can suffer from calibration drift, where their accuracy degrades over time due to aging and environmental exposure. They can also exhibit cross-sensitivity; for instance, a gas sensor's reading being affected by humidity and thermal sensors need frequent recalibration due to drift issues [298]. Without regular co-location and calibration against more expensive, *reference-grade* instruments, the data from low-cost sensors can be misleading, potentially leading to false alarms or missed events and eroding public trust in the technology. A comprehensive analysis of the challenges in deploying low-cost air pollution sensors emphasizes the critical need for robust calibration protocols, data quality control, and field performance evaluation to ensure data reliability [202].

Durability and Maintenance in Harsh Environments: Harsh environments can degrade sensor performance over time. Environmental sensors are, by definition, deployed in the environment, where they are exposed to harsh conditions. Extreme temperatures, high humidity, corrosive salt spray, dust, and physical impacts can cause rapid degradation. In aquatic environments, biofouling—the accumulation of microorganisms, algae, and plants on the sensor surface—can

quickly render optical and chemical sensors useless. Furthermore, maintaining a reliable power supply for hundreds or thousands of remote nodes remains a significant logistical challenge, requiring robust batteries or effective energy harvesting solutions (e.g., solar). The challenge of biofouling is a major focus in marine sensing. The state-of-the-art in anti-fouling strategies for marine sensors has been extensively reviewed, highlighting the critical impact of durability on the long-term operational cost and data quality of ocean observation systems [66].

Data Interoperability and Integration: Integrating data across different platforms and sensors is still challenging. A key vision of pervasive sensing is to create a holistic environmental picture by fusing data from many different sources. However, the current IoT landscape is highly fragmented. Different sensor manufacturers use proprietary communication protocols, data formats, and APIs. This lack of standardization creates *data silos*, making it extremely difficult to integrate data from a water quality network made by Company A with an air quality network from Company B. Overcoming this requires significant effort in developing data middleware or adhering to emerging interoperability standards. The Open Geospatial Consortium (OGC) has been actively working to solve this problem. Their *SensorThings* API standard is designed to provide an open, unified way to interconnect IoT devices, data, and applications over the web. The OGC SensorThings API has been introduced as an interoperable, geospatially-aware framework for the Internet of Things, directly addressing this challenge [69].

2.6. Case Studies: Real-World Applications and Impact Assessment

While the specific technical contributions of this thesis focus on material identification, autonomous plastic monitoring, and providing COTS computing infrastructure, it is crucial to first examine the broader landscape of environmental monitoring. The following case studies—spanning indoor air quality, agricultural water management, marine pollution, and acoustic wildlife tracking – are presented to illustrate the foundational role and proven impact of sensing across diverse ecosystems. More importantly, these varied domains expose universal technical bottlenecks: the reliance on static sensor networks, the struggle with heavy data transmission, and the need for sophisticated, localized edge computing. By analyzing these cross-domain applications, this section establishes the universal baseline challenges that necessitate the advanced, autonomous, and fog-enabled solutions (MIDAS, LIZARD, and Micro-clouds) developed in subsequent chapters.

2.6.1. Air Quality Monitoring: Community-Based Indoor Air Quality System

Case Study: IoT-Based Indoor Air Quality Monitoring for Behavioral Change

Proposed System: An IoT-based indoor air quality monitoring system designed to raise awareness and drive behavioral change in households regarding indoor

pollution sources. The system integrated multiple sensors to monitor PM2.5, CO₂, VOCs, and other indoor pollutants [161].

Technical Implementation: The system deployed wireless sensor networks throughout residential spaces, capturing real-time data on various indoor air quality parameters. Data was transmitted via cellular networks and processed using the COM-B (Capability, Opportunity, Motivation-Behavior) model to translate sensor readings into actionable behavioral insights for residents.

Results: The system successfully identified specific indoor activities that caused pollution spikes and provided personalized recommendations to households. Participants demonstrated measurable improvements in indoor air quality through modified behaviors around cooking, cleaning, and ventilation practices.

Sustainability Impact: This approach addresses the often-overlooked challenge of indoor air pollution, which affects vulnerable populations including children and elderly. By empowering households with real-time data and actionable insights, the system promotes healthier living environments. However, its reliance on static, fixed-location sensors highlights the spatial limitations of traditional monitoring, underscoring the need for the mobile, autonomous sensing paradigms explored later in this thesis.

2.6.2. Water Quality Monitoring: Precision Agriculture Water Management

Case Study: Smart Irrigation Systems with Multi-Parameter Water Quality Sensors

Proposed System: Deployment of comprehensive water quality monitoring networks in agricultural settings, integrating pH, dissolved oxygen, turbidity, and nutrient sensors to optimize irrigation water management and protect soil health [115].

Technical Implementation: The system utilized wireless sensor networks deployed across irrigation channels and groundwater sources. Sensors continuously monitored water quality parameters affecting crop health and soil sustainability. Edge computing nodes processed data locally to trigger automated irrigation adjustments based on water quality thresholds. *Results:* The monitoring system enabled precise water quality management, reducing fertilizer runoff by 30% while maintaining crop yields. Real-time detection of contamination events prevented soil degradation and protected groundwater resources.

Sustainability Impact: This approach directly supports sustainable agriculture by preventing soil acidification, reducing chemical runoff into waterways, and optimizing water resource utilization. The system contributes to long-term soil health preservation and sustainable farming practices. However, the system's reliance on edge computing to trigger automated local responses demonstrates the vital necessity of localized data processing—a concept that must be radically adapted for extreme environments, as addressed by the Micro-clouds.

2.6.3. Environmental Plastic Monitoring: From Traditional Assessment to Autonomous Marine Systems

Case Study: Evolution from Traditional Protocol- Based Monitoring to Autonomous Marine Plastic Debris Monitoring Systems

Proposed System: The integration of traditional monitoring protocols [246] with modern technologies has aided in the development of large-scale autonomous marine pollution monitoring systems that leverage IoT technologies, AGVs, and edge computing to create scalable, continuous marine pollution assessment networks [91].

Technical Implementation: The autonomous monitoring framework integrates multiple technologies – Autonomous Platforms such as AGVs and ASVs equipped with advanced sensors, IoT Integration involving connected sensor networks enabling real-time data transmission and processing, edge Computing for onboard processing capabilities for immediate data analysis and decision-making, scalable Architecture due to modular system design and multi-modal sensing involving integration of optical, acoustic, and chemical sensors for comprehensive pollution detection.

Results: The autonomous monitoring systems demonstrated the capability to provide continuous, large-scale surveillance of marine environments with minimal human intervention. These systems achieved broader spatial coverage, higher temporal resolution, and more cost-effective monitoring compared to traditional vessel-based approaches.

Sustainability Impact: The combination of traditional monitoring protocols with autonomous technologies creates a comprehensive marine pollution monitoring ecosystem that supports both scientific research and operational environmental management. This integrated approach enables proactive marine conservation strategies, supports international environmental agreements, and provides the continuous data streams necessary for adaptive ocean management. This specific evolution toward multimodal sensing and autonomous platforms forms the direct technological baseline for LIZARD and MIDAS systems proposed in this thesis.

2.6.4. Wildlife Monitoring: AI-Enhanced Acoustic Biodiversity Assessment Networks

Case Study: Systematic Review of Artificial Intelligence in Acoustic Wildlife Monitoring

Proposed System: A comprehensive analysis of artificial intelligence applications in acoustic wildlife monitoring, examining the evolution and effectiveness of AI tools and techniques for non-invasive biodiversity assessment across diverse ecosystems [259].

Technical Implementation: The systematic review analyzed 54 research works published between 2015 and March 2022, focusing on AI applications in wildlife acoustic monitoring. The methodology examined algorithm performance through

comparative analysis of various AI techniques including CNNs, machine learning classifiers, and deep learning approaches. The research assessed species coverage across different taxonomic groups, with particular focus on avian and mammalian species. Data processing evaluation included spectrogram-based analysis, pattern recognition, and automated classification systems, while deployment strategies analyzed sensor network configurations and data collection methodologies.

Results: The review revealed significant trends in AI-enhanced acoustic monitoring, showing exponential growth in AI utilization for wildlife acoustic monitoring over the study period from 2015 to 2022. The species focus demonstrated that birds represented the largest monitoring category with 26 studies, followed by mammals with 12 studies, demonstrating the effectiveness of acoustic monitoring for vocal species. Algorithm superiority analysis showed that Convolutional Neural Networks emerged as the most commonly used and effective AI algorithm, demonstrating superior accuracy compared to traditional categorization methods. Performance enhancement results indicated that AI techniques showed marked improvement over manual and semi-automated detection methods, enabling large-scale, continuous monitoring capabilities.

Sustainability Impact: AI-enhanced acoustic monitoring revolutionizes wildlife research by enabling non-invasive, continuous monitoring over vast areas with 24/7 coverage, supporting comprehensive biodiversity assessment, population tracking, ecosystem health evaluation, and anti-poaching efforts. The technology facilitates real-time detection of environmental changes and illegal activities while enabling automated analysis of massive acoustic datasets that would be impossible to process manually. Future development focuses on improving AI models for complex environments, expanding coverage beyond birds and mammals, and integrating with other sensor networks to create comprehensive ecosystem monitoring systems that support both conservation research and operational wildlife management at scale. The heavy computational burden of running these CNNs on continuous acoustic streams further emphasizes the need for resilient, high-capacity fog computing architectures in remote ecosystems.

2.6.5. Cross-Domain Analysis: Integrated Environmental Monitoring Framework

Synergistic Benefits: The integration of diverse sensor networks across air quality, water quality, plastic pollution, and wildlife monitoring creates powerful synergistic effects that exceed individual systems. Multi-parameter insights from combined monitoring provide comprehensive ecosystem health assessments revealing interconnections between atmospheric conditions, aquatic systems, pollution sources, and biodiversity patterns. Integrated sensors enable early warning capabilities that detect cascading environmental impacts, such as how air pollution affects water quality or plastic contamination influences wildlife behavior. Resource optimization through shared infrastructure and data processing reduces deployment costs

across monitoring domains, while comprehensive environmental data supports evidence-based policy development and regulatory enforcement.

Sustainability Outcomes: The integrated monitoring framework delivers transformative sustainability outcomes addressing environmental challenges at multiple scales. Ecosystem protection benefits from holistic monitoring that enables protection of interconnected environmental systems requiring coordinated management approaches. Early detection capabilities support proactive rather than reactive conservation strategies, enabling intervention before environmental degradation reaches critical thresholds. Accessible monitoring technologies empower local communities to participate in environmental stewardship, democratizing access to environmental data. Scientific advancement accelerates through high-resolution, multi-domain data supporting climate research and environmental modeling, providing researchers with unprecedented datasets for understanding complex environmental processes and developing predictive models for sustainable resource management. Ultimately, overcoming the siloed nature of these domains requires the exact convergence of innovative sensing, autonomous deployment, and decentralized computing that this thesis seeks to advance.

2.7. Summary

The growing integration of sensor technology, cloud/edge computing, and autonomous systems opens new frontiers in environmental sustainability. Existing environmental monitoring systems are fundamentally limited by prohibitive costs, sparse spatial coverage, and delayed responsiveness that render them inadequate for addressing contemporary environmental challenges at the scale and speed required. Pervasive sensing offers a transformative shift by enhancing scalability through modular deployment strategies, improving affordability via low-cost sensor democratization, and enabling real-time insight through continuous data streams that support proactive rather than reactive environmental management.

This literature review has comprehensively explored the multifaceted landscape of sensor-based environmental monitoring across five critical dimensions. The *ecological domains under threat* encompass air quality degradation in urban environments, water quality deterioration in agricultural and marine systems, plastic pollution contamination across terrestrial and aquatic ecosystems, and biodiversity loss requiring non-invasive wildlife monitoring solutions. *Computing methods tailored to these threats* include edge and fog computing architectures that enable local data processing, artificial intelligence algorithms for automated pattern recognition and species identification, and cyber-physical systems that create seamless interfaces between natural ecosystems and digital management tools.

The *technological and systemic challenges* reveal a complex landscape where sensor accuracy and calibration issues must be balanced against cost-effectiveness, where environmental durability requirements conflict with deployment scalability, and where data interoperability problems hinder the integration of multi-vendor

sensing networks. *Definitions and taxonomy of relevant technologies* established the foundational framework of pervasive sensing as ubiquitous sensor deployment, categorized sensor types including thermal, light, gas, acoustic, and biosensors, and autonomous platforms encompassing drones, AGVs, and mobile robots equipped with intelligent processing capabilities.

Real-world examples and opportunities demonstrated the practical viability of these technologies through diverse case studies spanning community-based indoor air quality systems that drive behavioral change, precision agriculture water management networks that optimize resource utilization, marine plastic debris monitoring systems that evolved from traditional protocols to autonomous IoT-based surveillance, and AI-enhanced acoustic biodiversity assessment networks that revolutionize wildlife conservation through non-invasive monitoring. The cross-domain analysis revealed synergistic benefits where integrated monitoring systems provide comprehensive ecosystem health assessments, enable early warning capabilities for cascading environmental impacts, optimize resources through shared infrastructure, and support evidence-based policy development.

The convergence of these technological capabilities with urgent environmental needs creates unprecedented opportunities for transformative environmental stewardship. The next phase of this research will propose specific implementation frameworks that address the identified challenges while leveraging the demonstrated synergistic benefits to create scalable, cost-effective solutions for comprehensive environmental monitoring and sustainable resource management at global scales.

3. MIDAS - THERMAL DISSIPATION FOR OBJECT CHARACTERIZATION

This chapter introduces MIDAS, an innovative thermal sensing solution designed to characterize everyday objects by analyzing heat dissipation patterns resulting from human touch. By leveraging contactless thermal imaging, MIDAS overcomes the limitations of traditional Radio-Frequency IDentification (RFID)-based and vision-based approaches, enabling reliable and accurate identification of materials across diverse object types and user interactions. Through rigorous experimental validation, the chapter demonstrates MIDAS's capability to identify materials under varying conditions, including multiple objects and mediation through other materials. By advancing contactless, data-driven material recognition, MIDAS not only enhances our understanding of daily human-object interactions but also contributes to environmental sustainability—supporting applications in more effective waste sorting, resource recovery, and responsible material usage, which are essential for sustainable product life cycles and smart resource management.

3.1. Introduction

Humans on daily basis touch numerous objects, which includes personal belongings, household appliances, food items, and clothing [348]. Acquisition of data based on these interactions holds a significant potential for generating valuable insights into human behavior [149]. These insights include both everyday activity monitoring, along with more intricate applications such as dietary tracking [33] and the identification of household practices [21]. However, the acquisition of such data is challenging due to limitations in the existing methodologies. Contact-based techniques, like RFID, necessitate either the instrumentation of all objects or close proximity between the object and the sensing device for a long period of time [110, 303]. On the other hand, non-contact methods like image-based object recognition have limitations in their ability to distinguish objects [330] and are highly affected by environmental factors such as changes in lighting, camera angle, and image resolution [255]. The aim of the current study was to develop MIDAS, an innovative sensing solution for characterizing everyday objects using thermal dissipation that occurs as a result of the human touch. MIDAS is based on principle that human contact with any object leads to the transfer of heat. After a certain period of time, this transferred heat dissipates as the object seeks to achieve thermal equilibrium with the environment, and the rate of dissipation varies according to the material properties of the object. MIDAS utilizes a COTS thermal camera to capture the changes in heat. Thermal cameras, which function without physical contact are resilient to variations in illumination and environmental context, which address the key limitations associated with current methods. Multiple smartphones are equipped with integrated forward looking infrared (FLIR) thermal cameras,

that are affordable external thermal cameras and can be connected with smart phones. These cameras enable the capture of thermal transfer and monitor the heat dissipation, thereby providing insights about everyday interactions with the material of these objects. We validated MIDAS through comprehensive experiments involving 14 different everyday objects that represent a wide array of common materials used in manufactured products. As part of the experimental evaluation, we also demonstrated the capacity of MIDAS to account for variations in human body temperature by assessing the robustness of thermal dissipation characteristics across 18 different individuals. The results demonstrate that human-emitted radiation can effectively characterize different materials and that this characterization remains consistent despite variations in individual differences and their methods of interaction with objects. MIDAS is capable of accurately identifying materials with up to 83% accuracy, and 16% improvement on a computer vision baseline that uses deep learning. Furthermore, we demonstrated that MIDAS is capable of characterizing materials through other materials and can characterize multiple objects at the same time. Thus, MIDAS represents a novel sensing technology that offers both accuracy and reliability in the characterization of everyday objects, enabling a broad spectrum of innovative applications.

3.2. Related Work

Thermal imaging: The usage of thermal imaging has been studied in different domains and applications with examples ranging from monitoring the manufacturing process of smartphone hardware components [322] to medical analysis [281, 113]. Other examples include facial recognition for bio-metric authentication [32], cognitive analysis [2], gestures [165, 3], and energy modeling of IoT devices [89]. Our work extends thermal imaging to material classification.

Material sensing: Materials have different characteristics different properties that can be exploited to categorize them. The most common material sensing approach is to rely on different parts of the light spectrum and measure either reflection or absorption at different frequencies. Examples range from the use of green light sensing to detect plastic waste [93] to the use of near-infrared sensing to facilitate medicine adherence [148] and the use of hyperspectral imaging for estimating sugar content in drinks [134]. Also, deep learning approaches for detecting different material types from reflection patterns at different wavelengths have been proposed [55]. Our work extends these by using thermal radiation in the infrared spectrum to estimate internal characteristics of materials through heat dissipation.

Sensorless sensing: Wireless signals can also be used to identify properties in materials. Examples include the use of variations in WiFi signal propagation characteristics to identify liquids [71], and the use of surface tension to characterize liquids [336, 310]. These methods generally require either close contact with the material or a transmitter - receiver pair to be placed on opposite sides of the

material. Our work offers a non-contact technique for material characterisation that piggybacks thermal radiation generated from humans.

3.3. Feasibility assessment

We performed two preliminary studies to demonstrate that human heat transfer can be effectively employed to characterize different materials and household objects. Thermal radiation is recorded using a commercially available smartphone (CAT S60), and the measurements are validated by utilizing a thermometer scanner, which serves as the reference instrument. All statistically significant differences are cross-validated using measurements from the reference device. We first provide a detailed description of the testbed setup before presenting the results.

3.3.1. Testbed

We captured thermal fingerprints by utilizing two devices: a handheld thermal imaging scanner (FLIR TG267) and a Caterpillar smartphone (CAT S60) equipped with an integrated FLIR thermal camera. During the course of our experiments, the smartphone was placed on a tripod positioned 30 cm to 35 cm from the object. The camera was manually calibrated after reaching thermal equilibrium with the environment, at a room temperature of 22 °C to 23.5 °C. The thermal video was recorded using the CAT S60, while reference thermal images were taken with the TG267 scanner. The Netatmo weather station was used to measure the room temperature. Dissipation times were automatically estimated from the thermal video and subsequently validated by comparison with ground truth data obtained through manual inspection of the video using a stopwatch.

3.3.2. Plastic thermal fingerprint dissipation

Experimental Design: We initially measured the dissipation time of thermal fingerprints on various plastic materials and correlated the captured fingerprints with the emissivity coefficient (ϵ) of each material. During this experiment, we focused exclusively on plastics to ensure that the emissivity of the materials is known. In the following sections, we validated the applicability of our solution to a wider array of materials. We examined several plastics commonly found in everyday items, including LDPE (Low-Density Polyethylene), HDPE (High-Density Polyethylene), PP (Polypropylene), PS (Polystyrene), and PVC (Polyvinyl Chloride). The material composition of an object is determined by its Resin Identification Code (RIC), with each material having a well-established emissivity coefficient (ϵ), generally ranging from 0.90 - 0.97. The plastic samples tested in these experiments were of identical shape, size, and were manufactured using the similar process. This uniformity ensures that any observed differences between the samples are only attributed to inherent material properties, rather than to external factors such as variations in shape or stiffness. During the experiments, the plastic



Figure 4: Selected waste materials for preliminary experiments: A (Beer Can), B (Ceramic Cup), C (Takeaway Box), D (Plastic bottle), E (Glass Bottle), F (Coffee Cup), G (Plastic Cup), H(Cigarette Butt), I (Glass Jar), J (Milk pack), K (Aerosol Can), L (Rubber glove), M (Metal spoon) and N (Face mask).

samples were placed inside a fridge with a constant temperature of 5 °C, to establish a baseline temperature for comparison. In order to heat the plastic samples we used a constant heat source (lamp bulb 60 W). A distance of 10 cm was maintained between the lamp and samples to avoid any damage caused by burning while ensuring that samples are exposed to sufficient amounts of thermal radiation. We examined various heating durations (1, 2, 3 and 4 minutes) to correspond to initial temperatures and monitored the dissipation of the thermal fingerprint. The most significant changes in dissipation were high during the first few minutes, so any period longer than 4 minutes was excluded. Throughout the experiments, the ambient temperature fluctuated between 22 °C to 24 °C.

Results: The results in Figure 5a show that the dissipation of thermal fingerprints differs between materials. A statistically significant Spearman correlation [254] was observed between the dissipation time and the emissivity coefficient of the materials ($\chi^2(2) = 48.83$, $p < .05$, $W = 0.93$), suggesting that the dissipation characteristics provides insights into the material properties of the object.

3.3.3. Other Thermal fingerprint dissipations

Testbed: In subsequent experiments, we demonstrated that the conclusions drawn in the previous section can be generalized to a broader range of objects and materials by evaluating the dissipation time of thermal fingerprints on various household items. We examined various objects commonly used in households, as illustrated in Figure 4, including: a beer can (A), ceramic cup (B), takeaway box (C), plastic bottle (D), glass bottle (E), coffee cup (F), plastic cup (G), cigarette butt (H), glass jar (I), milk carton (J), aluminum aerosol can (K), rubber glove (L), steel spoon (M), and a face mask (N). The dissipation of the thermal fingerprint is measured using a CAT S60 thermal imaging smartphone in conjunction with a certified CAT TG267 thermometer. A detailed description of the apparatus is




Figure 5: Dissipation time of thermal fingerprints in different plastic materials and waste material using two different devices: (a) Grouped by RIC code, (b) Thermometer scanner FLIR TG267 (baseline), and (c) Smartphone CAT s60.

presented in subsection 3.5. Following the methodology of our previous experiment, we assessed the thermal dissipation time after the objects are held for durations of 1, 2, 3, and 4 minutes. Throughout the experiment, the average body temperature of the individual holding the object ranged from 35 °C to 36 °C, and the ambient temperature varied between 22 °C to 24 °C.

Results: The results presented in Figure 5 are consistent with those observed for plastic objects, indicating that dissipation times vary among different objects and materials. The Friedman test, using the object materials as the experimental condition, revealed statistically significant differences ($\chi^2(2) = 48.83$, $p < .05$, $W = 0.93$) among the materials confirming that thermal radiation can effectively differentiate between various object materials.

3.4. MIDAS pipeline

The results detailed above demonstrated that the dissipation of thermal fingerprints provides critical information that can be utilized to characterize a diverse range of objects and identify their materials. We will now provide a brief overview of the sensing pipeline employed to characterize everyday objects. MIDAS processes a sequence of thermal images captured from the object’s surface and provides an estimate of the most probable material of that object (see Figure 6). Next, we outline the thermal image model utilized for classifying object materials based on the dissipation time of thermal fingerprints.

Preprocessing and Normalization: Since COTS thermal cameras lack cooling mechanisms, they are prone to inaccuracies due to overheating of the camera [184]. Factors impacting measurement quality include misalignment between thermal and RGB images, internal camera recalibration, and low image resolution. To address these issues, we preprocessed the thermal camera data by examining the background of consecutive images and filtering out those with significant differences. We also performed denoising on the images and normalized the thermal values to a consistent scale (0 and 255). This normalization allows us to process and analyze the images in grayscale. Figure 7 illustrates the outcome of

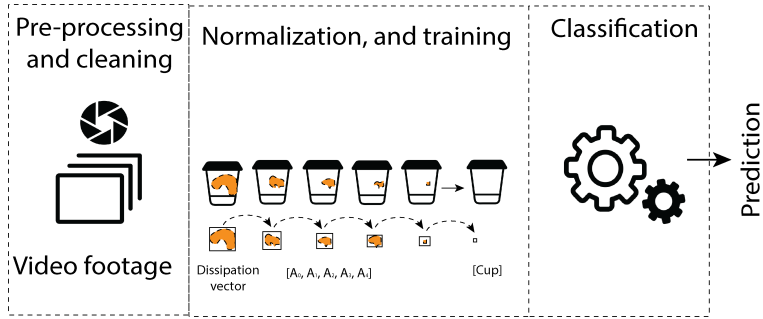




Figure 6: Processing pipeline of material classification based on dissipation time of thermal fingerprints.



(a)



(b)

Figure 7: Dissipation time of thermal fingerprint for two different objects: (a) Cardboard cup and (b) Cigarette butt.

the normalization procedure. The upper section displays a sequence of thermal images over time, while the lower section presents the corresponding images on the normalized scale for two objects: (a) a cardboard cup and (b) a cigarette butt. This normalization process enables the isolation of thermal fingerprints for each object.

Dissipation Rate: We estimated the dissipation rate of the thermal fingerprint from the normalized sequence of images as the function of reduction in the area of the thermal fingerprint, as defined by the following equation:

$$RA = (A_i - A_t) / A_i, \quad (3.1)$$

where RA corresponds to reduction in percentage of area, A_i refers to the initial area, and A_t is the reduced target area [84] (see Figure 7). The decrease in area

between consecutive images is employed to construct vectors that represent the dissipation time of thermal fingerprints for each object. To enable effective training of machine learning classifiers, we considered fixed sized vectors. It is important to note that once the thermal fingerprint has fully dissipated, the target area becomes zero, which effectively results in padding the vector with zeros.

Implementation: The vectors (1D arrays) representing dissipation time over sequence frames are used as feature vectors, while the object type serves as the label class. While state-of-the-art architectures such as Vision Transformers (ViTs) or deep CNNs dominate modern computer vision tasks, they were intentionally excluded from this pipeline for two reasons. First, the input to the classifier is not a high-dimensional image matrix, but rather a lightweight, fixed-size one dimensional vector of numerical dissipation rates. Highly parameterized models like ViTs are computationally excessive for this data structure and prone to severe overfitting. Second, MIDAS is explicitly designed for pervasive environmental monitoring on low-resource, edge-computing devices where energy and memory are heavily constrained. Therefore, classification models were constructed using computationally efficient, standard machine learning techniques: Random Forest (RF), SVMs, and a Multi-Layer Perceptron Classifier (MLPC). To ensure robust model evaluation, the dataset of thermal feature vectors was partitioned into an 80% training and 20% testing split. The training process utilized a stratified 5-fold cross-validation approach to maintain class distribution and prevent selection bias. Hyperparameter optimization was conducted using grid search for each algorithm:

- **Random Forest (RF):** Configured with 100 estimators (trees) using the Gini impurity criterion, with the maximum depth of the trees left unconstrained to allow node expansion until all leaves were pure.
- **Support Vector Machine (SVM):** Employed a Radial Basis Function (RBF) kernel, which is highly effective for non-linear temporal data. The regularization parameter (C) was set to 1.0, and the kernel coefficient (γ) was scaled based on the variance of the feature vectors.
- **Multi-Layer Perceptron (MLPC):** Configured as a lightweight neural network with a single hidden layer of 100 neurons. It utilized the ReLU activation function and the Adam optimizer, with an initial learning rate of 0.001 and early stopping implemented if the validation score did not improve over 10 consecutive epochs.

This configuration ensured that the resulting models remained lightweight enough for edge deployment while maximizing classification accuracy on the transient thermal data.

3.5. Robustness of Thermal Dissipation Fingerprints

The experiments detailed in section 3.3 demonstrated that the dissipation characteristics of thermal fingerprints differ among various materials. Human body




Figure 8: Experimental conditions: (a) Rigid interaction, used in Fixed-hold and (b) Free interaction, used in Natural and Quick holds.

temperature varies between individuals and can even fluctuate within the same individual at different times of the day [206] leading to variations in the initial thermal fingerprints. Robust operation of MIDAS despite these variations is essential for its practical applicability and overall effectiveness. In this section, we outline experiments in which 18 individuals interact with everyday objects. We analyze the resulting thermal fingerprints and their dissipation characteristics to identify and characterize the materials of the objects. The measurement setup is detailed in subsection 3.3.1. The experiments are conducted using an experimental testbed designed to capture video footage of thermal fingerprints for various object materials. To achieve this, we employed a commercially available smartphone thermal camera and a thermometer scanner to establish the reference baseline. We selected three different materials, each chosen for their distinct shapes and sizes, as well as their sufficiently rapid dissipation times to minimize the overall duration of the study.

3.5.1. Experimental setup

Experiment Design: We implemented a 3×3 within-subject design with holding pattern type and object type as independent variables. Each variable has three levels: Fixed-hold (FH), Natural-hold (NH), and Quick-hold (QH) for holding pattern type, and Plastic bottle (BOTTLE), Cardboard cup (CUP), and Cigarette butt (CIGAR) for object type. To control for order effects while maintaining a manageable number of combinations, holding pattern type was fully counterbalanced, and object type was counterbalanced using a Latin Square design. This resulted in nine experimental conditions: (1) BOTTLE-FH, (2) CUP-FH, (3) CIGAR-FH; (4) BOTTLE-NH, (5) CUP-NH, (6) CIGAR-NH; (7) BOTTLE-QH, (8) CUP-QH, and (9) CIGAR-QH. In the *Fixed-hold* condition, objects were grasped and maintained in a specific static position for one minute (see Figure 8a). In the *Natural-hold* condition, objects were held freely for one minute, replicating typical everyday interactions with the object (see (Figure 8b). For example, a participant might hold an empty bottle for one minute while searching for a trash bin. In the *Quick-hold condition*, objects were held freely by participants for a brief interval of 10 seconds.

Participants: We recruited a total of $N = 18$ participants (Males = 9 Females = 9) for the user study. Participants included students, administrative staff, and professionals from various fields and nationalities, all of whom had minimal or no prior knowledge of thermal imaging. The average age of the participants was 28 ± 7.8 years.

Task: Participants were instructed to hold the objects and simulate typical interactions with them. To gather data on natural interactions, participants were also asked to create a context for their actions. For the BOTTLE, participants were asked to mimic drinking from the bottle and then searching for a trash bin to dispose of the empty bottle. For the CUP, participants were instructed to stand while engaging in a brief conversation with an acquaintance or friend. In the CIGAR condition, participants were asked to simulate taking a cigarette from a cigarette box and then holding the cigarette by the filter while requesting a light. The cigarette was not actually lit during the experiment.

Procedure: Before beginning the experiment, each participant was seated in a comfortable chair for 10 minutes to allow their body temperature to adjust to the room's ambient temperature, which ranged from 22°C to 23.5°C throughout the duration of the experiments. During this time, participants were given a brief overview of the study and signed an informed consent form in accordance with local IRB regulations. When participants were ready to initiate the experiment, their body temperature was measured from forehead using a clinically certified contactless optical thermometer (DR CHECK FC500). Participants were presented with nine experimental conditions. For each condition, the object was first placed in an empty refrigerator maintained at 5°C for one minute (Figure 9a). This procedure eliminates residual thermal radiation in the material between experiments and establishes a baseline temperature for the material, ensuring that our results are comparable across participants. Kitchen tongs were utilized to transfer the object from the refrigerator to prevent any heat transfer from the human hands to the object. The object was then placed on a table for one minute to acclimate to the ambient temperature (see Figure 9b). Afterwards, participants performed the corresponding experimental condition. Upon completion, they placed the object on a fixed marker drawn on a table with a black background and surface. The researcher then used the CAT S60 to capture video footage of the dissipation of the object's thermal fingerprint. Simultaneously, a thermometer scanner captured thermal photos to establish a reference baseline (see Figure 9c). The black background was used to ensure clean video footage of the thermal fingerprints, minimizing any thermal interference from surrounding objects. Upon completion of the experiment, we recorded the participant's temperature at the palm and fingers relative to the objects using the thermal imaging scanner. The assessments were conducted in a university room over a span of two weeks, with sessions held between 11:00 AM and 7:00 PM. Due to the variability in human body temperature throughout the day [206], the experiment was scheduled during periods that corresponded with participants




Figure 9: Experimental testbed and protocol steps: (a) Object obtains baseline temperature, (b) Object habituates to ambient temperature, and (c) Participant performs the experiment and puts the object in the marked target to measure its thermal fingerprint.

working hours. The total duration of the experiment for each participant ranged from 40 min to 45 min.

3.5.2. Baselines

In the course of the experiments, we evaluate the recognition performance of MIDAS against two advanced techniques: deep learning-based automated computer vision [313, 243] and optical sensing [265].

Computer Vision: We trained a state-of-the-art CNNs model utilizing the publicly available TrashNet dataset [16]. Specifically, we focused on the plastic materials category, which contains 626 images of plastic objects used to train the deep learning model. Plastics are malleable, making their precise recognition susceptible to variations. The dataset features images of individual plastic items set against a white background. Since these images do not reflect realistic recognition conditions, we supplement the dataset with an additional 767 images sourced from the Japan Agency for Marine Earth Science and Technology (JAMSTEC) Deep-sea Debris Database. We manually annotated the collected images by drawing a rectangular box around the object material. The TrashNet plastic items were labeled as “trash,” while the JAMSTEC plastic items were labeled as “plastic.” Both datasets were augmented with various modifications, including noise, hue adjustments, blur, horizontal flips, and vertical flips applied to each original image. This augmentation increased the dataset to a total of 6985 images for training the model. We developed and trained the PlasticNet model using a Google Colab server equipped with a GPU. The training involved 100k iterations with a batch size of 12, running TensorFlow Lite 1.15. The base training model was *ssd_mobilenetv2_oidv4*, configured with default hyperparameters.

Light Reflectivity: As our second baseline, we assess material reflectivity [265] utilizing a photoresistor connected to the analog input pin of an Arduino MEGA ADK. The photoresistor detects changes in light intensity by measuring variations in its resistance, which are influenced by the reflectivity of the material. For

illumination, we use a red laser diode with a wavelength of 650 nm. The object was positioned 2 cm away from the light source, simulating practical applications of the sensor in transport belts and smart bins [313]. The measurements were taken with the sensor for various materials (as described in Section 3.3) for one minute, from two different random locations on each object.

3.6. Results

We demonstrated that MIDAS is capable of accurately characterizing different object materials based on the measurements obtained from the controlled user evaluation outlined in the previous section. We assess its robustness in relation to variations in human-object interactions and individual differences, evaluate the overall classification performance across various materials, and examined the robustness of the system in detecting materials even through other objects.

3.6.1. Differences in Thermal Fingerprints

To begin with, we have investigated the variations in thermal transfer from the human body and the subsequent dissipation across various objects. Our analysis examines the influence of both the type of object and the method of holding on thermal transfer, addressing each factor independently. The dissipation times for the three objects under various holding conditions are depicted in Figure 10. The data illustrates that the differences in dissipation times between the objects remain consistent across all holding conditions. The Friedman Test, conducted with dissipation time and object type as experimental variables, indicates that these differences are statistically significant for each of the three holding conditions: Fixed-hold ($\chi^2(2) = 20.33$, $p < .05$, $W = 0.56$), Natural-hold ($\chi^2(2) = 30.33$, $p < .05$, $W = 0.84$) and Quick-hold ($\chi^2(2) = 25.04$, $p < .05$, $W = 0.64$). Subsequent pairwise post-hoc analyses, performed using the Wilcoxon test with Bonferroni correction for multiple comparisons, have confirmed that the differences in dissipation times between all object pairs are statistically significant across the three holding conditions. Afterwards we evaluated the impact of holding type on thermal fingerprints. The Friedman test, using the dissipation times of each object under the three experimental conditions, reveals that the holding type significantly affects the thermal fingerprints for the plastic bottle ($\chi^2(2) = 12.44$, $p < .05$, $W=0.34$) and the cardboard cup ($\chi^2(2) = 16.48$, $p < .05$, $W=0.45$), but not for the cigarette butt. Post-hoc analyses indicated that the thermal fingerprints associated with the Quick-hold pattern differ significantly from those observed in the Fixed- and Natural-hold conditions. These results suggest that while thermal fingerprints exhibit considerable variation across different objects, the holding type influences these variations. However, the dissipation times are also dependent on the duration of contact with the object, which is expected since the length of time an object is held affects the amount of heat transferred and, consequently, the rate of dissipation.

We also conducted a separate analysis to determine whether variations in body temperature across different body parts influence the thermal fingerprint. This was done by comparing temperature measurements taken at three different locations: the forehead, the palm of the hand, and the fingertips. The Friedman test, with body part as the experimental condition, revealed significant differences in temperature ($\chi^2(2) = 29.66, p < .05, W = 0.82$). Post-hoc comparisons using the Dunn–Bonferroni method confirmed that these differences were statistically significant ($p < .01$) between the forehead and palm of the hand, as well as between the forehead and fingertips. The average temperatures recorded for the different body parts were as follows: forehead at 36.33 °C, palm of the hand at 30.16 °C, and fingertips at 30.87 °C. These results demonstrate that the fingers and palm of the hand typically induce comparable heat transfer, further validating the robustness of thermal fingerprints regardless of the manner in which individuals interact with objects. However, the observed deviation from forehead temperature implies that the palm and fingertips are more responsive to variations in ambient temperature rather than directly correlating with core body temperature, as observed with forehead measurements. This thermal sensitivity of the fingertips and palm suggests potential utility in extracting data on human activities. For instance, handling a smartphone may result in heat transfer from the device’s battery to the hand, while the temperature of a hot beverage could similarly alter the thermal readings of the hand. Overall, the findings suggest that human touch transfers a sufficient amount of heat, enabling the characterization of objects based on only touch, without the need for specialized technology. However, the results also reveal that dissipation times are influenced by the duration of user interaction with the objects, as well as by other factors to be discussed in the subsequent subsection. This implies that relative differences in dissipation characteristics, rather than precise dissipation times, should be employed when characterizing materials.




Figure 10: Thermal transferred conditions applied over three different objects (plastic bottle, cardboard cup and cigarette butt).

3.6.2. Other Factors Influencing Thermal Dissipation

Until now, we have demonstrated that the characterization of object materials through thermal radiation is dependent on exposure time, but is not influenced by the specific location where the object is touched. In this section, we will explore

additional factors that impact the transfer of thermal radiation from the human body to objects.

Gender and temperature: Body temperature typically influences the thermal radiation transferred to objects. In the subsequent analysis, we examine whether the gender of the participants affects thermal transfer, specifically investigating whether significant differences exist between female and male participants. The effects were analyzed separately for the palm and fingertips. Kruskal–Wallis tests, with gender and body part as experimental conditions, revealed significant differences in thermal transfer exclusively for the fingertips ($\chi^2(2) = 5.08$, $p < .05$). This disparity is most likely due to the differences in contact area size, with males generally possessing larger fingertip dimensions [5]. When considering gender and objects as experimental conditions, Kruskal–Wallis tests reveal significant differences in dissipation times across all three objects: ($\chi^2(2) = 3.94$, $p < .05$), plastic bottle ($\chi^2(2) = 12.17$, $p < .05$) and cardboard cup ($\chi^2(2) = 7.75$, $p < .05$). These findings suggest that the dissipation time of thermal fingerprints is influenced by temperature and that it may be possible to distinguish between female and male individuals who have interacted with the object. While this result does not alter the fact that objects can be characterized by thermal radiation, it is crucial to highlight the potential privacy implications. Indeed, the findings indicate that thermal radiation can reveal additional information about the individuals interacting with the objects. For instance, thermal fingerprints could potentially be used to compare household waste sorting practices between genders.

External temperature of Ambient environment: The surrounding temperature of an object directly influences the dissipation time of the thermal fingerprint. To quantify the impact of this factor, we conducted additional small-scale experiments. Initially, a BOTTLE was held by a human hand at different ambient temperatures of 22 °C to 23.5 °C. Subsequently, the BOTTLE was placed in a colder environment (a refrigerator with a temperature of 5 °C). This was followed by measuring the dissipation time of the thermal fingerprint during the transition from the ambient to the colder environment using both the CAT S60 and a thermometer scanner. Figure 11 presents the results. For comparison purposes, we also included the dissipation time of the thermal fingerprint in the ambient environment as a baseline. We observed that the total dissipation time of the thermal fingerprint is reduced by half when the temperature is transitioned to a colder environment. Despite this reduction, the overall patterns of change remain consistent across different objects, indicating that environmental conditions influence the thermal fingerprints. It is important to note that the magnitude of this change is directly proportional to the temperature differential and affects all objects uniformly. Therefore, accounting for ambient temperature in the analysis of thermal dissipation fingerprints is sufficient to address potential issues arising from temperature variations.

Internal Temperature absorbed from contents: In addition to the ambient temperature of the surrounding environment, objects can also be influenced by

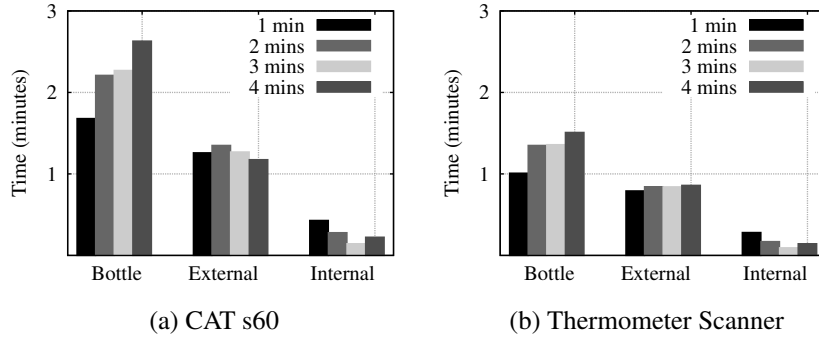


Figure 11: Influence of internal and external temperatures.

thermal radiation resulting from their contents. For instance, a cardboard cup may contain a hot or cold beverage. To explore this further, we filled the BOTTLE with water at a temperature of 21.2 °C to 21.5 °C. Prior to the experiment, the BOTTLE was placed inside a refrigerator (5 °C) to eliminate thermal radiation carryover effects between experiments. We then compared the dissipation times of an empty bottle and a filled bottle at ambient temperature (22 °C to 23.5 °C). Figure 11 presents the results. It was observed that internal radiation does indeed impact the thermal fingerprint of the BOTTLE. This finding is particularly relevant for identifying end-products that have not been fully consumed. In practice, these cases should be modeled as distinct (mixed) objects to ensure that the model can accurately differentiate between the pure material and instances where the object contains residual contents. The thermal fingerprint of the filled BOTTLE, in comparison to that of the empty BOTTLE, dissipates more rapidly due to the greater temperature differential with the environment.

Distance between object and thermal camera: In the experiments conducted so far, the distance between the thermal camera and the objects was consistently maintained at 30 cm to 35 cm, serving as the baseline. We subsequently analyzed the effect of increased distances by introducing three additional measurements: 70 cm (distance-1), 105 cm (distance-2) and 210 cm (distance-3). This analysis focused exclusively on the BOTTLE using a fixed-hold setup for one minute. At a distance of 70 cm distance, the dissipation time did not exhibit significant changes, with an average time of 1.13 minutes for the CAT S60. However, at longer distances, we observed greater variability in the dissipation time with the CAT S60. Specifically, at a distance of 105 cm, the average dissipation time decreased to 0.76 minutes, and at 210 cm, it further reduced to 0.26 minutes. The CAT S60 thermal camera, with a resolution of 80 × 60, appears adequate for distances up to one meter. The thermometer scanner, which has a higher resolution of 160 × 120 but a slower frame rate (6.67 Hz vs. 8 Hz), fails to capture a proper thermal fingerprint at a distance of 210 cm. Therefore, a higher resolution alone is insufficient to extend the operational range of MIDAS; the frame rate must also be taken into consideration.

Temperature Sensitivity: In the next step, we investigated the sensitivity of dissipation times to minor variations in temperature. To achieve this, we extended our analysis of the BOTTLE material under the fixed-hold condition, as previously described. A JANOEL18S incubator with adjustable temperature controls was employed to precisely regulate temperature changes. The object was placed within the incubator at a stable temperature for a sufficient duration to ensure uniform thermal equilibrium across the entire object. The material was exposed to temperatures ranging from 36 °C to 39 °C, corresponding to normal and increased human body temperatures. Upon reaching the desired temperature, the object was transferred to the testbed for recording. Subsequently, the dissipation time of the thermal fingerprint was measured. The ambient temperature in which the object acclimatized ranged from 23 °C to 23.5 °C. The dissipation times recorded at various temperatures were as follows: 3.33 min for 36 °C, 3.73 min for 37 °C, 4.23 min for 38 °C, and 4.34 min for 39 °C. The dissipation time is expected to be a function of the difference in temperatures of the object and the environment, and our results confirm that these slight temperature variations can be robustly detected using a commercially available thermal camera. While this method may not enable precise estimation of human body temperature, we envision that it could be effectively employed to detect irregularities in human temperatures. For example, instead of using thermal cameras to monitor individuals facial temperatures at airports, it is possible to detect abnormal temperature readings by examining the tangible objects that people touch while passing through security checks.

3.6.3. Dissipation Time Classification Performance

We subsequently demonstrated that our approach is capable of supporting the coarse-grained classification of object materials based on the dissipation time of thermal fingerprints, along with other contextual factors. As detailed in Section 3.4, we employed three classification techniques: RF, SVMs, and MLPC. The outcomes of the classification experiments are presented in Table 5. When utilizing only the thermal fingerprint, the maximum classification accuracy for material detection is approximately 83%. The incorporation of hold pattern type data does not result in improved accuracy, indicating that the manner in which individuals grasp the materials does not significantly influence performance. However, when incorporating gender information (i.e., whether the individual is male or female), there is a notable improvement in the accuracy of material detection, increasing to as much as 86%. Similarly, when predicting the gender of the user, the combination of dissipation time and material information yields a high accuracy estimation of approximately 78%.

3.6.4. Comparison Against Other Approaches

Afterwards, we made a comparison of MIDAS with two baseline approaches: computer vision and optical sensing. The computer vision method was evaluated

Table 5: Material classification accuracy (%) in different experimental conditions. Model data \rightarrow Predicted. Classification Method: Random Forest (RF), Support Vector Machine (SVM) and Multi-layer Perceptron (MLPC).

Test	RF	SVM	MLPC	Average
Predicting Material (M)				
(Vector) \rightarrow M	90.9	77.3	81.8	83.3
(Vector, Context) \rightarrow M	90.9	77.3	81.8	83.3
(Vector, Gender) \rightarrow M	90.9	86.4	81.8	86.4
(Vector, Context, Gender) \rightarrow M	86.4	81.8	81.8	83.3
<i>Average</i>	<i>89.8</i>	<i>80.7</i>	<i>81.8</i>	<i>84.1</i>
Predicting Context, Gender				
(Material, Vector) \rightarrow Context	77.3	81.8	72.7	77.3
(Material, Vector) \rightarrow Gender	77.3	77.3	81.8	78.8
<i>Average</i>	<i>77.3</i>	<i>79.6</i>	<i>77.3</i>	<i>78.1</i>

using 31 images and these images depicted real tossed plastic objects [218]. These images contained a total of 33 distinct plastic items, of which the deep learning model accurately identified 23, achieving an accuracy rate of 69.7%. It is important to emphasize that the purpose of this comparison was not to directly contrast our approach with computer vision, as the two methodologies address different problems. Specifically, the vision-based approach does not extract any internal characteristics of the objects, thus its utility in material recognition is limited to mapping specific item types to materials (e.g., drink bottles are typically made of PET). This approach works when the objects possess sufficiently distinctive features but tends to fail in more common scenarios where the objects have undergone changes in shape, such as tossed items that have lost their original form, or where other visual characteristics have been significantly altered. Additionally, computer vision is prone to issues such as occlusions and instances where objects are only partially visible, as highlighted in Figure 12a. In practice, computer vision and thermal dissipation methods can be complementary. For example, autonomous ground vehicles could employ computer vision to detect litter from a distance, then navigate closer to the object and use thermal dissipation techniques to ascertain the material composition of the object. In practice, computer vision and thermal dissipation techniques can complement one another. For instance, autonomous ground vehicles can employ computer vision to identify litter objects from a distance, maneuver closer to the object, and then apply thermal dissipation techniques to determine the material composition of the object.

Figure 12b presents the results for the second baseline, light reflectivity. The minimal variation in reflectivity values suggests that light reflectivity can effectively characterize different materials. However, it is also evident that different sections of the same material can exhibit significant variation, as seen with Cardboard Cup-1 and Cardboard Cup-2. This variation arises because objects are often composed of various materials and colors, which can influence their reflectance properties.

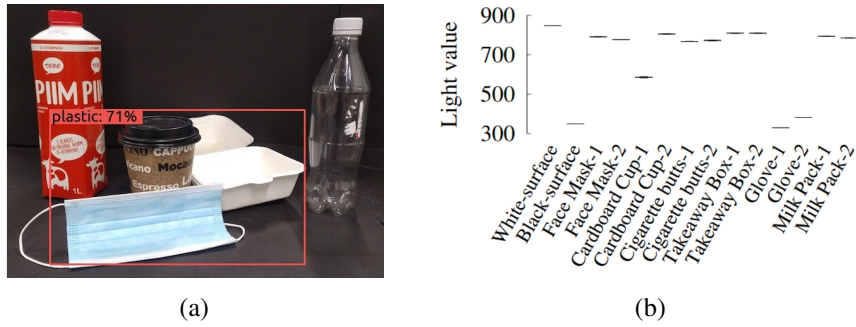


Figure 12: Baselines: (a) Trained model PlasticNet to identify and separate plastics from other object materials, and (b) Light reflectivity values of different materials measured with a photo-resistor.

Additionally, a notable limitation of this approach is the requirement for sensors to make contact with the material in order to classify it accurately.

3.6.5. Effect of thickness on thermal dissipation

The results discussed thus far have explored scenarios in which the camera has a direct line of sight to the objects with which the user interacts. However, in many practical applications, objects may be partially or fully occluded by other items. We demonstrated that thermal dissipation can still be effectively captured even when objects are covered by other materials. It is important to acknowledge that the thermal dissipation is inherently influenced by the properties of the covering materials. Consequently, the objective of these experiments is to demonstrate the potential of our technique for practical applications, rather than to assert its capability to detect materials through the cover.

Experimental Setup: Thermal radiation is absorbed from the outer surface of an object and dissipates towards its interior. To investigate this phenomenon, we conducted an experiment where thermal radiation from the exterior of an object is absorbed and subsequently dissipated inward, which allowed us to model the thermal behavior of the object. Considering that the material properties influence thermal dissipation, as previously demonstrated, we designed specialized samples to analyze the effect of material thickness on thermal dissipation. These samples were derived from common household objects. In this experiment, we selected a plastic container (Plastic), a cardboard box (Cardboard), and a face mask (Mask) as baseline samples for thickness measurements (as shown in Figure 13). We then layered multiple samples of the same object to create specimens with varying thicknesses. It is important to note that the overall penetration of thermal radiation is determined by two primary factors: (i) the thickness of the material and (ii) the characteristics of the material. These factors together define the thermal resistance of a material [182], which governs the amount of heat that passes through and, consequently, the extent of thermal radiation that can be captured on the opposite

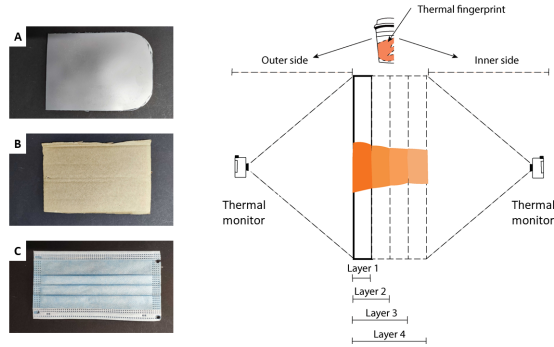


Figure 13: Thickness of object materials analyzed with thermal dissipation fingerprint.

Table 6: Material thickness in each layer for Cardboard, Plastic and Mask respectively.

Material Type	Number of Layers	Thickness (mm)
Cardboard	1 layer	2.801
	2 layers	5.62
	3 layers	9.183
	4 layers	11.93
Plastic	1 layer	0.44
	2 layers	0.91
	3 layers	1.53
	4 layers	1.95
Mask	1 layer	1.65
	2 layers	3.23
	3 layers	4.49
	4 layers	6.13

side of the covering material. In our experiments, we used a maximum of four layers, as we observed that thermal dissipation significantly diminishes beyond this point. The maximum thickness used in the experiments was approximately 2 mm, which corresponds to the typical thickness of many everyday consumer products. For instance, a plastic water bottle generally has a thickness between 1 mm to 2 mm. The thickness of each material layer is detailed in Table 6. We measured the thickness of the objects using a Fujisan FJS025 electronic micrometer screw gauge, which offers an accuracy of 0.001 mm and a measurement range from 0 mm to 25 mm.

Procedure: The thermal dissipation fingerprint was measured from the topmost layer of the material for a duration of one minute following touch. Afterwards, the thermal dissipation was measured from the inner side of the object. This experiment was conducted three times for each sample and thickness layer. Additionally, we

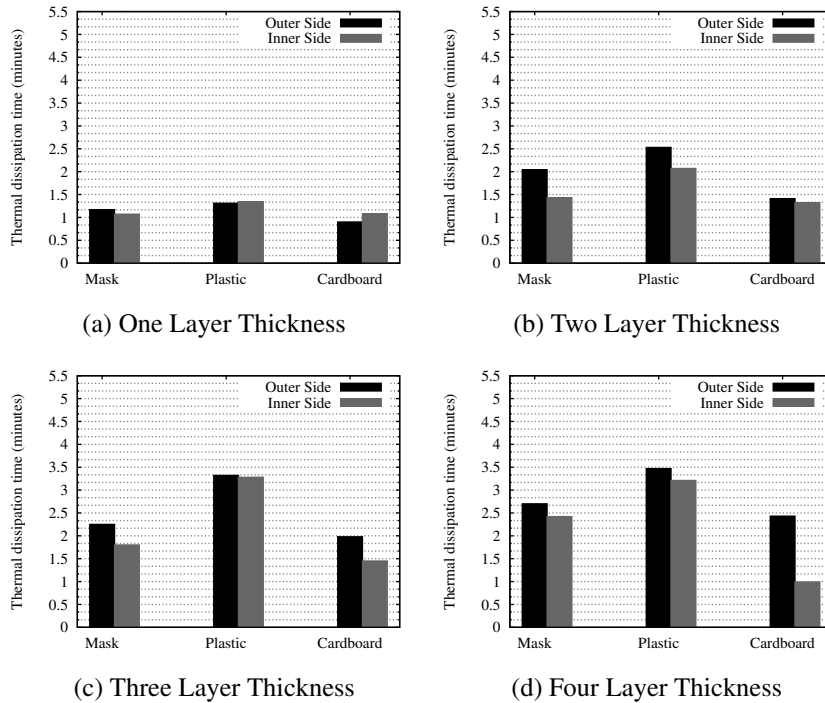


Figure 14: Effect on thermal dissipation with increase in the thickness of material.

measured the thermal dissipation time from the outer side of the object as a reference point. The average body temperature of the human subject holding the object ranged from 36.4°C to 36.7°C , while the ambient temperature varied between 22°C to 24°C . The experiments were conducted over six days, between 12:00 PM and 2:00 PM each day.

Results: Figure 14 illustrates the thermal dissipation measured from the inner side of the object, alongside the thermal dissipation fingerprints obtained from the outer side. The results reveal two key observations. First, the dissipation time increases with a larger surface area. Second, the dissipation time recorded from the inner side of the object is shorter than that from the outer side, with the magnitude of the difference being dependent on the material type and thickness (i.e., more thickness results in a larger difference). Indeed, the difference in dissipation time is smaller for the plastic bottle compared to the face mask or cardboard, highlighting that the effect of thickness is also influenced by the material covering the object. Nevertheless, the results demonstrate that thermal radiation from human touch can still be detected even in samples with multiple layers. In addition to demonstrating the robustness of dissipation fingerprints against partial or complete occlusion, these results suggest the potential for integrating MIDAS technology into everyday objects.

3.7. Thermal Dissipation Fingerprint for Multiple Objects

Our current prototype and experimental setup are designed for the individual characterization of objects using a single thermal dissipation fingerprint. However, in practical applications, humans frequently interact with multiple objects in their surroundings, necessitating the simultaneous analysis of thermal dissipation fingerprints from several objects. This section outlines an extension of the MIDAS pipeline to facilitate the prediction of multiple objects and validates the extended pipeline through benchmark experiments.

3.7.1. From Single to Multiple Objects

When analyzing the thermal fingerprints of multiple distinct objects, the primary challenge is to accurately separate and identify the various objects within the captured image. In thermal imaging, this task is notably more easier than in traditional computer vision. The dissipation of thermal fingerprints is not sensitive to the orientation or shape of an object, as these fingerprints are readily distinguishable from the background. However, as depicted in Figure 15a, the arrangement of objects can significantly influence the extraction of thermal dissipation vectors. When objects are not in contact with one another (dispersed), the thermal fingerprint of each object is easily identifiable. Conversely, when objects are in contact (agglomerated), the thermal fingerprint may spread across multiple surfaces.

We have extended the MIDAS system to support the analysis of multiple objects by incorporating an additional step in which the captured image is segmented into Region of Interest (ROIs). The current implementation processes normalized grayscale images and identifies different ROIs by employing the *Counter Approximation Method*¹. An illustration of this process is provided in Figure 15b. Once the ROIs are detected, each one is mapped to a dissipation vector corresponding to a single object. Given that thermal radiation spreads and transfers across the surfaces of objects, it is crucial to focus on the centroids of the hottest regions to ensure that the ROIs correspond to distinct objects rather than to the transfer of dissipating heat along the object's surface. MIDAS employs deterministic segmentation rather than machine learning (ML) to isolate the thermal ROI. The high natural contrast of thermal fingerprints makes traditional thresholding highly robust. This avoids the severe computational overhead of ML segmentation, conserving edge-device resources for the actual predictive task: classifying the thermal dissipation vectors.

3.7.2. Evaluation

We evaluated the performance of our extended approach in recognizing multiple objects simultaneously. Specifically, we measured the degradation in response time

¹Available in python CV2 package

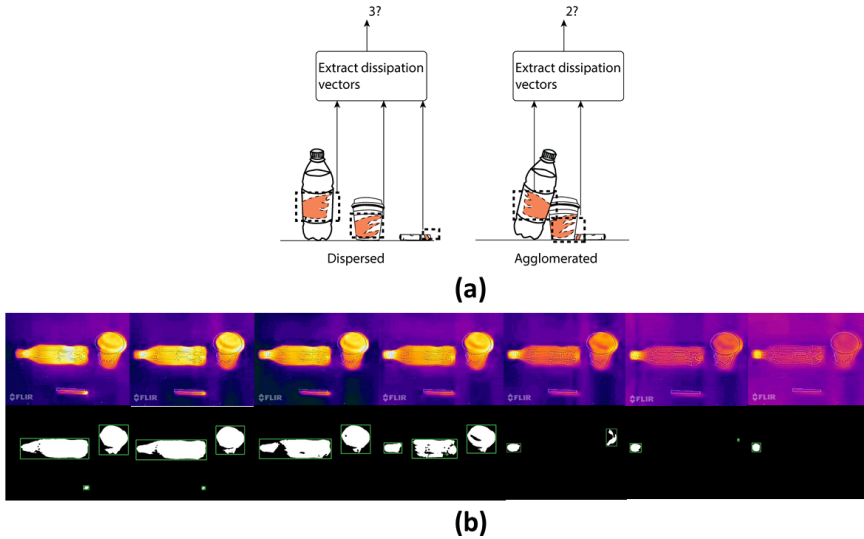


Figure 15: Detection of multiple thermal dissipation fingerprints. a) Object arrangements, b) Extraction of thermal dissipation vectors.









when detecting multiple objects within the same images and assessing the accuracy of the detection.

Experimental Setup: The duration and frame rate of video footage capturing thermal dissipation are critical factors that significantly impact the accurate identification of multiple objects. In this experimental setup, we fixed the frame rate at 30fps (i.e., 1800 frames per minute) and conducted experiments on videos with different length i) four videos of one-minute length; (ii) four videos from 30 seconds of length to 120 seconds in 30 second increments. In parallel, we also examined the impact on performance as the number of objects to be identified increased. Specifically, we analyzed the identification of two to four objects simultaneously, arranged in various configurations. Table 7 provides details on the grouping of object materials used in these experiments. We utilized the same household objects from our previous studies, which included a Coffee Cup (CC), Cigarette Butt (CB), Plastic Bottle (PB), and a Face Mask (Mask).

Results - Dispersed case: Figure 16 presents the results of our initial experiment, in which videos of consistent length and frame rate were used. The systems processing time naturally increases as more dissipation vectors are provided as input; however, this increase is negligible overall. Regarding the arrangement of objects, the results remain stable, indicating that the response time of MIDAS is not significantly influenced by the objects arrangement—provided that the dissipation vectors of the materials in the arrangement can be extracted without interference or noise.

Figure 17 presents the results for experiments conducted with varying video configurations. As the length of the video feed increases, the response time of MIDAS also increases. For videos up to 90 seconds in duration, the video length

Table 7: Table explaining grouping arrangement for prediction of Multiple Objects

Arrangement Group	Description
A	One Object at a time. Coffee Cup  , Cigarette Butt  , Plastic Bottle 
B	Two Objects at a time. "Coffee Cup + Plastic Bottle"  , "Cigarette Butt + Plastic Bottle"  , "Coffee Cup + Cigarette Butt" 
C	Three Objects at a time. "Coffee Cup, Plastic Bottle and Cigarette Butt" 
D	Four Objects at a time. "Coffee Cup, Plastic Bottle, Cigarette Butt and Face Mask" 

predominantly influences performance. It is only when the video length reaches 120 seconds that the impact of segmenting multiple objects becomes noticeable. In practical terms, thermal dissipation fingerprints remain stable across different video lengths, indicating that a shorter video duration of 30 seconds is sufficient for most scenarios. The results further demonstrate that the response time does not always increase proportionally with video length. This occurs because the response time is typically governed by the object with the longest dissipation time. Consequently, when one object has a significantly longer dissipation time than others, the overall response time tends to remain fairly consistent. The response time is largely dependent on the calculation of the area reduction in the image. Once the thermal fingerprint has fully dissipated, the area effectively reduces to zero, and the computations to return fast. In our experiments, both the plastic bottle and coffee cup exhibited longer dissipation fingerprints compared to the cigarette butt, resulting in the response time leaning towards longest vectors.

Results - Agglomerated case: When objects or their thermal dissipation areas intersect, the resulting overlap can cause dissipation vectors to become mixed. To evaluate MIDAS's performance under these conditions, we conducted experiments with materials in contact. Figure 18 illustrates the processing time of MIDAS for these scenarios. In this analysis, we focussed exclusively on arrangement groups B, C, and D, which involve multiple objects in contact. The results reveal an increased computational overhead, primarily attributable to the continuous need to separate thermal dissipation vectors in order to characterize each individual object.




Figure 16: Response time of multiple object identification when using MIDAS. (same video configuration). The bars within the figure correspond to different individual objects (A) or combination of objects (two objects: B, three objects: C, four objects: D) as given in Table 7.

Effect on Accuracy: Afterwards, we evaluated the impact of multiple objects on the accuracy of MIDAS. Our assessment is limited to the dispersed case, as performance is highly dependent on the quality of segmentation, which is challenging to control systematically. Indeed, in rare instances, segmentation may fail when the thermal fingerprints of two objects merge into a single "blob" in the image, rendering it impossible to separate them. When segmentation fails, the recognition process also fails. The accuracy for predicting single material objects is 83%. To evaluate the effect on accuracy when detecting multiple materials simultaneously, we begin with combinations based on the arrangement groups outlined in Table 7. We provided the model with combinations from the various arrangement groups using five different input video footages, each containing the same combinations, to assess the average prediction accuracy for multiple material detection. We then compared the predicted labels with the true labels to calculate the average accuracy of the model. For the configuration involving two materials arranged according to setup B, the model achieved an average accuracy of 79.97%. For arrangement C, which included three materials, the average accuracy was 77.75%. Finally, for arrangement D, which involved four objects, the average accuracy was 61.33%. Overall performance is influenced by the quality of segmentation, the resolution of the image, and the type of object being analyzed. As more objects enter the field of view, the size of each object decreases compared to scenarios where only a single object is present. This negatively impacts segmentation quality and the area of thermal dissipation, leading to a decline in overall classification accuracy. Similarly, the four-object configuration is the only scenario that includes the face mask, and part of the observed performance decline is due to the challenges associated with accurately recognizing the mask.

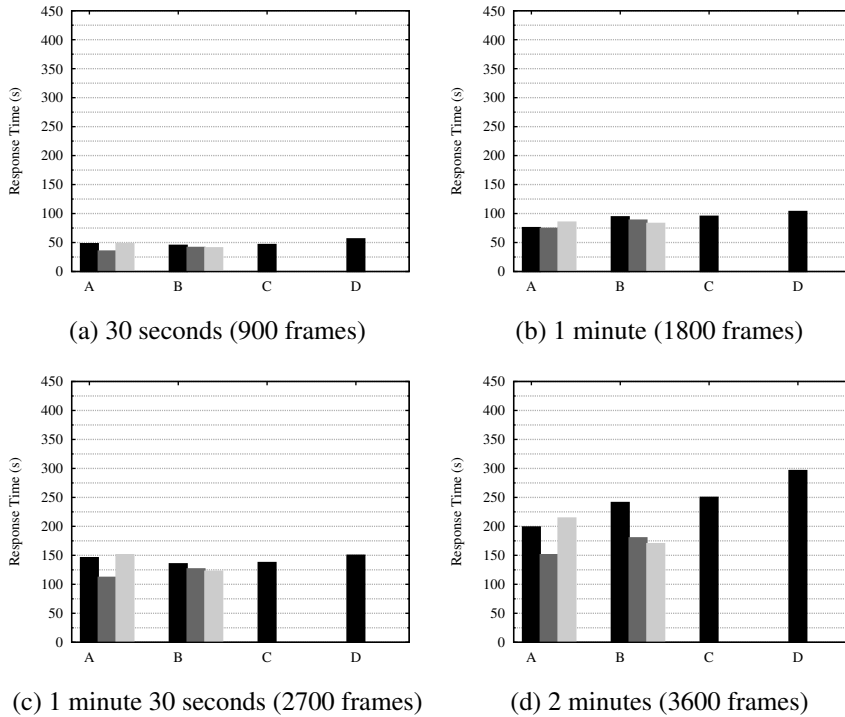


Figure 17: Impact on response time when using videos of different length and frame rate (Disperse objects case). The bars within the figure correspond to different individual objects (A) or combination of objects (two objects: B, three objects: C, four objects: D) as given in Table 7.

It is important to note that in such cases, objects that are poorly detected can be individually flagged for separate analysis, allowing for focused evaluation and potentially improving the overall accuracy of the analysis.

3.8. Discussion

Human temperature: Human temperature changes in cycles, being at its highest during hours of activity (day) and lowest during sleep (night) [206]. We demonstrate that interactions with objects can be used to characterize materials. The best results are obtained when the body temperature is stable, but the relative differences in thermal fingerprints are consistent across variations in body temperature. Conversely, interactions with an object of known material and in a stable environment can be used to detect relative differences in body temperature.

Room for improvement: Naturally, there are further challenges that need to be addressed to make our solution more robust across diverse environments. Adapting our approach to continuous monitoring requires accurate and noise-free thermal images, e.g., using calibration [184]. Not all materials can be characterized using




Figure 18: Impact on response time when using videos of different length and frame rate (Agglomerated objects case). The bars within the figure correspond to different combinations of objects (two objects: B, three objects: C, four objects: D) as given in Table 7.

our solution as thermal cameras have different emissivity ranges and some materials may reflect too much – or too little – thermal radiation. Such materials are usually used to preserve user’s privacy, e.g., ATM pin codes [4]. The experiments were conducted using thermal cameras integrated into smartphones, but in the future it may be possible to use cheaper alternatives, such as low-cost thermal array sensors [240]. There are also challenges arising from practical deployments as environments with oscillating temperatures can result in unstable fingerprints and as there can be situations where only partial dissipation fingerprints are available (e.g., due to the use of gloves). While outside the scope of our present work, there is also room for developing application areas, e.g., validating temperature differences with patients as part of clinical studies.

Other material properties: Dissipation time of thermal fingerprint gives insights about material types and correlates with emissivity. Thermal imaging could be used to potentially infer other material properties, such as thickness and elasticity. Potential use cases include detecting the pollutant type of marine plastics [91] and monitoring the decay in organic materials using differences in thermal dissipation characteristics.

Controlled baselines vs. real-world dynamics:: In our experimental setup, human touch was utilized as a reliable, controlled heat source to establish the thermal baselines necessary to prove the MIDAS concept. We demonstrated that the transient heat transfer from a stable body temperature can successfully characterize materials based on their thermal dissipation fingerprints. However, translating this to pervasive environmental monitoring requires shifting the thermal source from human interaction to ambient environmental dynamics.

Robots and autonomous devices: Thermal radiation analysis of objects touched by humans can be used to inform and train different robots and autonomous devices, e.g., UAVs [176], about the material properties of objects. New sensing and interaction modalities can also be envisioned as part of robotic systems, e.g., incorporating heat sensation to detect the material of an object and to enable autonomous devices to adjust their operations with objects in the surrounding environment [171]. For instance, a robotic arm can rely on a camera to detect an empty bottle on a table. However, the arm's pressure to lift and put the bottle should be proportional to the bottle material. Otherwise, the robotic arm can break apart the bottle, e.g., plastic vs. glass bottle. Besides, since it is possible to differentiate thermal radiation emitted by different genders, autonomous devices can adjust their interactions accordingly.

Application scenarios: Transitioning outdoors:: The core focus of this chapter has been establishing the fundamental viability of thermal dissipation as a material-sensing modality. While the validation occurred indoors using human contact, the most significant potential for this technology lies in outdoor, autonomous applications. Rather than relying on human touch (which is impractical for large-scale waste sorting), future deployments must leverage environmental heat sources—specifically solar radiation. As everyday objects and litter absorb sunlight, their subsequent cooling periods (e.g., when a cloud passes or the sun sets) create the exact thermal dissipation curves modeled by MIDAS. Integrating this passive thermal sensing with AGVs or drones enables the large-scale classification of litter thrown in terrestrial and aquatic environments [333]. Exploring and scaling this transition from controlled indoor thermal baselines to dynamic outdoor solar baselines forms the exact foundation for the autonomous systems explored in the next chapter.

3.9. Applications for Environmental Sustainability: Outcomes and Implications

The MIDAS sensing modality, leveraging thermal dissipation characteristics for material identification, represents a novel and promising approach that directly supports environmental sustainability objectives. By detecting unique thermal fingerprints left on objects through human interaction or environmental exposure, MIDAS enables accurate characterization of waste materials, which is vital for enhancing recycling efficacy and reducing environmental pollution. For example,

in household waste management, MIDAS can facilitate automated sorting by identifying plastics, glass, and organic matter based on residual thermal signatures, thereby decreasing contamination in recycling streams and contributing to circular economy initiatives [246, 91].

Beyond static waste sorting, MIDAS holds significant potential when integrated with autonomous environmental monitoring platforms such as ground vehicles or robotic systems. These autonomous devices can leverage the thermal dissipation signals to classify litter dynamically in real-time, enabling scalable and cost-effective monitoring of plastic pollution in both terrestrial and aquatic environments [333]. Such applications address critical global concerns, including microplastic contamination and ecosystem degradation, advancing practical solutions aligned with sustainable development goals.

Furthermore, MIDAS's ability to infer additional material properties – such as thickness and elasticity – from thermal dissipation patterns offers valuable insights for ecosystem health monitoring. Detecting the decay or transformation stages of organic materials can inform timely conservation interventions and improve environmental management strategies [91] [334]. This capability enhances the responsiveness and precision of ecological monitoring systems, which is essential for maintaining biodiversity and ecosystem services.

Technological trends toward miniaturization and reduced costs of thermal imaging sensors also play a crucial role in expanding MIDAS's applicability. Embedding such sensing capabilities into everyday objects – from medical containers supporting controlled medication management to interactive toys – opens avenues for pervasive environmental sensing with minimal intrusiveness [85, 325]. The anticipated evolution of low-power, compact thermal sensors facilitates widespread deployment, enhancing real-time material identification and environmental data collection on larger scales.

While challenges remain – such as ensuring robustness under variable environmental conditions, dealing with sensor calibration and thermal noise [184, 121, 340], and adapting to diverse material emissivities [120] – ongoing research and development efforts promise to overcome these hurdles. Addressing these issues is critical for transitioning MIDAS from controlled experimental settings to practical, long-term deployments in dynamic real-world environments.

In conclusion, MIDAS exemplifies the integration of fundamental thermal sensing science with applied environmental sustainability. It provides a versatile and scalable toolset that enables improved waste management, proactive plastic pollution monitoring, insightful ecosystem health assessment, and intelligent autonomous system behavior. By enhancing material characterization capabilities within pervasive sensing frameworks, MIDAS contributes substantively to the advancement of sustainable environmental monitoring and management technologies.

3.10. Summary

The MIDAS presents an innovative sensing solution that leverages thermal dissipation patterns resulting from human touch to characterize and identify everyday object materials. The system operates by capturing thermal dissipation pattern changes over time using thermal imaging and modeling these variations to infer material properties. Through comprehensive empirical validation, MIDAS demonstrates remarkable robustness to variations in human interaction patterns and achieves material classification accuracy of up to 83% based solely on thermal dissipation fingerprints.

The research establishes that the thermal dissipation characteristics provide a reliable material signature that transcends shape-dependent or visual-based classification approaches. MIDAS operates effectively within a practical range of approximately one meter, making it suitable for real-world deployment scenarios. The system's ability to extract material-intrinsic properties rather than relying on object-specific features enhances its generalizability across diverse applications.

The validation process included extensive controlled experiments and multi-participant user studies, demonstrating consistent performance across different interaction modalities and user behaviors. Comparative analysis with state-of-the-art computer vision and light reflectivity approaches revealed MIDAS's superior robustness due to its focus on fundamental material characteristics rather than superficial object attributes.

MIDAS significantly advances environmental sustainability by enabling automated waste classification at the point of disposal, facilitating more efficient recycling processes and reducing contamination in waste streams. The technology's ability to identify materials through natural human interactions eliminates the need for specialized sorting procedures, thereby reducing energy consumption and operational costs in waste management systems. By accurately distinguishing between different material types, MIDAS contributes to improved circular economy practices, enhanced resource recovery rates, and reduced environmental burden from improper waste disposal. Furthermore, the system's potential for integration into smart waste management infrastructures positions it as a valuable tool for advancing sustainable urban development and environmental stewardship.

4. LIZARD - PERVASIVE SENSING FOR PLASTIC LITTER MONITORING

This chapter introduces LIZARD, an advanced autonomous sensing system designed specifically to address the detection and monitoring of plastic litter in natural and urban environments. By integrating thermal imaging with light reflectivity sensing, the LIZARD system enables reliable detection of a wide range of litter sizes, including macroplastics as well as smaller micro- and meso-particles that are often missed by conventional detection methods. Unlike traditional litter monitoring approaches that depend on labor-intensive manual inspection or costly, stationary laboratory instruments, LIZARD enables in situ, scalable, and energy-efficient mapping of plastic pollution. Its lightweight machine learning models and seamless integration with autonomous ground vehicles facilitate high-resolution monitoring over large areas with minimal human intervention. Through extensive field and laboratory validation, LIZARD has demonstrated significant improvements in accuracy, especially for smaller, fragmented plastics, which pose serious long-term threats to ecosystems and human health. By empowering automated, targeted clean-up and advancing early detection of persistent pollutants, this chapter makes a substantial contribution toward advancing sustainable waste management practices and protecting both terrestrial and aquatic environments from the enduring impacts of plastic litter.

4.1. Introduction

Littering, caused by human activity, is a major concern, especially in densely populated areas. It arises from inadequate waste management practices, inefficient garbage collection systems, and a general lack of awareness regarding the severity of the issue [37, 178]. The detrimental effects of litter are numerous, impacting both natural ecosystems and human health [270, 132]. For instance, animals such as birds and dogs may ingest litter, ultimately introducing these pollutants into the global food chain [304, 238]. One of the primary challenges associated with litter pollution is its persistence in the environment; once introduced, it becomes difficult to completely remove. The diverse types of litter, including plastics, can remain in the ecosystem for extended periods, often blending with soil and other natural elements. Additionally, such materials do not naturally decompose but instead fragment into smaller particles due to environmental degradation, human activities, and weather conditions. Consequently, litter pollution constitutes a significant and multifaceted problem that demands immediate attention and action.

To effectively monitor and mitigate this pollution, it is critical to first establish the physical classifications of plastic debris targeted in this research. Environmental plastics are generally categorized by size into three primary groups: macroplastics (large, easily identifiable items typically > 15 mm), mesoplastics (fragments

between 5 mm and 15 mm), and microplastics (particles < 5 mm). The detection of these smaller litter fragments is vital for reducing the long-term negative effects of pollution. However, scalable, cost-effective, and accurate solutions are currently lacking. Present state-of-the-art detection methods rely on specialized instruments, such as spectroscopy, to analyze samples from specific locations. While these methods are effective, they cannot be integrated into consumer devices or made accessible to end-users, as they require advanced technical knowledge and expertise [208]. As a result, the most common method for litter removal relies on human effort, either through volunteers or contract workers. This approach is most effective for recently deposited litter, as identification typically depends on manual visual inspection. This method is best suited for litter that remains intact and has not yet blended into the surrounding environment [246]. These activities can be effective when conducted regularly, particularly before litter objects begin to break down. However, these activities are expensive and require substantial planning and logistics, making them unsuitable as a scalable long-term solution. Some technological solutions have been proposed to aid in litter cleanup, such as autonomous sweepers [289] or aerial UAVs [177] for detecting litter. Unfortunately, these technologies are limited to identifying large, easily visible objects. The effectiveness of such solutions could be improved by incorporating additional sensing technologies, such as cameras mounted on vehicles that use object recognition techniques to detect litter. However, these solutions also face challenges in scalability and are ineffective at detecting smaller, hazardous litter fragments [209]. Moreover, current object detection techniques have limited ability to discriminate between objects, particularly when the objects have lost their original shape, such as fragments of broken bottles, further complicating the detection of smaller or degraded litter.

We present LIZARD, a novel and innovative pervasive sensing solution designed for monitoring plastic litter. LIZARD is capable of detecting both larger (macro) litter and smaller plastic fragments, such as micro- and meso-plastics. Our focus on plastics stems from their identification as one of the most significant environmental hazards due to their extensive use and resistance to decomposition. Microplastics, in particular, have recently been recognized as a serious health concern, with a significant portion originating from ground litter. LIZARD employs a cutting-edge sensing approach that integrates two complementary modalities: thermal imaging and optical sensing based on light reflectivity. These modalities are particularly suited for autonomous monitoring because they utilize lightweight machine learning models, unlike computer vision systems, which require complex and energy-intensive models. LIZARD's initial step is to analyze the thermal characteristics of objects induced by sunlight using thermal imaging. As thermal radiation decreases with the reduction in object size, LIZARD incorporates a second sensing modality based on light reflectivity to detect smaller litter fragments in the environment. However, since light reflectivity is sensitive to ambient light conditions, it requires close proximity to the litter to be effective. To mitigate this

limitation, LIZARD uses thermal imaging analysis to identify optimal areas for sampling, and subsequently applies close-contact light reflectivity sensing to those regions. This approach reduces the time spent using light reflectivity and lowers the overall energy consumption by minimizing the mapping effort. We conducted an extensive and rigorous evaluation of LIZARD through a combination of controlled laboratory experiments and real-world deployments in three real world location scenarios. The results demonstrated that LIZARD accurately detects three categories of litter: micro, meso, and macro plastics. Additionally, we illustrate how LIZARD can be easily integrated with commercially available autonomous ground drones, enabling autonomous litter monitoring. Our work lays the foundation for innovative approaches to improve environmental sustainability and mitigate the harmful effects of littering. LIZARD provides a comprehensive solution that surpasses existing methods by detecting both large plastic litter (macro) and smaller litter fragments (micro and meso) while offering compatibility with autonomous ground vehicles for scalable deployment.

4.2. Related work

Table 8 shows a summary of relevant work that investigates the identification of litter, and the use of personal and autonomous devices to achieve it.

Pervasive sensing: Different smartphone sensors have been re-purposed to identify materials and evaluate their inherent properties on the spot [319, 80]. Determining sugar content in liquids has been explored using near-infrared sensors [134]. Soil and moisture salinity has been explored through wireless signals [327]. Characterization of liquids using wireless also has been explored [71]. Several other sensors have been used for liquid characterization, e.g., acoustic [328], vibration [335], and light sensors [124]. Optical sensing has also been adopted, e.g., to measure water quality [88], to characterize underwater waste [93], and to detect ripeness of organic produce [349]. Cameras have been used extensively to analyze liquid's tensions [336] and identify materials via material reflection [330]. RFID stickers have been used to learn the quality of food [110]. Near-infrared spectroscopy devices have been developed to identify medical pills [148], non-invasive blood glucose level monitoring [285], and brain activity monitoring [331]. Near-infrared spectroscopy is also used to estimate the quality of fruits [257]. Thermal imaging also has been proposed to detect different material types [55]. Unlike others, our work combines different sensing modalities to identify plastics at different size scales, including, meso, micro and nano.

Plastic identification: Manual cleaning campaigns have been adopted as a solution to overcome the litter removal problem. Automated solutions also have been developed to tackle the problem. Computer vision methods (object recognition) are the most common autonomous approaches to identify plastics [37]. Methods for identifying plastics in oceans [17] and removing plastic litter from public spaces [279] at macro level have been investigated. Optical solutions have been studied to identify

Ref	Plastic size	Integration with AV	Techniques	Area Mapping	Off-the-shelf components
[55]	Macro	No	Ambient Light	Full	Yes
[80]	Macro	Yes	Human touch	Partial	Yes
[333]	Macro	Yes	Sunlight	Full	Yes
[224]	Macro	No	Object recognition	Full	No
[104]	Macro	No	Sunlight	Full	No
[279]	Macro	Yes	Object recognition	Full	No
[37]	Macro	No	Object recognition	Full	No
[152]	Macro	No	Object recognition	Full	No
[178]	Macro	No	Object recognition	Full	No
[227]	Macro	No	Object recognition	Full	No
[313]	Macro	No	Object recognition	Full	No
[316]	Macro	Yes	Object recognition	Full	No
[91]	Macro	Yes	Light	Full	Yes
[93]	Macro	Yes	Light	Full	Yes
[330]	Macro	No	Light	Full	Yes
[346]	Micro	No	Light	Full	No
Our work	Macro & Micro	Yes	Light	Partial	Yes

Table 8: State-of-the-art for identification of plastics and potential integration with autonomous vehicles (AV).

different types of plastics underwater [333, 91]. Macro-plastics (more than 200 mm in size) are easy to detect by the human eye [81]. Despite several approaches available to detect plastics, a key problem that has been partially overcome is the detection of micro-plastics. Indeed, unlike macro plastics, the micro plastics (< 5mm in size) are difficult to detect due to their size. Microplastics are generated through the degradation of macro and meso plastics, as they break down in response to environmental conditions. [209]. The most common method to detect micro plastics are based on spectroscopy techniques, like fourier transformation infrared spectroscopy (FTIR) and Raman spectroscopy [208]. While these methods are highly accurate, they can only be utilized in controlled environments. Moreover, these methods are expensive and require bulky instruments. Thus, they cannot be applied in the wild contexts or integrated into autonomous vehicles with ease. Our work explores the identification of different plastic particles in the wild, affected by different environmental factors, e.g., sunlight, and different surfaces. Moreover, we also demonstrate how our solutions can be integrated into autonomous ground vehicles.

4.3. Motivation

The most commonly employed automated method for supporting litter detection is vision-based object detection [16]. However, this approach exhibits limited accuracy when deployed in real-world environments, as it primarily detects larger and intact litter objects. In addition, it is energy intensive, making it unsuitable for

use in resource-constrained devices, such as drones [345]. We demonstrate that plastic identification can be integrated into AGVs using lightweight thermal sensing. Thermal imaging has proven to be a promising technology for characterizing the material composition of various litter objects [80], and it can utilize opportunistic heat sources, such as sunlight, to lower the cost of the sensing process [333]. Our results indicate that these methods can improve plastic detection and removal by accurately identifying the material, thereby ensuring that proper recycling processes can be applied.

Testbed: We conducted experiments to measure the thermal dissipation time of plastics across different size categories: macro-scale (diameter > 10mm), meso-scale (5 to 10mm), and micro-scale (< 5mm). For the macro-scale analysis, we focused on plastic objects and used specialized samples produced through a standardized manufacturing process. This approach ensures that the objects are comparable and that any variations in thermal radiation are attributed to differences in the polymer material. A 60 W bulb was used as the heat source to induce thermal radiation. For the micro and meso-scale plastics, we used samples collected directly from the environment, as detailed in Section 3.5. Thermal radiation in these samples was induced using a Janeli Incubator, which provides a consistent amount of thermal energy. A commercially available FLIR CAT S61 smartphone was employed as the thermal sensing unit for all experiments.

Procedure: Plastic samples were first exposed to ambient temperatures ranging from 22 °C to 24 °C prior to the start of the experiment. Each sample was then placed on a black surface, with a lamp positioned at a fixed distance of 10 cm to prevent burn damage while ensuring adequate thermal radiation exposure. The black surface was chosen to ensure consistent measurements by absorbing additional thermal radiation. We evaluated four heating durations (1, 2, 3 and 4 minutes), corresponding to varying initial temperatures, and measured the dissipation of the thermal footprint over time. Longer durations were not considered since a 4-minute duration was sufficient to reveal the relationship between dissipation time and emissivity coefficient. Similarly, micro- and meso-scale samples were placed in an incubator for approximately 15 minutes at a temperature of 30 °C. To verify that the incubator induced sufficient thermal radiation in the plastic fragments, macro-scale samples were also included in the experiment. Figure 19 presents the arrangement of the samples. Furthermore, two types of plastic fragment arrangements were considered: agglomerated (AA) and dispersed (DA). In the agglomerated arrangement, the samples overlapped, allowing thermal radiation transfer between them, whereas in the dispersed arrangement, the fragments were separated from one another.

Results: The thermal dissipation time of plastic samples, as shown in Figure 20, can serve as an effective method for identifying different material types. A Friedman test, using plastic type and thermal dissipation time as variables, revealed statistically significant differences between the plastic materials ($\chi^2(2) = 61.47$, p




Figure 19: Combination of macro-plastics and micro-plastics (agglomerated and dispersed).

$< .05$, $W = 0.80$), confirming that thermal radiation can successfully distinguish between them.

Although thermal radiation can be easily measured for macro plastics, Figures 20(e-f) demonstrate that this analysis cannot be reliably applied to plastic fragments. As shown in Figure 20e, when fragments are agglomerated (i.e., overlapping), there is sufficient thermal radiation to generate a thermal footprint. However, accurately identifying the material becomes difficult, particularly when the fragments consist of different materials. In contrast, when fragments are dispersed, thermal radiation is insufficient to produce a detectable footprint. Additionally, we evaluated the thermal dissipation times of the macro samples using an incubator as the heat source. The results, presented in Figure 20f, indicated that the dissipation times exhibited a similar relative trend when compared to those using a bulb as the thermal radiation source. A Kolmogorov-Smirnov test confirmed that the differences in thermal dissipation times were not statistically significant between the two heat sources ($KS=0.2$, $p > 0.05$). This finding suggests that uniform heat distribution across the object is not necessary; even a small region with consistent heat points is adequate for detecting plastic materials.

Insights: Our results indicate that while thermal dissipation time is effective for characterizing macro plastic samples, it falls short in identifying smaller plastic fragments. Micro and meso samples can only be reliably characterized when they are agglomerated, a condition that is rarely observed in natural environments, as will be demonstrated in our main experiments in Section 3.5). Consequently, relying solely on individual sensing modalities is inadequate for detecting plastics across different scales. The LIZARD system addresses this challenge by combining two complementary sensing modalities, offering a more comprehensive solution to the issue of scale variability.

4.4. LIZARD Sensing Pipeline

LIZARD has been developed for autonomous litter monitoring, a task that presents considerable challenges. The operations must be lightweight enough to avoid interfering with the AGVs primary functions, such as autonomous navigation.

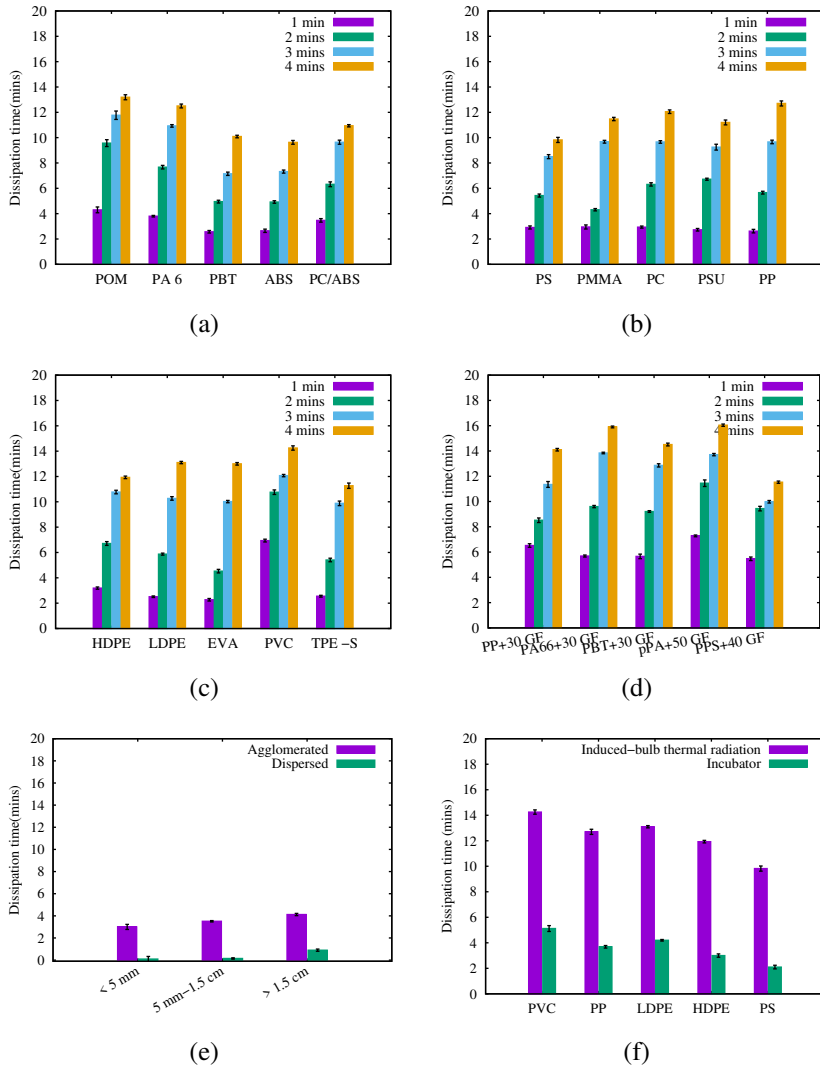


Figure 20: Thermal dissipation times for (a-d) 20 plastic samples, (e) Micro and meso plastics, and (f) Selected macro plastics.

Moreover, the litter detection system must function reliably across various surfaces and weather conditions. Since AGV's are primarily focused on navigation, any auxiliary tasks, like litter monitoring, must operate with minimal resource usage. LIZARD addresses these challenges by employing two complementary sensing modalities and a two-phase sensing pipeline that minimizes unnecessary processing. The overall pipeline, outlined in Figure 21, begins with a thermal camera that monitors thermal dissipation patterns to identify larger macro fragments, generating a region of interest (ROI). This ROI is then further analyzed using light sensing with an adaptive sampling scheme, designed to minimize excess processing. The

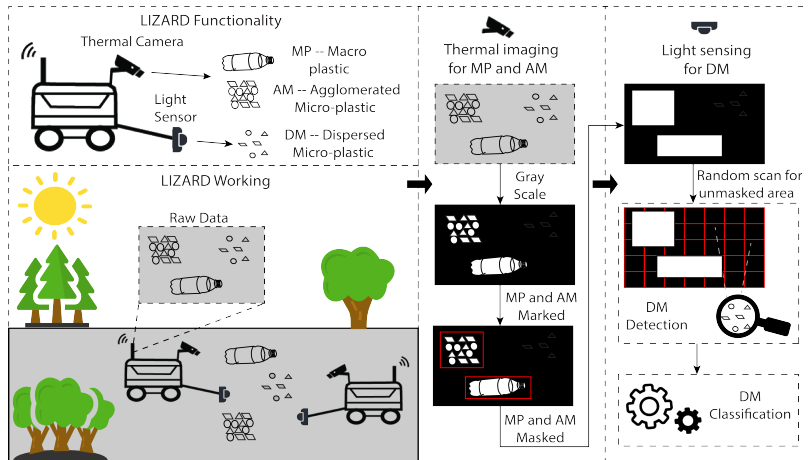


Figure 21: Plastic fragment sensing and monitoring pipeline.

following sections provide detailed descriptions of each phase.

4.4.1. Phase 1 - Thermal imaging

Pre-processing: In the first phase, video footage is captured with a thermal camera and then converted into a sequence of grayscale images for preprocessing. This grayscale conversion aids in isolating the background and segmenting the image into distinct objects. In order to reduce thermal noise, caused by the camera’s low resolution and inherent inaccuracies, a Gaussian Blur is applied to smoothen the image. Following this, the grayscale image is converted into a binary format (0-255) to estimate the thermal dissipation time, allowing the isolation of visible heat sources as thermal footprints. To further enhance the visibility of these thermal footprints, each image is processed using the Adaptive Histogram Equalization technique, which sharpens edge definitions. These steps are summarized in Algorithm 1 (steps 1-3) and adhere to best practices in thermal image processing [80, 239].

Reducing sampling area: Since macro litter is easily detectable through thermal imaging, our focus shifted more towards challenging task of identifying smaller plastic fragments in the environment. Mapping an entire area for these fragments is resource-intensive, making it unsuitable for devices with limited resources. Therefore, we first focussed on reducing the area to be explored. It is important to highlight that, in practice, moving the AGV consumes the most resources. The sensing components of LIZARD are designed as a separate payload with their own power source and computing resources, allowing us to focus on optimizing the resource consumption of the sensing pipeline. Given that most current AGVs do not readily support external component integration or offer extensive programmability, using a separate payload is the most feasible approach.

The main technical challenge is how to reduce the area under detailed analysis, as the size of the area directly affects resource consumption. To address this, we

first exclude regions where macro litter has been identified, as these areas will already be designated for cleaning. Our primary focus is on regions where smaller plastic fragments may remain in the environment, even after cleaning efforts. To refine the analysis, the binary images are further processed by removing the white regions representing heat sources. These heated areas are detected and extracted using OpenCV's Contour Approximation Method. Once macro litter is identified through contour detection, we create a patch that masks these larger samples, as detailed in Algorithm 1 (lines 4 to 13). Since these areas do not contain plastic fragments, sampling is unnecessary, thereby reducing the sampling area.

Algorithm 1: Reduce sampling area of sensing and calculate ROI coordinates

Data: ThermalImage containing plastic fragments
Result: *patch*; *coordinates* of plastic fragment ROIs

- 1 *image* \leftarrow Conversion to Grayscale;
- 2 *image* \leftarrow Apply GaussianBlur;
- 3 *image* \leftarrow Convert to Binary(0 or 255);
- 4 *c* \leftarrow Detect all contours from image;
- 5 initialize *contourarea*;
- 6 **for** *each contour* in *c* **do**
 - 7 *contourarea* \leftarrow area of each contour;
 - 8 *area90thquantile* \leftarrow 90th quantile of *contourarea*;
 - 9 **if** *contourarea* > *area90thquantile* **then**
 - 10 *patch* \leftarrow *contour* Bounding Box Coordinates;
 - 11 apply bounding box *patch* on *image* ;
 - 12 **end**
- 13 **end**
- 14 *ROI* \leftarrow calculate ROIs from remaining sampling area;
- 15 *RandomSample* \leftarrow Random sample from *ROIs*;
- 16 *BoundingBox* \leftarrow Bounding box of *RandomSample*;
- 17 *centroid* \leftarrow Centroid of *BoundingBox*;
- 18 *coordinates* \leftarrow Coordinates of *centroid*;
- 19 return *coordinates* of *centroid*

4.4.2. Phase 2 - Light reflectivity

Defining ROI mapping area: Next, we define regions of interest (ROIs) within the remaining unexplored areas. To accomplish this, we divide the area into multiple ROIs of uniform height and width. While this approach has been successfully applied to generate images at different resolutions [321], in our case, it is used to define specific areas for sampling. We calculated the coordinates of the small ROIs designated for sensing. The key principle here is that sampling is conducted within these small areas. To accomplish this, we first employed a Canny Edge Detection

method in OpenCV on the thermal image to identify ROIs from the remaining unexplored regions. This technique isolates the boundaries of plastic fragments. Subsequently, we performed Morphological Transformation (dilation) to enhance the boundaries by adding pixels, which is essential because the boundaries may sometimes be only partially visible. Dilation helps interpolate and complete any missing boundary sections. As a result, the plastic fragment boundaries become more prominent and thicker. We then mask the large plastic samples using the patch generated previously in Algorithm 1. Finally, Contour Detection is used to extract the ROI area after the plastic fragment boundaries are clearly defined.

Scheduling light sampling and classification: The longer the monitoring system can operate, the larger the area it can cover, leading to more accurate identification of litter. However, movements of the ground vehicle and its actuators, such as robotic arms, result in significant power consumption, limiting the duration of monitoring based on the area being sampled. One of the key challenges for LIZARD is optimizing the balance between monitoring accuracy and the cost of ground drone movements. We address this challenge by employing an adaptive sampling strategy that analyzes the ROI to identify areas where light reflectivity sampling should be conducted. Since the extracted ROIs may include noise due to varying backgrounds, random sampling is performed within the extracted ROIs (contours in this case) to reduce the cost of sensing and monitoring. After random samples are generated, the centroid of the bounding box enclosing the region of interest is calculated, and the coordinates are passed to the light reflectivity sensor for sensing and monitoring, as detailed in Algorithm 1.

Sampling robustness: Given that microplastics can be present in any environment, another significant challenge for LIZARD is ensuring reliable operation across various environmental backgrounds. Certain locations may be more challenging to map than others due to factors such as dirt, humidity, and wind. To enhance the identification of sampled data, we performed wavelet transform [166], on the sensor data for enrichment and denoising. This process involves filtering, thresholding, and scaling of light values. Scaling is performed to generate wavelet coefficients at various scales, enabling multiresolution analysis. Following this, we used change point detection to analyze the reconstructed light sensor data, as illustrated in the Figure 22. For this, we employed the *roerich* technique, which performs change point detection based on density ratio estimation [125]. By specifying the sliding window size, the outcome of the change point detection includes both the detected change points in the light values (y-axis) and their magnitudes, as shown in the Figure. These detected change points indicate the presence of plastic fragments, while the magnitude represents the type of plastic. The PR-AUC for the true and predicted change points was approximately 0.94 (precision: 0.88; recall: 1.0). Additionally, as part of our pipeline, after light samples from a ROIs are denoised and preprocessed, they are fed into classical machine learning models to classify the type of plastic fragment. LIZARD employs




Figure 22: Wavelet transformation of light sensor data (top), change point scores (middle) and change point detection (bottom).

RF and SVMs to facilitate the classification of plastic fragments.

Technical challenges: LIZARD integrates thermal imaging to detect large plastic fragments with light reflectivity for identifying smaller particles, addressing the challenge of high costs associated with using sensors to explore entire areas. Additionally, LIZARD implements an enrichment method to improve the robustness of microplastic identification in diverse environmental conditions. Further challenges include hardware design, sensor placement, field of view, and sensor calibration. Another critical issue to overcome is the dynamic adaptation to changing environmental factors for effective operation in the wild. Despite these challenges, LIZARD currently operates efficiently across various environments, providing core functionality that is easily extendable.

4.5. Experimental Setup

We conducted a comprehensive evaluation of LIZARD through rigorous experiments to demonstrate its performance in identifying plastic fragments. The evaluation begins with controlled experiments, illustrating how light reflectivity can be utilized to characterize plastic fragments collected from real-world environments. Following this, we demonstrate how our method can be effectively applied to detect plastic fragments dispersed in natural, uncontrolled settings.

Apparatus: We utilized an off-the-shelf Caterpillar (CAT) smartphone for our experiments. To induce thermal radiation in the plastic samples, we employed a JANOEL18S incubator as the heat source. For the light reflectivity source, we used red light diodes and photoresistors with a wavelength of 650 nm. Additionally,

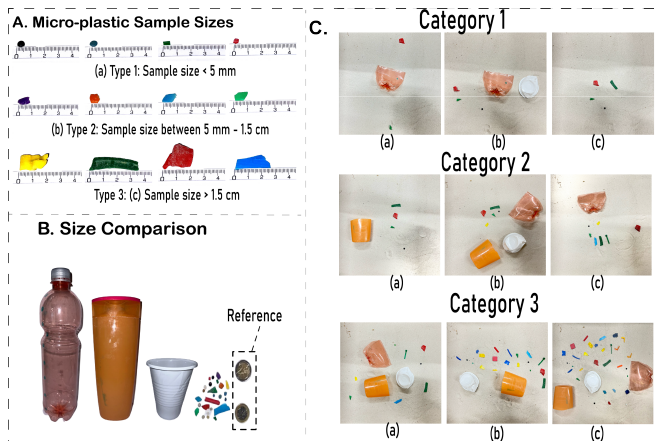


Figure 23: A. plastic types: (a) 1: Micro (< 5 mm), (b) 2: Meso (5 mm to 15mm), and (c) 3: Macro plastics (> 15 mm); B. Plastic size comparison using coins as a reference; C. Arrangements of micro, meso, and macro plastic samples used in our analysis.

we developed an array of 15 light sensors to simultaneously map an area from an autonomous vehicle.

Plastic samples: Plastic fragment samples were collected from the wild using a specialized process [34] and categorized into three types, as illustrated in Figure 23a: Type 1 (< 5 mm), Type 2 (> 5 mm and ≤ 15 mm) and Type 3 (> 15 mm). Type 1 represents microplastics, while Type 2 represents mesoplastics. These two categories are the most concerning, as microplastics can be easily ingested by animals or transported into aquatic environments, where they enter the food chain [15]. Mesoplastics are typically in the process of breaking down into smaller fragments, serving as a source of future microplastics. Type 3 corresponds to macroplastics. The experiment also included macroplastic samples from end products, as depicted in Figure 23b, such as a plastic bottle (PET), a shampoo bottle (HDPE), and a disposable plastic cup (PP). These samples were arranged together to simulate plastic litter, enabling simultaneous analysis.

Sample arrangements: To evaluate the feasibility of detecting plastic fragments in the environment, we begin by analyzing their natural distribution across different settings. This is achieved through the analysis of TACO dataset [283], a publicly available image dataset depicting litter in natural environments. The dataset contains over 5,000 images of litter, ranging from macro to micro sizes. To investigate how micro and meso litter is distributed, we extract the relevant images from the dataset. The primary goal of this analysis is to establish typical arrangements of plastic fragments, which can then be analyzed by our method in a controlled setting. The arrangements used in the experiments are presented in Figure 23c and categorized into three groups (see Section 3.6 for further details on the categories).

Procedure: After identifying the arrangements of plastic fragments, they are

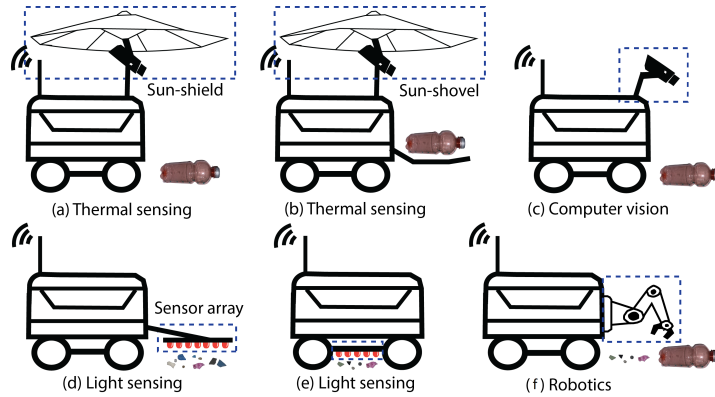


Figure 24: Lizard design alternatives.




Figure 25: Lizard prototype, (a) Main Components; (b) Underlying Sensor array; (c) Lizard in action

replicated in a controlled testbed. An exploration area with different backgrounds is developed, and the arrangements are mapped onto it. The arrangements are normalized to fit within an area of 20 by 14 cm. Once positioned in the exploration area, the setup is placed in an incubator to induce thermal radiation on the samples. Afterward, the exploration area is removed from the incubator, and the LIZARD pipeline is applied for further analysis.

Design: There are a wide range of options for integrating LIZARD onto AGVs. Figure 24 presents potential designs and highlights the components that must be added to the AGVs. One alternative is to install a simple sun-shield and use thermal imaging [333], while another relies on computer vision (Figure 24(c)). Alternatively, light sensors can be integrated into the lower part of the AGV to take light reflectivity measurements (Figure 24(d-f)). In our study, we used a fixed sensor array (Figure 24(d)) to focus on evaluating sensing performance. However, in practical applications, a robotic arm would be preferable as it could also aid in removing plastic fragments. We employed light-dependent resistors (LDRs) as sensors, which were sufficient for our purposes.

Implementation: We have developed a proof-of-concept pipeline, which has been integrated with a commercial off-the-shelf AGV. The working prototype

is illustrated in Figure 25, where its various components are highlighted. The prototype utilizes a sensor array comprising 15 light sensors, spaced approximately 1.5 cm apart, capable of mapping an area of 20x14 cm². This sensor array is mounted on an arm equipped with a linear guide slider, enabling it to move back and forth via a stepper motor (NEMA17). Each light sensor output is connected to the analog input pins (pins 1 to 15) of the Arduino Mega ADK microcontroller, and the photoresistors linked to the light sensors measure the light intensity by detecting resistance. Algorithm 1 maps the coordinates between the ROIs in the image and the area covered by the sensor array. The sensor array moves to the specified locations, activating the corresponding light sensors to collect data from each target area. The movement of the sensor array is controlled by the stepper motor driver (A4988), which is connected to the Arduino microcontroller. After the data samples are collected, they are uploaded to a web server with a timestamp for further analysis.

4.6. Results

Light reflectivity performance: Figure 26 demonstrates the performance of the proposed light reflectivity method in identifying different types of plastic litter fragments from varying distances. Figure 26a presents the results for sampling conducted at distances of 2 cm and Figure 26b at 5 cm distance respectively. The results indicate that light reflectivity can effectively characterize individual plastic fragment samples. A Kruskal-Wallis test revealed no significant differences between the plastic samples ($\chi^2 = 10472.83$, $\eta^2 = 0.99$, $p < 0.05$). Although slight variations in sample characterization are observed when using different sampling distances, the relative differences between the types of characterization are preserved. Additionally, a Kolmogorov-Smirnov test confirmed that the light reflectivity values from both distances were similar (KS=0.235, $p > 0.05$). It is also evident from the figure that the ability to differentiate between samples decreases as the distance increases, implying that detection becomes infeasible when the sensor is positioned too far from the fragments. Therefore, to ensure optimal performance, light sampling should be conducted from a distance no greater than 5 cm.

Plastic fragment sample arrangement analysis: Afterwards we quantified the number of plastic fragments found in various contexts in the wild captured by the TACO dataset. Figure 27a presents the results. It is clear that most of these fragments are nearly undetectable by current detection methods. Approximately 90% of the cases where micro or meso plastics are identified contain between 5 and 10 fragments, which are typically dispersed across a given area. In contrast, around 10% of cases feature a significant number of plastic fragments (more than 15 samples at once). Detection of areas with fewer plastic fragments is more challenging than those with larger quantities, as higher concentrations of micro and meso plastics can readily trigger cleaning or removal efforts. However, areas with lower fragment densities often go unnoticed, allowing pollutants to accumulate

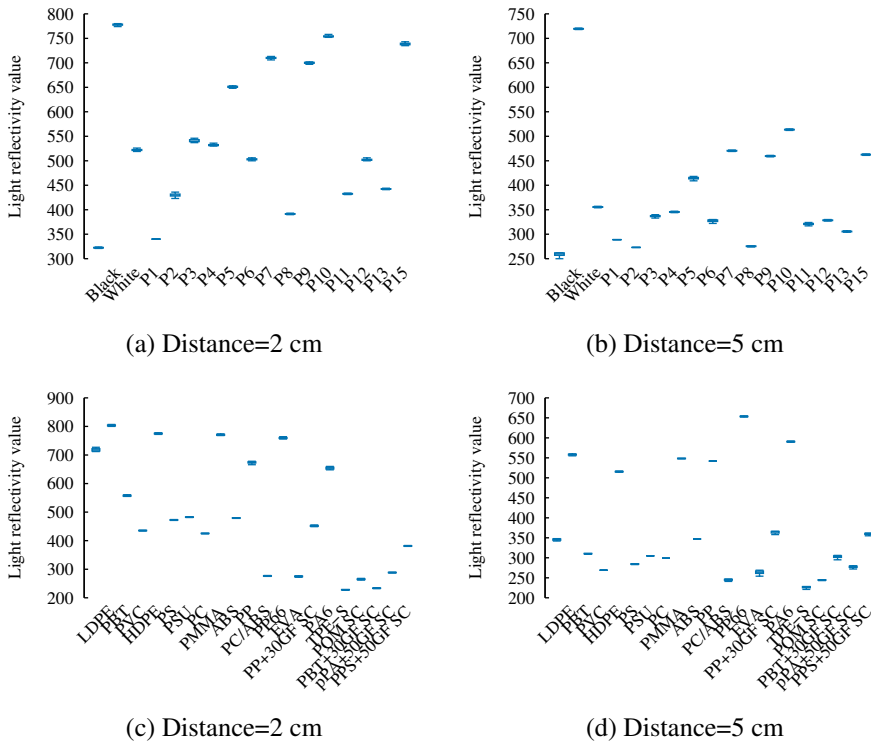


Figure 26: Light reflectivity values: (a),(b) plastic fragments of different sizes; (c),(d) 20 different plastic fragment samples.

gradually until they reach detectable levels over time.

Ratio analysis between plastic fragments and macro samples: Plastic fragments are seldom found in isolation; they are typically accompanied by macro-sized litter, which likely serves as the source as it degrades over time. To investigate this further, we analyzed the ratio between macro plastics and plastic fragments. Table 27b presents the results. The data shows that in 25% of the images from the dataset, the ratio of plastic fragments to macro litter is 2 : 1. The median ratio (50th percentile) is 3 : 1, while for the 75th percentile, the ratio is approximately 5 : 3. The maximum observed ratio of fragments to macro litter is 82 : 13. Based on this analysis, we categorized the images into three groups: Category 1 includes images with litter distributions below the 50th percentile, Category 2 encompasses images between the 50th and 75th percentiles, and Category 3 includes images above the 75th percentile. Further analysis revealed that 97% of the fragments are dispersed, while less than 3% can be classified as agglomerated. This indicates that detection methods like thermal imaging are ineffective for identifying most of the plastic fragments found in natural environments.

Individual modalities (baseline): Next, we analyzed the thermal dissipation times and light reflectivity values of the plastic objects used in our experiments.


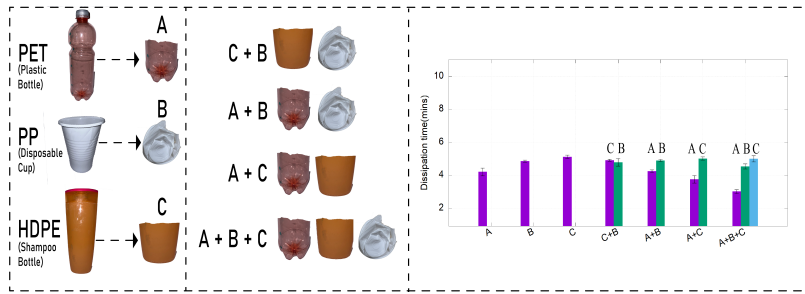



Figure 27: Litter samples distribution. (a) Amount of plastic fragment samples identifiable in the wild (TACO dataset). (b) Litter distribution in extracted images

Figure 28a presents the results of macro plastics detection through thermal imaging. The data indicates that macro plastics exhibit relatively consistent thermal dissipation times, even when broken into smaller pieces of varying sizes (A, B, C) as demonstrated in the figure. A Kolmogorov-Smirnov test comparing the thermal dissipation times between intact plastic products and their broken macro plastic parts (from the same products) confirmed that the thermal dissipation characteristics remained unchanged ($KS=0.141$, $p > 0.05$). The primary limitations of thermal imaging as a standalone solution are its limited resolution and the difficulty in inducing thermal radiation in small objects. Therefore, the LIZARD system combines thermal imaging with an additional modality to enhance the detection of smaller particles. Simultaneously, we derived light reflectivity profiles from the intact products by taking measurements from different points along their surfaces. Figure 28b shows the results, which demonstrates that light measurements for individual products fall within a narrow range, with minimal variation. This suggests that individual products can be reliably characterized using light reflectivity alone. The specific light reflectivity values recorded were: Plastic bottle (1066.96 ± 107.56), Shampoo bottle (1734.06 ± 202.97), and Disposable plastic cup (2440.33 ± 165.92). However, light sensing requires sampling across multiple locations and struggles to identify specific areas contaminated by small particles. These findings suggest that while both methods—thermal imaging and light-reflectivity can function independently, their combined use significantly enhances detection by focusing on specific areas for more accurate identification.

LIZARD performance: We assessed the performance of identifying macro plastics and plastic fragments using arrangements derived from real-world contexts. To conduct this analysis, we positioned our samples in locations based on the extracted synthetic contexts, enabling the application of our LIZARD sensing pipeline. Figure 29 presents the results for detecting plastic fragments, as macro plastics are easily detectable. Given that the background in which the samples are



(a)



(b)

Figure 28: Light reflectivity and thermal dissipation values of plastic samples (a) thermal dissipation times for macro plastics, (b) Light reflectivity from different sampling points

placed can influence the detection outcomes, we first examined the effect of using a consistent background across all selected arrangements (Category-1, Category-2, and Category-3). A black background was chosen as the baseline for this analysis. A Kruskal-Wallis test comparing the light reflectivity values between the baseline background and various plastic fragment types and arrangements demonstrated that plastic fragments can be effectively distinguished from background measurements ($\chi^2 = 270.68$, on average $\eta^2 = 0.90$, $p < 0.05$). In Category-1 (Figures 29(a-c)), 100% of the plastic fragments were successfully detected. Similarly, in Category-2 (Figures 29(d-f)), the detection rate for plastic fragments also reached 100%. Category-1 and Category-2 exhibit a lower distribution of plastic fragments (below the 75th percentile), and we mostly rely on Type-2 and Type-3 fragments. As a result, all fragments in these categories were successfully detected. In contrast, the analysis of the Category-3 arrangement (Figures 29(g-i)), which has a litter distribution above the 75th percentile, shows an overall detection accuracy of up to 89%. The reduction in accuracy is attributed to the increased presence of Type-1 plastics (less than 5 mm), which are more difficult to detect, suggesting that Type-1 plastic fragments can sometimes evade detection. Figure 30a summarizes the number of plastic samples detected by the LIZARD system compared to the actual number of

plastic fragments present in the selected arrangements. The results indicate that a single area mapping is insufficient to detect all types of plastic fragments, as it provides accurate results primarily for larger fragments. To reliably identify micro and meso plastics, multiple mappings of the area are necessary. Given that the LIZARD system already narrows the focus area to target micro and meso plastics, it is well-suited for repeated sampling without adding significant overhead to the process.

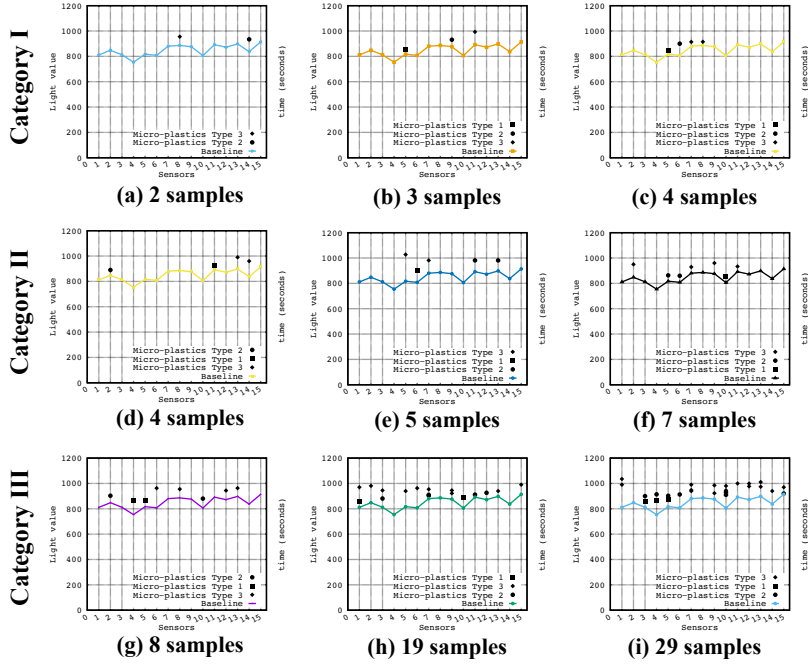


Figure 29: Light reflectivity for different categories.

LIZARD performance on more realistic background: In addition to utilizing a black generic background as a baseline for plastic fragment detection, we also assess the performance of our method using more realistic backgrounds typically found in public spaces. From our analysis of the TACO dataset, we identified the most common backgrounds where litter is present, which include concrete, sand, soil, wood, and grass. A bumpy background was also incorporated into the evaluation. The baseline light values of these backgrounds without plastic samples are presented in Figure 31. As anticipated, certain light backgrounds exhibit greater variation in light values compared to others, prompting us to include these backgrounds in our experiments. The results of light reflectivity values across different backgrounds are shown in Figure 32. A Kruskal-Wallis test performed between the light values of these backgrounds and various plastic fragment types reveals statistically significant differences ($\chi^2 = 89.5$ to 120.0 , $\eta^2 = 0.85$, $p < 0.05$), indicating that the light measurements of plastic fragments differ based on the background location. Thus, it is feasible to distinguish the

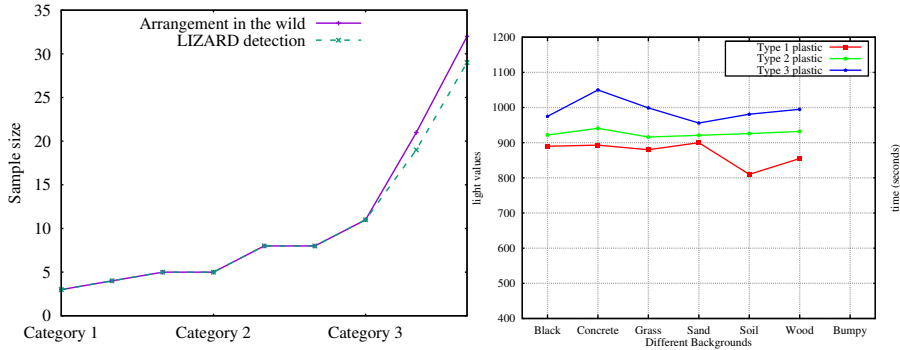


Figure 30: a) Comparison with ground truth b) Light values of same plastic fragments in different contexts

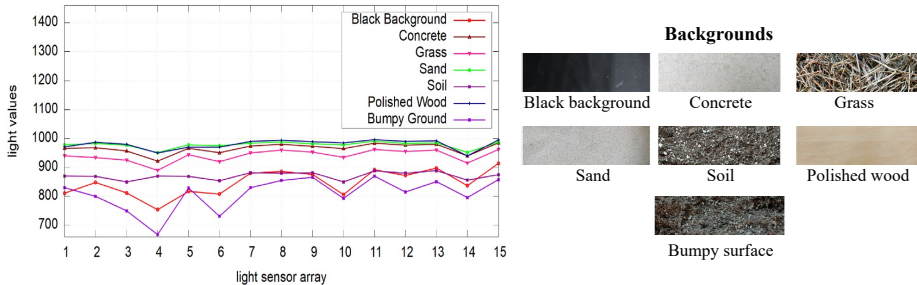


Figure 31: Baseline light reflectivity of different backgrounds

micro samples from their respective backgrounds. This is further demonstrated in Figure 30b, where we present the light values of specific plastic fragments in various contexts. The figure shows that the light values of plastic samples display minimal variation. The mean and standard deviation of light reflectivity values for Type 1 fragments across all backgrounds are 869 ± 31.54 , for Type 2, the values are 924.28 ± 9.81 , and for Type 3, values are 992.28 ± 29.24 . Additionally, Figure 32f illustrates how background deformations, such as bumpy ground, impact detection performance. Despite significant fluctuations in light values due to background irregularities, detection remains consistent. A Kruskal-Wallis test performed on the light values in bumpy backgrounds and different plastic fragments also indicates statistically significant differences ($\chi^2 = 118.5$, $\eta^2 = 0.883$, $p < 0.05$).

Classification performance: We first applied the k-means clustering algorithm to the light values of different plastic samples, setting the number of clusters to 3, to demonstrate that the light values corresponding to the three different types of plastic samples fall into distinct groups. The results, shown in Figure 33a, indicate that plastic fragment types are more easily detectable when using a black background. The centroids of the three clusters, represented by black circles, are well separated, with the centroid for Cluster 1 at 870, for Cluster 2 at 914, and for Cluster 3 at 970. However, when analyzing the impact of more realistic backgrounds, the

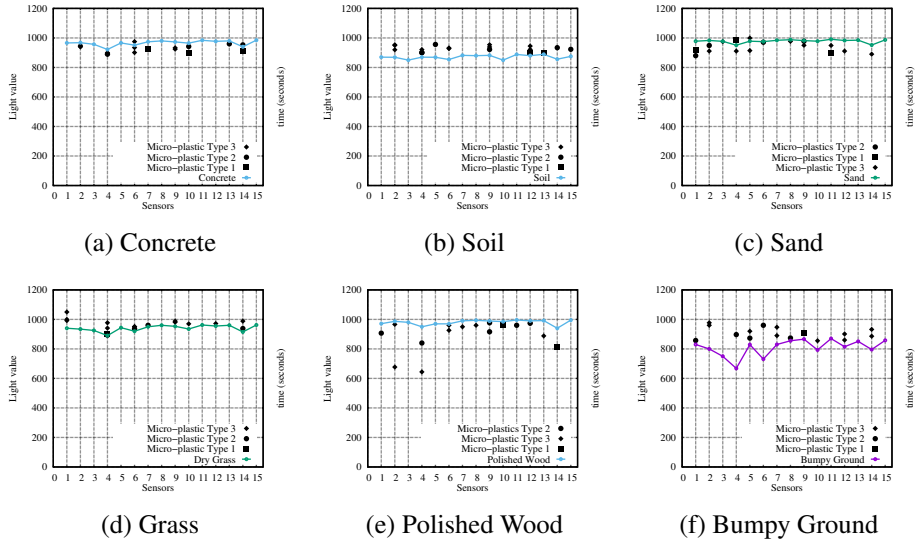


Figure 32: Light reflectivity values of plastic samples.

light values become more dispersed. This is particularly evident with lighter backgrounds such as polished wood or sand, where more light from the diode is reflected back to the photoresistor. In other words, the variation in light values depends on the background location, making the background an important factor in plastic sample classification. Based on this information, we then constructed a classical machine learning model. We preprocessed the categorical background data using One-Hot Encoding. When employing the RF classifier, we obtained a 10-fold cross-validation accuracy of approximately 77% to 80%. A similar score of around 70% to 75% was achieved using the SVM classifier. Figure 33b shows the confusion matrix of performance.

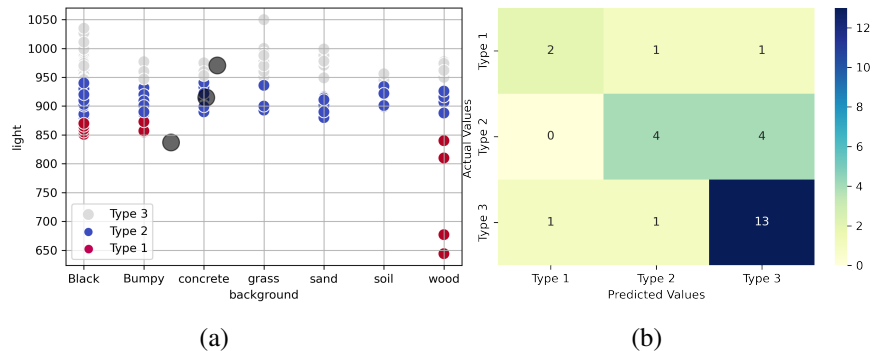


Figure 33: Clustering and classification results - (a) k-means clustering, (b) Confusion matrix for random forest classifier

Comparison to baseline: The performance of the LIZARD system is compared

against visual inspection, which is the most commonly used method for identifying and removing plastic fragments [246]. Other available methods often require bulky, expensive, and complex instruments, making them impractical for large-scale applications. In contrast, visual inspection can be applied in a variety of contexts. LIZARD, unlike traditional methods, requires minimal or no human intervention and can also enhance the accuracy of inspections conducted by the naked eye. To establish a baseline for comparison, we conducted a user study to quantify human performance in identifying plastic fragments. A testbed was designed to evaluate various locations (backgrounds), with Sand, Grass, and Rocky backgrounds selected to align with our primary method evaluation. The test used the Category-3 arrangement, which includes the largest number of plastic fragments (29) commonly found in natural environments. This arrangement covers a 20x14 cm area and includes micro (9), meso (13), and macro (7) plastics (see Figure 34a). We surveyed 45 participants (15 per background) from a university campus. Participants were randomly asked to perform the task while walking through the campus, and no personal data was collected. The task required them to remove all visible plastic fragments from the testbed, stopping when they believed the location was clear. A stopwatch and an Android application were used to record the timestamps between the removal of plastic fragments.

Comparison results: Figure 34b presents the results of the experiment. The data shows that at the beginning of the task, plastic fragments are removed quickly. On average, across all locations, 18 plastic fragments are removed at a rate of 20 seconds, before the time required for further removals increases exponentially. The average time taken for each location is as follows: Sand = 70.4 seconds, Rock = 81.04 seconds, and Grass = 113 seconds, indicating that the type of location influences visual performance. Regarding removal performance, 92% of the plastic fragments were removed from the sand, 86% from the rocky background, and 76% from the grass, underscoring the challenge of completely eliminating plastic fragments from the environment. When compared to the baseline, LIZARD demonstrates superior removal performance, achieving up to 98% effectiveness, and its identification speed surpasses that of human visual inspection.

4.7. Real-world practicability

Our experiments yielded promising results using real-world litter footage; however, they did not fully assess the system's performance in actual environmental conditions. In this section, we present LIZARD's robustness when tested in real-world environments.

Experimental setup: LIZARD was evaluated across three distinct locations, each featuring a different background: soil (park), gravel, and cast iron (manhole cover). Figures 35(a-c) show the deployment and Figures 35(d-f) provide a detailed view of these background locations. The experiment was carried out over a period of five days at various times (morning, afternoon, and night) to assess its performance

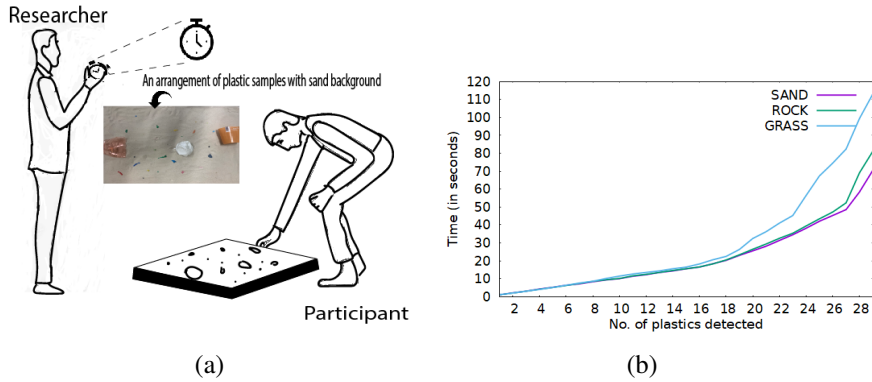


Figure 34: Removal of plastic fragments via visual inspection. (a) Illustration of the method, (b) Removal rate.




Figure 35: LIZARD deployment. Sampling locations - backgrounds: (a) & (d) rocky (b) & (e) iron, (c) & (f) soil.

under different lighting conditions.

Procedure: We selected an arrangement of 15 plastic samples, comprising the three types of plastics evaluated in previous tests: 4 Type 1 samples (microplastics < 5mm), 6 Type 2 samples (mesoplastics 5mm to 1.5cm), and 5 Type 3 samples (macro plastic fragments > 1.5cm). LIZARD utilizes light reflectivity to detect these different fragment types. Additionally, we collected ambient luminosity data using the LUX Meter Application for further analysis.

LIZARD performance: Figure 36(a-c) displays the luminosity values across different backgrounds as measured by the LUX Meter Application, while Figure 36(d-f) provide a characterization of the backgrounds where the plastic fragments were located. Each background presents a distinct profile, with gravel showing the highest variability due to its uneven surface. Additionally, the background characteristics change significantly based on luminosity, with morning and afternoon measurements differing substantially from those taken at night. Figure 37 illustrates the light intensity values for the various plastic fragments at different times

of day. These results align with our controlled evaluations, confirming that all plastic fragments can be detected, though there is a slight decline in performance due to the uneven surfaces and varying luminosity over time. The best detection results are achieved during the morning and afternoon, when ambient sunlight positively influences the sensors, whereas the worst performance is observed at night. It is worth noting that this issue could be mitigated by equipping the AGV with a separate light source. The primary errors occur between Type 2 and Type 3 fragments, with microplastics being easily distinguished. Most of the classification errors arise between mesoplastics and macroplastic fragments. To verify the consistency of our real-world data with controlled tests, we conducted statistical testing. The soil background exhibited the closest match between the two environments. We performed a Kolmogorov-Smirnov test to compare the measurements in the soil condition, confirming that the data is similar between the two experiments (KS=0.428, $p > 0.05$). Finally, we assessed LIZARD classification performance using the data collected from the light. Unlike controlled experiments conducted under constant luminosity conditions, the real-world data reflects varying light levels. By incorporating luminosity as an input feature, the overall model accuracy improved to approximately 85% (Figure 38).

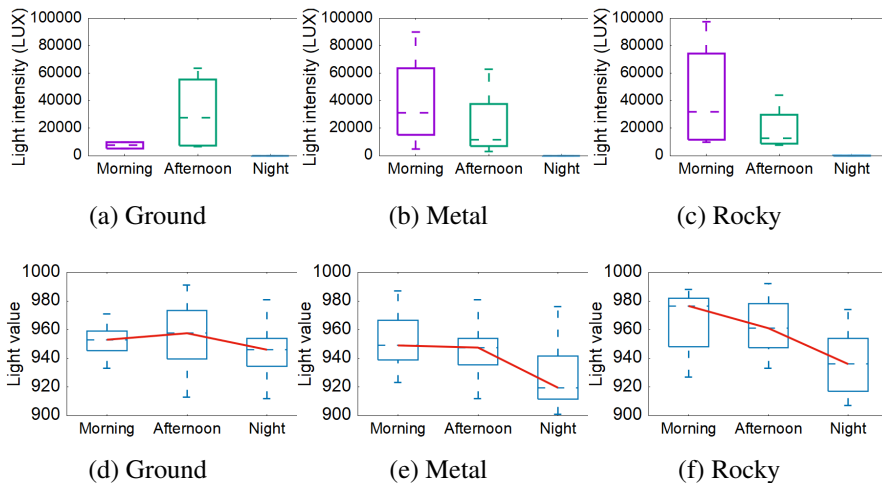


Figure 36: a)-c) Ambient luminosity during the day, d-f) Baseline background light values during the day.

4.8. Discussion

Stakeholders and Adoption: The experiments demonstrated that LIZARD is a promising solution for identifying macro, meso, and micro litter in the environment. This capability to obtain information about different litter particle sizes is the key feature setting LIZARD apart from existing solutions. Naturally, the accuracy of LIZARD is lower than with dedicated scientific instruments, such as

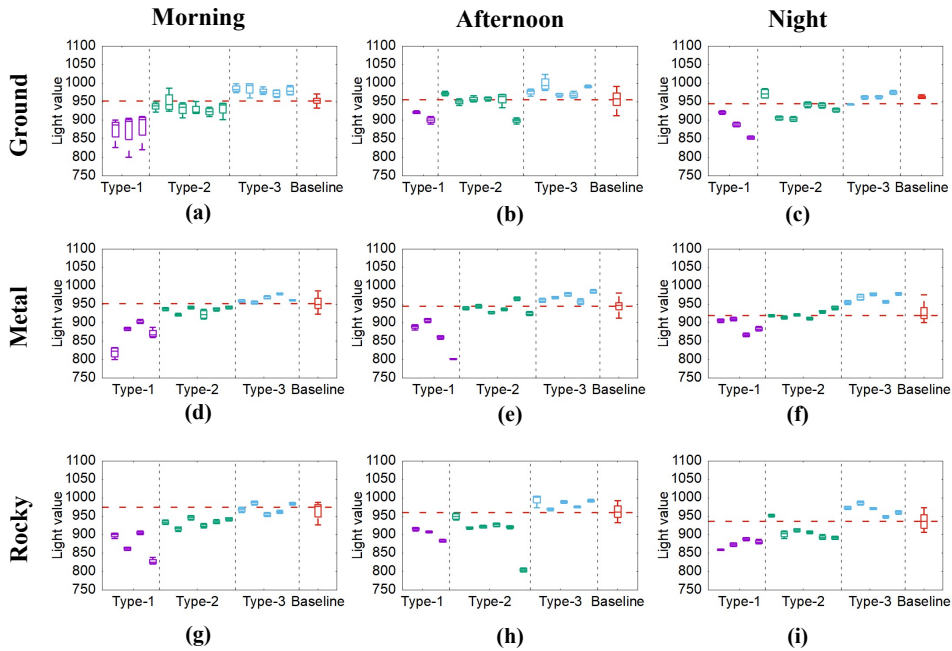


Figure 37: Plastic fragment detection in the Wild




Figure 38: Improvement on classification accuracy.

FTIR or Infrared microscopy, and thus the main use for LIZARD is to use it to autonomously identify areas that are badly affected and to help coordinate cleaning efforts. Another possible use would be to integrate it as part of drones that are operating within a city (e.g., delivery drones).

Context and sensor calibration: A limitation of our technique is background influence on micro-particle detection. Darker backgrounds (e.g., soil) ease detection as they reflect less light than lighter backgrounds (e.g., sand). Overall, our results were promising even with varying backgrounds. Highly reflective ground might benefit from more advanced sensor designs, such as using a small array of light sources to project a pattern for material estimation. Other options involve closer sensor placement or models considering background material. We are also interested in nighttime applications, leveraging consistent lighting, often from onboard AGVs, in low-light conditions.

Environmental Adaptability: Rain and snow can challenge micro-particle detection, but our solution remains reliable for macro particles, leveraging thermal imaging’s environmental insensitivity. While it offers insights into potential micro-particle locations, accuracy may be reduced. To address these, we can explore extendable covers to counteract weather effects and reflections. LIZARD uses thermal imaging for macro-plastic detection, acknowledging the impact of weather conditions. Shielding the thermal camera is a potential remedy. In aquatic or aerial settings, computer vision provides an alternative. For instance, LIZARD can identify plastic fragments in water, complementing established methods. While integration with underwater drones is feasible, it requires thoughtful design. However, implementing computer vision on resource-constrained devices demands extra resources.

On energy consumption: LIZARD adds sensors to AGVs with minimal operational impact. Unlike underwater and aerial vehicles, AGVs handle extra weight better, without navigation instability or increased energy use. In our experiments, LIZARD’s added weight doesn’t burden AGVs, avoiding extra energy consumption. AGV sensors include their own batteries, adding to the total weight. We opt for classical machine learning for classification, as deep learning, requiring ample data, proves resource-intensive for constrained devices [219].

AGV Integration: The experimental prototype combined off-the-shelf components and an affordable AGV. Integration challenges need to be addressed for larger-scale operations. The sensing units have their own power sources and don’t affect the AGV’s resource consumption. Plastic recognition uses a smartphone with a thermal camera and a sunshield, adding 350 grams. Additional light sensors and batteries add 50 grams. The total payload is 400 grams. This has indirect effects on energy [176] and may require re-calibration for navigation [291]. Future drones are expected to be modular. Currently, integration is difficult. Light sensors and a stepper motor are placed on the front of the AGV, requiring weight compensation at the back. Re-configurable components can overcome this issue.

4.9. Applications for Environmental Sustainability: Outcomes and Implications

The LIZARD system introduces a novel, autonomous approach to detecting and mapping plastic litter across a wide range of particle sizes—including macro, meso, and micro plastics—in real-world environmental contexts. By utilizing a compact combination of thermal imaging and light reflectance sensing, LIZARD can be seamlessly integrated with AGVs and, potentially, with city-operating delivery drones. This enables scalable, cost-effective monitoring across urban, terrestrial, and potentially aquatic environments, making large-area, high-frequency mapping of plastic pollution feasible for both citizen science initiatives and municipal stakeholders.

A primary application of LIZARD is its capability to autonomously identify and

localize regions afflicted by significant plastic litter accumulation. Unlike costly laboratory-based analysis systems such as FTIR or Infrared microscopy, LIZARD delivers in-situ, near-real-time insights, which can efficiently guide and prioritize networked or human-led clean-up interventions. Its ability to distinguish plastics of diverse sizes on various backgrounds—soil, sand, and urban surfaces – addresses the pressing need for field-deployable plastic pollution mapping, supporting rapid responsiveness to new pollution hotspots and adaptive management strategies. This supports sustainability goals by enabling resource-efficient deployment of remediation teams, reducing the persistence of plastics in natural environments, and underpinning targeted policy enforcement for pollution prevention.

The low payload and modest energy requirements of LIZARD mean it is readily adoptable on commercially available AGVs or drones without significant operational trade-offs, supporting broad deployment even in resource-limited regions [176]. By relying on classical machine learning for plastic classification, the system avoids the high hardware and energy overhead of deep learning approaches, which may be impractical for lightweight, battery-operated autonomous systems [51]. This ensures that plastic monitoring can be democratized and maintained for long durations without prohibitive costs or frequent maintenance.

LIZARD's sensing principles are also extendable. While current deployments focus on terrestrial plastics, the core methodology can be adapted – given suitable waterproofing and mechanical constraints—to aquatic environments, supplementing or partially automating laborious manual sampling and dye-based detection methods commonly used in water quality surveys. In addition, LIZARD can potentially be upgraded to monitor plastic nano-particles through integration with specialized filters and high-sensitivity light sensors, aligning with emerging concerns over nano-plastic contamination in both ecosystems and the food chain [102, 170].

From a systems and sustainability perspective, LIZARD's modular sensor array design facilitates easy upgrades and contextual calibration, ensuring detection accuracy even as underlying backgrounds, lighting, or weather conditions change. For example, by using arrays or multi-spectral (red, green, blue) sensors and incorporating weather shields, LIZARD can maintain reliability in both dry and moderately adverse environmental conditions. Future research will enhance its robustness through multimodal sensing—combining thermal, visual, and potentially chemical modalities—to minimize errors due to occlusions or environmental interferences [224, 349].

LIZARD's targeted AGV integration ensures energy efficiency remains high. Studies show that payload additions up to approximately 400g—such as those associated with LIZARD's sensor package—introduce negligible impact on AGV operation and battery longevity in most terrestrial vehicles, allowing for extensive, routine monitoring without frequent recharging or significant mission planning changes [176, 153, 320]. The use of sensor arrays and stepper motors for detailed area sampling supports improved spatial coverage and finer mapping, which is crucial for monitoring plastic pollution trends over time.

In summary, LIZARD enables scalable, high-resolution, and cost-effective environmental plastic monitoring, significantly enhancing the operational scope of clean-up efforts, prevention strategies, and public policy. By making plastic detection accessible to a broader set of stakeholders—including cities, environmental agencies, researchers, and even citizen science communities—it contributes directly to the reduction of environmental plastic loads, the improvement of ecosystem resilience, and the fulfillment of sustainability targets such as those outlined in Sustainable Development Goals (SDG) 14 (Life Below Water) [18] and SDG 15 (Life on Land) [253].

4.10. Summary

LIZARD presents an innovative pervasive sensing approach for detecting plastic fragments in environmental settings through dual-modality sensing combining thermal imaging and light reflectivity. Designed for integration with autonomous ground vehicles (AGVs), LIZARD enables continuous monitoring of public spaces and natural ecosystems without manual intervention. The system addresses the significant challenge of detecting small-scale plastic pollution that is difficult to sample using conventional methods. Through comprehensive evaluation across diverse realistic environments, LIZARD demonstrates up to 80% detection accuracy, representing the first autonomous technology capable of field-deployable plastic fragment identification.

LIZARD significantly advances environmental sustainability by enabling proactive, large-scale plastic pollution monitoring that was previously impossible with manual methods. The system's autonomous operation reduces the human resources required for environmental surveys while dramatically increasing coverage and frequency of monitoring efforts. By facilitating early detection of plastic accumulation in ecosystems, LIZARD supports timely intervention and cleanup activities, preventing further environmental degradation. The technology's integration with AGVs creates scalable monitoring networks that can operate continuously across urban and natural environments, contributing to more effective pollution control strategies, enhanced ecosystem protection, and data-driven environmental policy development.

5. MICRO-CLOUD: FOG COMPUTING FRAMEWORK FOR UNDERWATER PERVASIVE DATA SCIENCE

This chapter introduces the concept of *Micro-clouds* – a fog computing framework for underwater environments – offering a decentralized, modular, and cost-effective approach to enabling real-time data processing at or near the source of aquatic data collection. By leveraging micro-clouds built from COTS hardware, the framework addresses key challenges in underwater data science, including limited communication bandwidth, high latency, and the lack of computational resources in remote or infrastructure-poor marine and freshwater settings. Through extensive experimental validation with representative input data for object detection, pollution monitoring, and biodiversity assessment, the chapter demonstrates that micro-clouds can deliver agile, energy-efficient, and scalable analysis of underwater sensor data. Moreover, they enable broader, more inclusive monitoring capabilities across vast and previously inaccessible regions. Ultimately, this advancement supports environmental sustainability by enabling faster responses to ecological threats, fostering data-driven marine management, and promoting proactive protection of aquatic ecosystems in alignment with global environmental goals and sustainable development practices.

5.1. Introduction

Underwater environments are increasingly becoming as a new frontier in data science. A variety of underwater sensors—such as hydrophones [221] video cameras [222], and those measuring salinity, pH, or other water parameters are being deployed to support a range of applications. These applications include oil pipeline monitoring [131, 7], fishery management [267, 231], reef and fish school assessments [128, 42], and harbor safety monitoring [187]. However, operating computational systems underwater poses challenges, and the lack of high-speed communication technologies restricts the available computing infrastructure. Consequently, significant delays often occur between data collection and analysis, limiting the potential scale and impact of these applications [91, 224]. To address these limitations, it is crucial to provide computing resources close to data sources. Such proximity enables timely insights and broadens the scope for real-time, context-aware underwater applications. These applications range from monitoring underwater pollution [59] and forecasting litter dispersion [212] to analyzing pollutions impact on marine life in real time [264]. Furthermore, local computing resources can simplify communication infrastructure by enabling direct device-to-device connectivity underwater, reducing reliance on surface-based or cloud-based communication. The primary method for enhancing the computational capabilities of underwater sensors currently involves utilizing surface-based infrastructure, such as ships or buoys, which provide computing resources or serve as

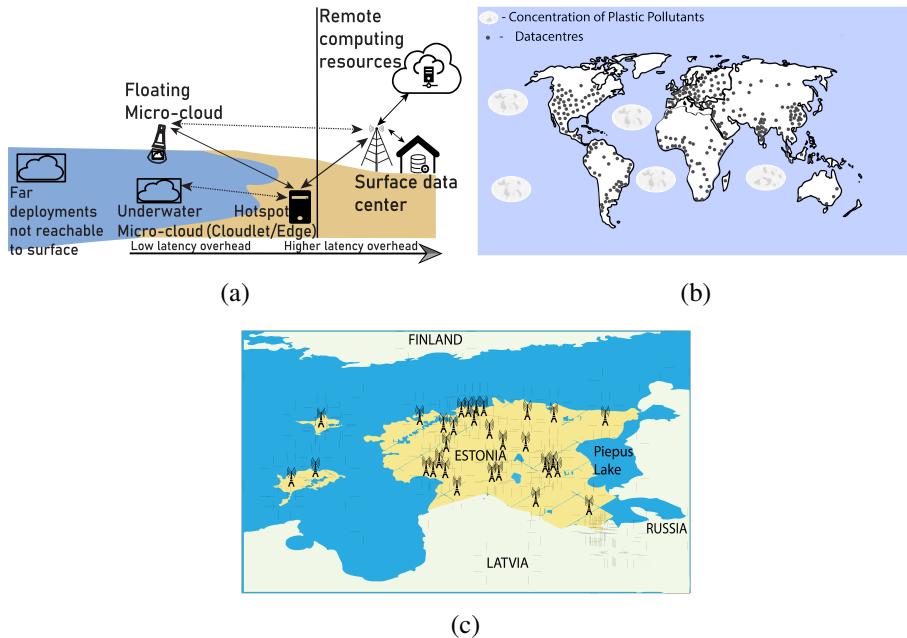


Figure 39: Complexity of moving underwater collected data to computing infrastructure. (a) Surface-based deployments can only support underwater infrastructure that is close to coastal areas, (b) Major areas of marine litter are far from data centers and also in areas that are far from coastal regions, (c) Inland regions, such as lakes, rivers and streams, also can be poorly covered by communication infrastructure.

gateways to land-based systems. This approach is limited by its support for shallow depths and short distances from land-based infrastructure. Additionally, it relies on specialized communication interfaces, such as laser-based optical communication [275, 47, 263] to transmit data from underwater sensors to the surface. The use of surface-based infrastructure for relaying data to remote systems is feasible only in coastal regions near populated areas, as high-speed, high-bandwidth communication requires base stations to be located within approximately 10 km, or at most a few tens of kilometers, of the gateways (Figure 39a). As a result, this approach is unsuitable for most areas of interest, such as regions with high concentrations of marine litter that are far from computing and networking infrastructure [Fig. 39b]. Similarly, inland water bodies like rivers, lakes, and streams often lack access to appropriate computing resources (Fig. 39c). Although recent initiatives aim to bring data centers closer to marine environments [61], these efforts primarily target coastal areas near densely populated urban regions, making them unlikely to serve as a broadly adopted solution for underwater data processing.

It is important to acknowledge that the foundational paradigm of the micro-cloud as a decentralized edge computing architecture was first introduced by Elk et

al. [79]. Building upon this established paradigm, we propose a *fog computing*¹ solution for underwater data science that leverages submerged COTS devices to provide cost-effective and decentralized access to computing and storage resources. While the Micro-cloud is fundamentally designed as a fog framework, its modular nature allows it to flexibly act as an edge device or engage in cooperative processing depending on the specific physical constraints of the underwater deployment. To demonstrate this concept, we developed a proof-of-concept offloading framework and two prototypes utilizing COTS microcontrollers, such as Raspberry Pi, as fog nodes to support underwater applications. By integrating these inexpensive, compact, and energy-efficient COTS components, fog computing capabilities can be seamlessly embedded into AUVs, buoys, ships, and other underwater infrastructure. Our solution enables on-demand processing support and facilitates scaling underwater data science. We adopt a micro-cloud architecture, where multiple devices collaborate to deliver fog computing capabilities. This approach is particularly well-suited for resource-constrained devices and offers a cost-effective implementation [162, 79]. Additionally, our method utilizes standard communication interfaces; for instance, we demonstrate that Wi-Fi interfaces can facilitate interactions with the micro-cloud, provided the devices are within a sufficiently short range. Notably, our solution achieves networking performance comparable to established techniques like underwater LoRa [174], while also performing underwater computations. Communication range and bandwidth can be further enhanced using advanced interfaces, such as optical communication systems, which can achieve data transmission rates of 10 Mbps over distances of up to 40m [145].

We demonstrated the feasibility and practical advantages of our solution through extensive benchmarking, which includes both surface-based and underwater computing tasks under various water conditions. These tests even extend to deployments on the seabed during a recreational scuba dive. Our evaluation focuses on the solution’s ability to support diverse applications, using tasks that reflect the specific requirements of underwater data science while also aligning with established fog computing benchmarks [188]. Following these experiments, we conduct additional tests in underwater environments, emphasizing object detection from camera footage. This task is widely used in underwater data science and holds significant relevance for applications such as pollution detection, biodiversity estimation, and pipe leakage detection [224]. The results of our experiments demonstrate that micro-clouds can provide general-purpose support for a wide range of underwater computing tasks and are capable of operating effectively in complex underwater environments. Standard communication interfaces, such as Wi-Fi, are sufficient to maintain connectivity among the devices forming the micro-cloud, as long as they are within the same container. Clients can communicate with the micro-cloud as long as they are within a short distance of the container (i.e., a few centimeters),

¹We define fog computing as decentralized support for computing and data processing that is offered close to the source of data. This definition is adapted from [332] and aligns with the original definition of fog computing by Cisco.

although specialized underwater communication technologies can be used to extend the range for client requests. Despite the limited range, these findings have significant practical implications, suggesting that even simple COTS underwater drones without dedicated communication interfaces can support underwater data science. By attaching a separate fog container that uses COTS technologies, these drones can avoid the need for complex communication interfaces. In terms of computational performance, we find that COTS devices, such as Raspberry Pi, offer sufficient computational power for most underwater data science tasks. However, performance slightly declines as the depth of the container increases. The most significant drop occurs at shallow depths due to heat accumulation inside the container, with performance stabilizing at deeper depths. We also demonstrated that optical communication systems are a promising option for extending the communication range, although the stability of connectivity is influenced by the calmness of the water. Lastly, we showed that accelerometer-based motion analysis can optimize micro-cloud performance by regulating resource usage and identifying conditions that enhance the likelihood of successful communications.

5.2. Related Work

Our work draws inspiration from edge and fog computing research and from underwater IoT and data science. We next review relevant works in these fields to highlight key requirements and challenges for providing computing support. Table 9 provides a summary of this comparison, providing the proposed name of the solution (Proposed); the type of deployment supported by the solution (Deployment); the type of hardware used as underlying processing infrastructure per (fog) processing unit (Underlying hardware); the available communications that the deployment provides (Available communications); the flexibility of the solution to be used underwater (Ready for underwater deployment) and whether the solution exploits COTS components to provide high replication rate and large-scale adoption (COTS components). The key novelties of our work are solutions to several practical challenges resulting from the underwater environment, and the provisioning of an underwater computing infrastructure (micro-cloud) that is low-cost and easy to implement. Achieving this is critical for ensuring that computing support can be easily deployed and used to support diverse underwater (pervasive) data science applications. Indeed, unlike our work, existing solutions are limited to offering access to external, surface-based infrastructure, or augmenting the capabilities of individual devices, without offering a general purpose platform that can simultaneously support multiple devices and serve the needs of a broad range of underwater (pervasive) data science applications.

Table 9: Summary of most relevant work for fog provisioning. Our work is the first to offer fog computing support consisting of multiple devices that is ready for underwater deployment and that uses low-cost COTS components.

Proposed	Deployment	Underlying hardware (per node)	Available communications	Ready for underwater deployment	Off-the-shelf components
VM-Based Cloudlets [251]	edge/fog	individual	Wireless LAN	no	no
OREO [323]	edge	multiple	Wireless LAN	no	no.
Fog micro datacenter [1]	datacenter	multiple	Wireless LAN	no	no
Pocket Cloudlets [154]	mobile-device	multiple	Wi-Fi/Cellular	no	yes
Pervasive Data Science [163]	edge	multiple	Wireless LAN	no	yes
Smartphone cluster [40], MISCO [74]	mobile devices	multiple	Wi-Fi	no	yes
FemtoClouds [111]	edge	multiple	Wi-Fi	no	yes
Collaborative processing methods [172, 290, 306]	edge	multiple	Wi-Fi	no	yes
Cloudrone [249]	edge	individual	Wi-Fi	no	yes
Geographical relocation methods [294], [14], [57]	data center	multiple	co-axial/fiber optic	yes	no
Natick [61], [214]	data center/ cloud	multiple	co-axial/fiber optic	yes	no
IoUT [146], [73], [141]	wireless sensor networks	multiple	Wi-Fi	yes	no
Underwater exploration and monitoring [7], [267], [271], [119]	wireless sensor networks/AUVs	individual	Wi-Fi	yes	no
Underwater infrastructure [131], Aqua-Fi [263]	WSN, drones	individual	Wi-Fi/optical	yes	no
Penguin [93]	micro-cloud/AUV	individual	Wi-Fi	yes	yes
POSEIDON [221], [43], [135]	edge	individual	acoustic/wired	yes	yes
Deep learning in Oceans [224]	AUV/ROV	individual	Wi-Fi	yes	yes
Our work	fog	multiple	Wi-Fi	yes	yes

Computational Augmentation: Edge and fog computing are the main paradigms for augmenting computing by offering processing, intelligence, storage, and other functionality close to the data source [332]. Edge computing provides services that are in the vicinity of the data sources, such that there are no oscillating changes in communication latency that hamper the battery and performance of applications. For example, improving battery saving by caching data on the edge has been investigated [323]. A key limitation with edge computing is the lack of dense and ubiquitous deployments to provide continuous support to end-applications. Fog computing, in turn, assumes the support covers data storage and is able to integrate intelligence [332]. Fog computing can be delivered from any device with enough processing resources that is blended within the environment. For example, common devices acting as fog nodes, include, access points, IoT devices, cloudlets and edge servers. Our work explores a new frontier for fog computing, developing and deploying micro-clouds in underwater environments to increase the ubiquity of access to processing resources for underwater applications. Existing works cannot be directly adopted in underwater environments as there are unique challenges when operating underwater. Our work addresses some of the key challenges, including water motion, poor wireless propagation, and the need for sufficient waterproofing, and highlights their impacts on delivering computational support. We also demonstrate how off-the-shelf devices can be harnessed for underwater

needs. Our work serves to pave the way for real-time data analytics in underwater environments.

Cloudlets and Micro-clouds: Cloudlets provide computing power close to users [251] and are the foundations of edge and fog computing paradigms. A micro-cloud, in turn, depicts an extended form of a cloudlet, whose underlying resources are aggregated using multiple distributed devices [79, 162]. Since smart and IoT devices have increased computational capabilities, approaches to create dynamic micro-data centres with them have been proposed [154, 1, 163]. It has been demonstrated that a rack of smartphones can be used to create a cloud computing-like infrastructure [40, 74]. Moreover, collaborative processing and federated learning between devices can be used to create dynamic and elastic computing infrastructures on the edge [111, 172, 290, 306]. Also, micro-clouds can operate on aerial drones at the edge [249]. Our work draws inspiration from the possibilities offered by micro-clouds, addressing key challenges and developing the necessary support to deploy and operate them in an underwater environment, as well as demonstrating the benefits micro-clouds can bring to underwater data science.

Data Centre Deployments: Among numerous ecological challenges, reduction of data centres emissions is an issue that has been investigated widely to overcome the impact on climate change [20, 294]. Data centres have been moved to different geographical locations, in order to improve cooling and provisioning of services to end-users [14, 57]. While dunking the data centers has been explored by Microsoft in its Naptik [61], underwater data centres can take advantage of low temperature sea floor and reduce cooling power, improving energy efficiency [301]. Optimal deployment of marine cloud computing have been investigated [214]. These all suggest there is significant potential in underwater computing infrastructure that warrants further investigation. Our work explores the design of (small size) micro-clouds that can be deployed underwater near where the computations are performed. These micro-clouds can provide localized computing support for underwater data science applications, complementing the deployment of dunked data centres and providing broader coverage of computing resources over targeted areas.

Internet of Underwater Things: IoT has increasingly large scope in underwater scenarios, and many of the application scenarios have been covered in previous works [224, 146, 73, 141]. Most of the existing work has focused on developing applications for specific purposes, including marine pollution monitoring [7], aquaculture [267, 271] and study of marine life [141, 62, 22]. Other work has focused on adopting existing technologies and developing new ones for underwater usage [146]. Also, underwater sensors networks have been studied in detail [119, 87] and several other technologies have been integrated into them, e.g., drones [131], and floating infrastructure [263]. In particular, different communication technologies have been explored [47, 247, 222, 223]. Despite the increasing amount of research, very little work exists on augmenting the processing resources of Internet of Underwater Things (IoUT) applications. Other works have attempted to facil-

itate data gathering from oceans using Underwater Sensor Networks (UWSNs) by proposing efficient routing protocols [344]. A programmable IoUT project called SEANet has also been developed to make it convenient to add , remove and replace both hardware and software [68]. Other works have focused on making the IoUTs scalable [199]. Our work addresses the gap in the availability of processing resources in underwater environments, developing a practical solution that addresses challenges in operating external computing infrastructure in an underwater environment and that is beneficial for augmenting the processing capabilities of these types of IoT applications and deployments.

Underwater Data Science: The most common method to analyze underwater data is to use passive analysis, where data is collected from underwater and then taken out to be processed by surface-based infrastructure [63]. Large amounts of data are collected underwater that require on-site data analysis, e.g., deep learning [224]. Multiple tools have been developed to support deep learning applications for underwater data analysis [43, 135, 221], e.g., whales with CurvRank [41], dolphins with finFindR [286], NNPool [181] and PhotoID Ninja [39], and turtles with MY-DAS [45], among others. Despite these solutions targeting analysis of underwater data, they do not operate underwater and require surface based infrastructure. In parallel to this, autonomous vehicles have been equipped with augmented computing payloads to perform processing underwater using deep learning [93, 224]. Autonomous vehicles have limited operational time underwater and its difficult to extend their design for augmentation of computing resources [176]. Thus, alternative deployment to obtain additional computing support is required. Our work provides one such alternative, developing micro-clouds that can operate underwater and demonstrating how they can serve the localized processing needs of underwater applications. This makes it possible for applications to augment their processing capabilities without relying on surface-based deployments.

Summary of Literature Review: Underwater environments remain highly challenging for computing and, currently most underwater data science applications rely on scenarios where either surface-based computing support is reachable (e.g., buoys close to access point stations or tethered drones) or there is a significant delay between the collection and analysis of data [224]. Our work addresses technical challenges in developing a general-purpose solution for improving access to computing resources, offering an easy-to-implement and low-cost solution for delivering computing support to underwater applications. Our work is the first to provide such a platform and to solve technical challenges in deploying and operating the platform underwater. At the same time, our work is firmly grounded on the current state-of-the-art in fog and edge computing, including the delivery of micro-clouds, but extending these solutions to highly challenging underwater environments. While fog and edge computing technologies have been explored previously, translating existing solutions from surface-based infrastructure to underwater computing support is non-trivial and requires addressing challenges resulting

from this shift in the environment. We address these challenges, including limited wireless connectivity, water motion, and need for water-proofing, developing solutions and analyzing their effects on the computing support that can be offered.

5.3. Feasibility Experiment

Our work addresses the need for general-purpose solutions capable of supporting the processing requirements of a broad range of underwater applications. Micro-clouds composed of commodity devices offer a promising solution, as they are cost-effective, widely available, and capable of handling common processing tasks [163]. The affordability of the components mitigates the economic impact of potential equipment failures and facilitates denser and more rapid deployments. Unlike traditional fog computing scenarios, underwater applications can benefit significantly even from low-end devices, as the primary requirement is the availability of dedicated computational support. This is particularly relevant because underwater platforms generally possess minimal computational capacity, as their primary functions focus on navigation and maneuvering. Simultaneously, reliable computational support remains indispensable in such environments. Integrating computational support into these devices presents challenges due to the need for waterproofing both hardware and software components [224]. We propose deploying micro-clouds either as independent submerged components or as modular attachments to devices operating in underwater environments. For instance, a micro-cloud could be housed within a separate container mounted on an AGVs tasked with collecting and analyzing underwater measurements [93]. Alternatively, it could function as part of buoys or other observation platforms [296].

Feasibility Experiment: To evaluate the feasibility of deploying functional micro-clouds, we conducted controlled benchmark experiments to assess the impact of submersion on their computational performance and resource utilization. A micro-cloud was constructed using a Raspberry Pi 4 (RPi4) micro-computer contained within a waterproof glass container (refer to Section 5.5 for details of the experimental setup). The RPi4 was chosen due to its widespread use as a rapid prototyping platform for IoT applications, despite its default unsuitability for underwater environments. While a single RPi4 offers limited computational resources, it is inadequate for large-scale applications. However, collaborative processing can be employed to enhance computational capabilities by interconnecting multiple devices [144]. In practical fog computing scenarios, sufficient computational power is required to process data captured from the surrounding environment. Common data types include images, videos, and environmental parameters such as salinity, temperature, and pH. The RPi4 is capable of handling such data and can even perform deep learning-based object detection on images [224]. Consequently, the initial focus of our experiments was on benchmarking the individual processing capabilities of the RPi4. We developed a lightweight fog service on the RPi4, designed to operate within a client-server architecture and provide computational

resources upon request. After submerging the micro-cloud in shallow water (a few centimeters as shown in Figure 42c), we analyzed the effects of water on connectivity and computational performance. These initial benchmarks were conducted at shallow depths to mitigate the risk of equipment loss. In subsequent sections of this chapter (see Section 5.8), we explore the performance of these devices in more realistic operating conditions, where they are submerged to depths of several meters.

Experimental task and setup: For the feasibility benchmarks, we selected a computing task that ensures consistent and uniform response times while incorporating resource-intensive processing. By controlling these parameters, we aimed to minimize the influence of non-deterministic computing behaviors on the experiment. The chosen task involves a primality test and search, specifically identifying all prime numbers within a given list of integers. In the experiment, a client sends a request to the micro-cloud-based fog service, providing a list of 20 integer numbers within a specified range of 100000 – 105000. The service processes the request by identifying the prime numbers in the list and returns the results to the client. To simulate multiple clients submitting requests, we utilized JMeter² to generate various user workloads. We progressively increased the workload from 100 to 500 users to assess the micro-cloud’s capacity to handle varying demands. Additionally, we simulated different connectivity conditions by adjusting the depth at which the micro-cloud is submerged. For communication between the devices in these controlled experiments, we used the standard Wi-Fi interface.

Results: Figure 40 presents the performance outcomes of the micro-cloud when handling different workloads from concurrent users. To evaluate performance, we estimate the round-trip time (RTT) based on successful request completions. We assess the workload handling performance both on the surface (baseline) and underwater, using the same experimental setup. Figure 40a and d show the results under baseline conditions. Subsequently, we analyze the impact of submersion on the ability to handle different computational workloads. Figures 40b, c, e and f illustrate the performance results when the micro-cloud operates underwater. From these results, we observe that the micro-cloud can manage workloads while submerged, with minimal overhead compared to the baseline (Figures 40b and e). We found that when the micro-cloud is submerged at depths between 1 and 5 cm from the surface (referred to as Depth-1 in the figure), it can successfully complete all workloads. At depths between 6 and 12 cm (Depth-2 in figure 40c and f), the transmissions begin to experience reliability issues, and the micro-cloud drops some requests. To quantify this signal attenuation, we continuously monitored the Received Signal Strength Indicator (RSSI) and packet transmission metrics. While the surface baseline maintained a strong RSSI of approximately -10 dBm with a near-zero packet loss rate, the RSSI at Depth-2 degraded significantly to -45 dBm. This severe attenuation resulted in a packet re-transmission rate exceeding

²<https://jmeter.apache.org/>

18%. The loss of connectivity negatively impacts performance primarily due to this increased overhead caused by failed transmission attempts and continuous packet re-routing. The loss of connectivity negatively impacts performance, primarily due to increased overhead caused by packet loss and the limitations of the Wi-Fi interface. As Wi-Fi is known to suffer from poor underwater propagation [60], the client and micro-cloud must be positioned within a few centimeters of each other to ensure reliable communication. It is important to note that this connectivity issue pertains only to communication between the client and the micro-cloud. Internally, the devices forming the micro-cloud can communicate using Wi-Fi or other standard communication interfaces if they are housed within the same container. However, integrating all components into the same container can lead to heat accumulation, which may become problematic at shallow depths (around 1 meter or less). At deeper depths, the effects of sunlight are reduced, the water temperature is cooler, and the increased water pressure outside the container assists with cooling. To extend the communication range, alternative solutions such as acoustic or optical communication technologies, which can operate at depths of several meters, could be explored [263, 47]. One such solution is discussed in Section 5.7.

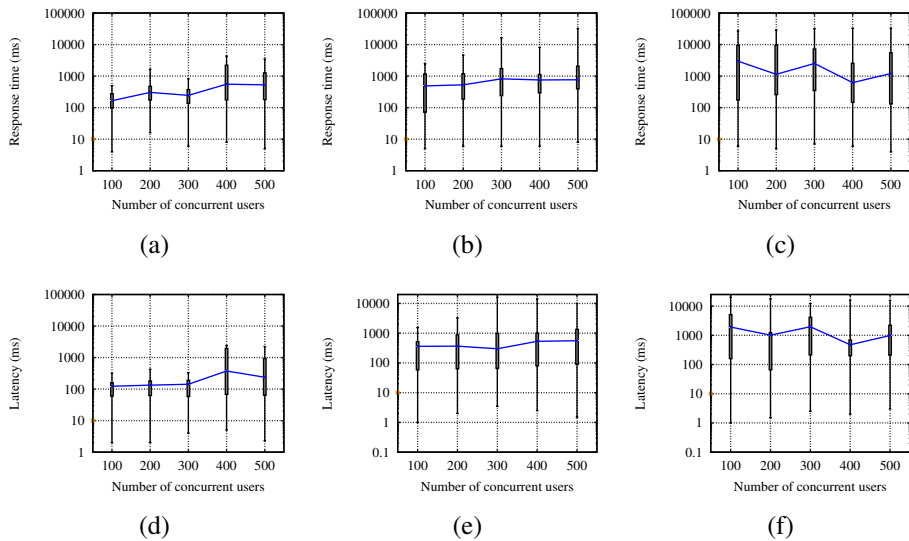


Figure 40: Capacity results of submerged micro-clouds when handling multiple users, a) Response time (Baseline - no water), b) Depth-1, c) Depth-2, d) Latency (Baseline - no water), e) Depth-1, f) Depth-2.

Finally, we assess the impact of sudden water motion on the submerged micro-cloud. Our analysis focuses on the first depth class, where the micro-cloud remains accessible without dropping requests. To investigate the effects of water motion—such as waves, currents, and tides—on access to the micro-cloud’s computational resources, we place an accelerometer sensor in a glass container on the

water’s surface, with the micro-cloud submerged. This setup allows us to capture surface water motion while the micro-cloud handles user workloads. Figure 41 illustrates the results, which highlight different types of water motion. Specifically, Figure 41a and 41b show conditions with minimal water movement. Under these circumstances, the micro-cloud can successfully complete user workloads without issues. In contrast, Figure 41c and 41d depict the effects of sudden disturbances on the water surface caused by nearby water vehicles. These disturbances result in significant motion, which destabilizes connectivity with the micro-cloud. When such high motion occurs, connectivity to the micro-cloud becomes unreliable, emphasizing the need to identify optimal transmission conditions to mitigate these challenges.

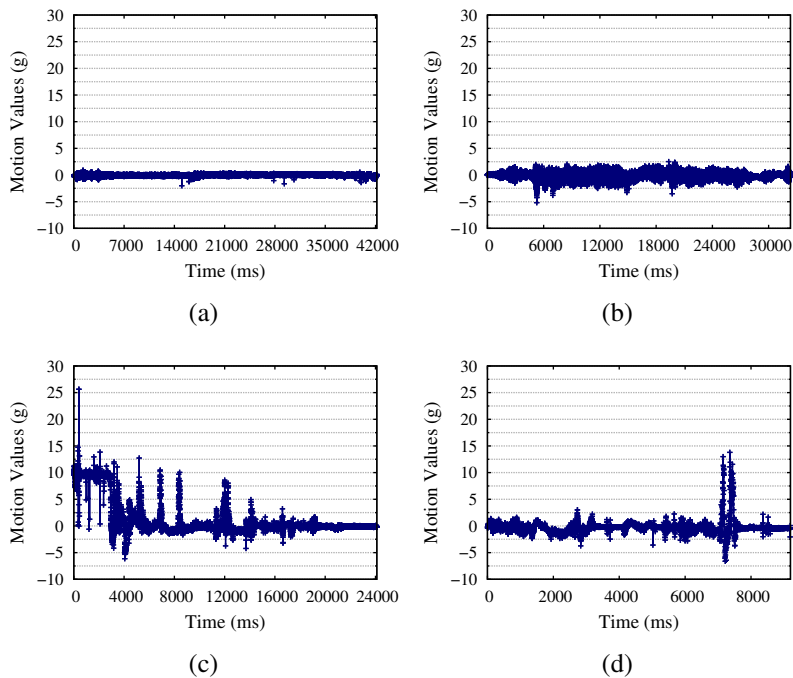


Figure 41: Different types of water motion influence connectivity to the submerged micro-cloud, a-b) Low, c-d) High (induced by water vehicles operating nearby).

5.4. Underwater Micro-Cloud Design

We subsequently present the design of submersible micro-clouds (cloudlets) engineered to detect water stability and enhance transmission reliability. Figure 42 illustrates the comprehensive prototype design of the micro-clouds, detailing their internal components and their deployment in real-world environments using waterproof enclosures for experimental testing and evaluation.




Figure 42: Micro-cloud prototype: (a) Internal components, (b) Waterproof encasing, and (c) Deployment in the wild.

5.4.1. Architecture

Figure 43a illustrates the overall architecture of a micro-cloud, comprising modules that integrate functionalities for computing, sensing, energy monitoring, and other essential operations. A modular design was adopted to ensure extensibility and adaptability. The micro-clouds utilize off-the-shelf components that are portable, compact, and lightweight. These components offer flexibility for rapid prototyping and easy replacement, facilitating integration with other devices and infrastructure, as demonstrated in Figure 43. For example, micro-clouds can be attached to underwater drones or existing aquatic infrastructure, such as buoys, to enable pollutant monitoring (Figure 43b). Additionally, micro-clouds can collaborate to form high-computing infrastructures through cooperative processing and analysis (Figure 43c). They can also be deployed at specific locations to support nearby devices (Figure 43d). By maintaining the micro-cloud infrastructure within close proximity to underwater IoT devices, the computational burden on these resource-constrained devices is significantly reduced. This optimization extends the operational duration of underwater IoT devices and supports the implementation of advanced fog analytics techniques, thereby enhancing exploration capabilities.

5.4.2. Components

Computing: IoT devices operating underwater face significant limitations in computing resources, restricting their ability to perform resource-intensive analyses, such as those involving machine learning or deep learning algorithms. To address this challenge, submersible micro-clouds serve as auxiliary computing units, augmenting the computational capacity of these devices. These micro-clouds incorporate separate microcontrollers or smartphone devices as processing units. Due to their compact size, lightweight design, and portability, micro-clouds can be easily assembled, waterproofed, and attached to underwater equipment, such as remotely operated vehicles (ROVs) or AGVs.

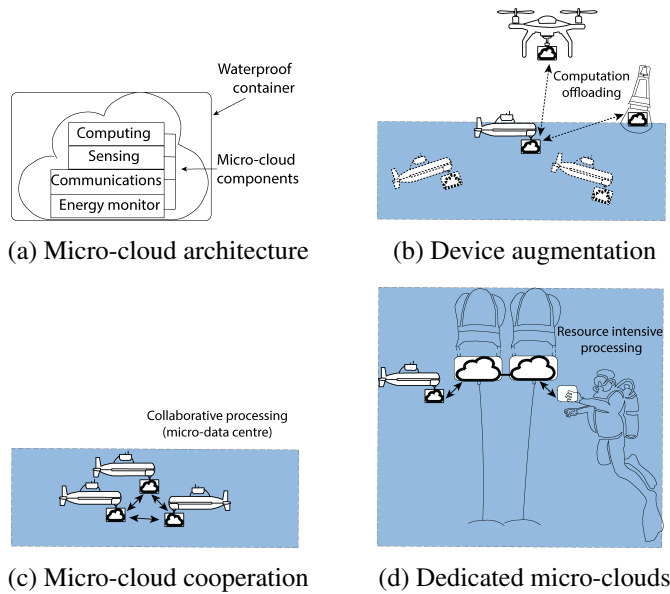


Figure 43: Initial micro-cloud deployment, (a) Micro-cloud architecture, (b) augmenting individual devices, (c) collaborating to improve fog processing performance, (d) dedicated micro-clouds in a location that can be used by underwater IoT devices in a contact-based manner.

Communications: Micro-clouds are equipped with interfaces compatible with widely used communication technologies. These interfaces enable micro-clouds to respond to computational requests from other devices, establish cooperative networks, collaborate on task execution, and offload data and computation to external fog or cloud infrastructures on the surface. Although underwater wireless signal absorption presents challenges [337], wireless communication remains viable for devices within close proximity to each other or near the sea surface. In this study, Wi-Fi is utilized for communication among devices within a single fog node (i.e., components forming the micro-cloud) and between nearby devices, such as ROVs, (AUVs), or divers operating in the vicinity of the fog node. For connections with buoys, ships, or infrastructure located farther away, a dedicated underwater communication interface is required [263]. As demonstrated in Section 5.7, optical (light-based) communication can extend the communication range and enhance data transfer rates for devices interfacing with the fog or collaborating to form it [144, 94]. While optical communication is a promising solution, other underwater communication methods, such as acoustic and electromagnetic technologies, can also be explored for longer-range connectivity. However, these alternatives currently require specialized equipment and are not readily available as low-cost, off-the-shelf options for enabling fog computing.

Sensing: The micro-cloud prototype incorporates sensors integrated into smart devices to assess water motion (turbulence) affecting the micro-cloud. For example,

accelerometer and gyroscope sensors are employed to detect significant water movement. This data is used to identify optimal communication conditions by recognizing periods of low turbulence, thus optimizing the timing for data sampling. This is particularly important as turbulence can disrupt communication and lead to increased resource overhead. Motion sensors also play a role in estimating the submersion depth of the micro-cloud. For instance, high motion is typically observed near the surface, whereas significantly lower motion occurs near the seabed. Additionally, we envision the use of external sensors deployed outside the micro-cloud to regulate its submersion process. Temperature and pressure sensors, for example, can provide depth information, enabling the device to initiate intensive processing that benefits from natural cooling provided by the environment. Oscillations in wireless signals could also be leveraged to detect submersion, allowing devices to adjust their duty cycling operations accordingly, such as reducing automatic Wi-Fi discovery underwater. Furthermore, surface sensors can inform the submerged micro-cloud to ascend when solar power availability is suitable for recharging its battery resources.

Energy Monitoring: Although micro-clouds effectively enhance the processing capabilities of underwater devices, isolated micro-clouds are prone to significant energy consumption. To address this challenge, processing tasks can be outsourced through collaboration with other micro-clouds or by offloading workloads to dedicated fog and cloud infrastructures. To facilitate this, the micro-cloud prototype includes an energy monitoring component that profiles the energy demands of specific tasks. This enables the efficient distribution of running processes, optimizing energy usage and ensuring sustained operation.

5.4.3. Fog Provisioning Underwater

The current prototype of the micro-cloud for underwater edge deployments utilizes wireless signals for service provisioning. Since micro-clouds are composed of aggregated devices, the off-the-shelf service discovery mechanisms integrated within smart devices are employed. These mechanisms leverage Peer-to-Peer functionalities embedded in the default implementation of direct Wi-Fi (Wi-Fi P2P API) to facilitate service discovery. In this study, we demonstrated that light-based communication (see Section 5.7) can overcome the limited signal coverage imposed by Wi-Fi (within 10 – 12 cm range). However, as light communication technologies are not yet sufficiently advanced for widespread off-the-shelf adoption, Wi-Fi remains the primary communication technology for our micro-cloud. Consequently, we envision two primary use cases for the micro-cloud. The first deployment involves fixed IoT Device with LAN Access. The micro-cloud can be deployed in a fixed location, anchored to a floating buoy, for PAM. In this scenario, it is accessible on the water surface and can support applications such as real-time analysis of underwater acoustic signals [221]. The second use case being Mobile Deployment with AUVs wherein the micro-cloud can operate in collaboration with AUVs,

positioning itself near underwater devices requiring additional computing power. For instance, cooperative networks of multiple AUVs equipped with cloudlets could enable advanced underwater operations [176].

5.5. Computational Benchmarks

To illustrate the potential of submersible fog as a versatile solution for enhancing the computational capabilities of underwater applications, we conduct comprehensive computational benchmarks. These benchmarks are designed to represent the processing demands of underwater data science tasks, while also aligning with commonly used fog benchmarks. In Section 5.8, we further validate the feasibility of deploying the micro-cloud underwater by implementing it on the seabed. In these experiments, we assess the data processing and transfer requirements of various applications by selecting tasks from the DeFog benchmark suite [188] which exhibit characteristics similar to those of underwater data science tasks. We quantify both computation and communication latency, evaluating the ability of submerged micro-clouds to manage the workload of users accessing processing resources. By incorporating a broader range of applications with varying computational demands, we gain insight into the processing stress experienced by the micro-clouds. Additionally, we measure the energy consumption of the micro-clouds while processing intensive computational tasks under submerged conditions. Since submerged micro-clouds are designed to complement, rather than replace, other infrastructure, we also conduct an experiment to measure data transmission with external cloud and micro-cloud systems. This emulates a scenario in which devices can offload processing when opportunistic connectivity to external sources is available. The overall prototype and experimental testbed are shown in Figure 42.

Experiments and Metrics: We evaluated several performance aspects of the micro-cloud. Initially, we measured both computation latency (RTT) and communication latency (CL). Computation latency was assessed based on three factors: the time required to access the resource, the execution time of the task, and the time taken to transmit the result. Communication latency (CL) refers to the data transfer time involved in back-and-forth interactions, excluding any resource access time. Additionally, we examined the micro-cloud’s ability to handle the workload generated by concurrent users submitting tasks (multi-tenancy). Finally, we measured the energy consumption (EC) of the micro-cloud during the processing of computational workloads.

Apparatus: We utilized a Raspberry Pi 4 (RPi4) and an LG G4 mobile phone (LGP) as the processing units for the micro-cloud computer board, with each being used independently in separate instances. The RPi4 is equipped with up to 4GB of RAM and a Quad-Core Cortex-A72 (ARM v8) 64-bit SoC running at 1.5 GHz. The LGP operates on Android version 6.0 and features a removable Li-Ion 2540 mAh battery. To measure energy consumption underwater for both processing units, we developed an energy monitoring system using an Arduino board in combination

with an Adafruit INA260 current sensor³. The current sensor was placed in line with the positive wire of the USB cable connecting the RPi4/LGP to the external battery pack. The sensor was then connected to the I2C pins (SCL—I2C clock pin, SDA—I2C data pin) of the Arduino MEGA ADK development board to measure the current flow. An application running on the Arduino board captures the current sensor data every 100 ms and stores it on an SD card mounted on the board, enabling real-time tracking of energy consumption while the micro-cloud is submerged. To enhance the accuracy of the energy measurements, a DS1307 Real-Time Clock (RTC)⁴ module was incorporated, providing real-time timestamps for the sensor readings. The accuracy of our energy measurements was found to be comparable to those obtained using a commercial multimeter, such as the Peaktech 3430⁵. DeFog was executed on the RPi4, while the LGP was used to run offloading applications.

Setup: Two sets of experiments were conducted: baseline and underwater experiments. The baseline experiments were performed to benchmark the micro-cloud when deployed outside of water. Subsequently, we conducted the same experiments to assess the impact of water on the micro-cloud’s performance when submerged. To achieve this, the micro-cloud prototype was enclosed in a waterproof glass container to safeguard the processing, energy, and sensor components from potential water damage. The micro-cloud was then submerged, and the previously described metrics were collected. As the primary goal of these experiments was to benchmark computational performance rather than demonstrate practical feasibility, we limited the experiments to a Depth-1 level (i.e., 1~5 cm depth) and ensured minimal water motion at the surface.

Tasks: To obtain performance metrics of the micro-cloud underwater, we used four distinct applications from the DeFog benchmark suite [188]. Each application processes different types of input assets to trigger task execution. Additionally, we developed two offloading modes to migrate computation to external sources. The first mode offloads a long data stream, while the second mode offloads computing tasks at the code level. Below, we provide a brief description of the four DeFog applications and provide examples of underwater data science applications in various fields, where the computational tasks align with those in the benchmark applications.

- **YOLO** A deep learning-based object classification application utilizing the YOLOv3 dataset. In this experiment, asset images with an average size of 223 KB are used. Deep learning-based object classification is commonly applied in various underwater data science applications, such as marine plastic monitoring [97, 309] and reef ecosystem monitoring [297].
- **PocketSphinx** A speech-to-text conversion engine that processes audio files

³<https://learn.adafruit.com/>

⁴<https://www.adafruit.com/product/3296>

⁵<https://www.peaktech.de/>

as assets. In this experiment, the audio files have an average size of 207 KB. A pertinent example of an underwater data science application using acoustic signals is the discrimination of marine mammals from vocal calls [192].

- **Aeneas** A text-audio synchronization application that enforces alignments of text-audio entries. In this context, audio files with an average size of 400 KB are used as assets. Forced alignment is applied to assess man-made noise pollution levels in marine environments, such as mapping the noise generated by vessel propellers [324].
- **iPokeman** A latency-critical GPS application for VR Online Mobile Games. It processes asset files with an average size of 131 KB. Georeferencing marine mammal trajectories is used in citizen science applications, allowing users to report marine mammal sightings while aboard sea vessels [315].

To generate the workload of users executing these applications, we employed JMeter, a load-testing tool designed to generate dynamic and concurrent user workloads. Using this tool, we examined the effects of increasing workloads on the underwater micro-cloud. The user workload ranged from 1, 2, 5, 10, 25, 50, 100, to 250 users. To further analyze underwater offloading, we developed two specialized offloading applications, which are described below.

- **Stream app** The first application represents long data streaming to an external source. It consists of an image processing application that applies a Box Blur filter to images. The application retrieves an image from the local device storage and processes it with the filter, creating a blurring effect to obscure image features. Images of varying sizes—0.5 MB, 1 MB, 3 MB, and 5 MB—are used for this task.
- **Code App** The second application demonstrates computation offloading at the code level. It involves a chess game application⁶ based on the MinMax algorithm optimization. The application transmits the current chessboard state to the MinMax algorithm, which computes the optimal next move. Unlike the Stream App, this application transmits smaller data packets with an average size of 170 KB, which are used to trigger resource-intensive processing.

For both offloading applications, we evaluated the offloading process to execute tasks on cloud and fog infrastructures, respectively. The next section presents an in-depth evaluation of our micro-cloud through the rigorous experiments described in Section 5.5.

5.6. Results

We conducted comprehensive benchmarking to evaluate the performance of submersible micro-clouds, focusing on key metrics such as computation latency, communication latency, energy efficiency, and multi-tenancy capacity. Additionally, we compare the performance of submerged computing systems against a

⁶<https://github.com/huberflores/CodeOffloadingChess>

traditional above–surface deployment baseline.

5.6.1. Underwater Computing Performance

We started by quantifying the impact of submersion and low water motion on computation and communication latency within the micro-cloud. Figure 44 presents results for three applications: YOLO, PocketSphinx, and Aeneas, alongside baseline measurements obtained from above-water deployments for comparison. As anticipated, water introduces a communication overhead. For example, YOLO experiences a threefold increase in asset transmission time, while PocketSphinx and Aeneas incur an additional 1–second delay. Furthermore, we observed that submersion also impacts computing latency, with task execution times increasing when the micro-cloud operates underwater. For instance, YOLO’s average computing latency rises from 8 seconds above water to 12 seconds underwater. Given that communication latency accounts for only 2 seconds of the delay, the remaining 2 seconds are attributed to the device’s data handling processes. Similar patterns are evident in PocketSphinx and Aeneas. We further demonstrated that this effect is more pronounced at shallow depths. As depth increases beyond 1 meter, improved heat transfer and cooling mitigate the latency overhead caused by submersion.

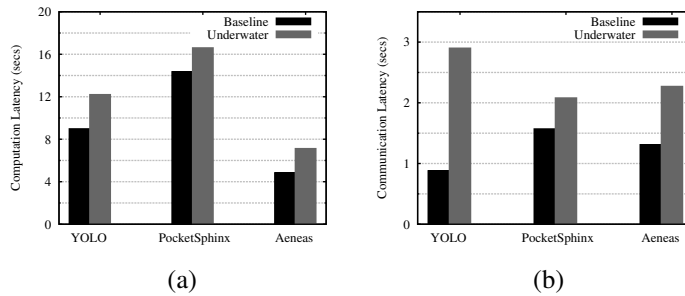


Figure 44: Latencies of different applications, a) Computation Latency, b) Communication Latency

In summary, our findings suggest that applications within proximal range can leverage the external resources of the micro-cloud, although the performance gains may be lower compared to above-surface deployments. Nonetheless, the elimination of data transmission to the surface—or to a remote cloud—offers distinct advantages, even if the submerged micro-cloud does not achieve the same performance level as its surface-based counterpart. The results further demonstrate that these benefits can be realized using standard wireless communication technologies rather than specialized interfaces, highlighting the potential for scalable, modular, and easily deployable computing solutions. For advanced applications, our analysis indicates that positioning the micro-cloud close to the application can significantly reduce transmission power requirements, thereby extending the operational lifespan of underwater system.

5.6.2. Underwater Capacity Performance

Submersible micro-cloud deployments are designed to support multiple concurrent users or devices, enabling the execution of simultaneous tasks through shared data processing resources (i.e., multi-tenancy). To evaluate the multi-tenancy capabilities of submersible micro-clouds, we analyzed their capacity to manage workloads generated by concurrent users submitting computing tasks for processing. Figure 45 illustrates the results, including baseline measurements for computation and communication latency, shown in Figures 45(a) and 45(c), respectively. While an overhead in communication latency due to concurrent transmissions is observed when comparing Figures 45(c) and 45(d), the increase is relatively small. For instance, with a workload of 50 users running the Aeneas application, communication latency increases by only 3 seconds. In contrast, computation latency exhibits a larger overhead. As shown in Figures 45(a) and 45(b), a workload of 50 users requires an average computation time of 20 seconds in the baseline setup, compared to 30 seconds underwater. This suggests that the overhead introduced by submersion is comparable to the overhead caused by multi-user workloads. On a per-user basis, the total delay is three times higher, with submersion contributing an additional 0.5 seconds to communication latency. Similar patterns are observed across other applications. The computation latency overhead percentages when comparing baseline and underwater results are 13% for YOLO, 6.63% for PocketSphinx, and approximately 25% for Aeneas. For communication latency, the overhead percentages are higher: 41% for YOLO, 21% for PocketSphinx, and 42% for Aeneas. An ANOVA test, with baseline and underwater deployment as experimental conditions, confirms significant differences in both computation latency ($\chi^2 = 11.27$, $p < .001$, Kendall's $W = 0.99$) and communication latency ($\chi^2 = 58.05$, $p < .001$, $\eta^2 = 0.996$). The observed overhead is likely due to the device's internal thermal management, as the Raspberry Pi throttles performance when temperatures rise significantly. This observed overhead is directly linked to the device's internal thermal management. During the 50-user workload tests at shallow depths, onboard telemetry indicated that the Raspberry Pi's internal CPU temperature rapidly escalated, peaking at 82° Celsius. At this threshold, the microcontroller automatically initiates thermal throttling, systematically reducing CPU clock speeds to prevent hardware damage, which explicitly explains the sharp increase in computation latency. These findings indicate that submersible micro-clouds are most suitable for small to moderate-scale deployments (e.g., supporting up to 10 devices or users). For larger-scale multi-user scenarios, surface-based edge computing is better suited. However, in real-world aquatic deployments, the internal heat issue can be mitigated by submerging the micro-cloud to greater depths, as demonstrated in Section 5.8).




Figure 45: Results of computation and communication latency for concurrent users, (a) Baseline computation latency, (b) Underwater computation latency, (c) Baseline communication latency, (d) Underwater communication latency.

5.6.3. Underwater Energy Consumption

We also evaluated the energy consumption of the micro-cloud during the execution of applications underwater. Monitoring energy usage is crucial for effectively off-loading computation and distributing the processing cost of tasks across available micro-clouds. Energy consumption was measured using the energy monitoring implementation detailed in Section 5.5, and its accuracy was validated against a multi-meter, which served as the baseline. The results, presented in Figure 46, demonstrate that the energy monitor closely aligns with the multi-meter, providing fine-grained and accurate energy consumption measurements. A non-parametric ANOVA test, using the multi-meter and energy monitor as experimental conditions, confirmed that there were no significant differences ($\chi^2 = 1.628$, $p = 0.21$, Kendall's $W = 0.999$) between the two measurement methods.

Although previous studies have shown that the energy consumption of the RPi4 remains consistent between idle and active modes [198], an energy overhead is still observed when the RPi4 operates underwater. Figure 46(c) illustrates the results for underwater energy consumption, highlighting the overhead caused by increased transmission efforts and the heavier processing load on computing resources. When analyzing energy monitor deployment as experimental conditions, a non-parametric ANOVA test confirmed statistically significant differences in energy consumption between surface and underwater scenarios ($\chi^2 = 15$, $p < 0.001$, Kendall's $W = 0.999$)

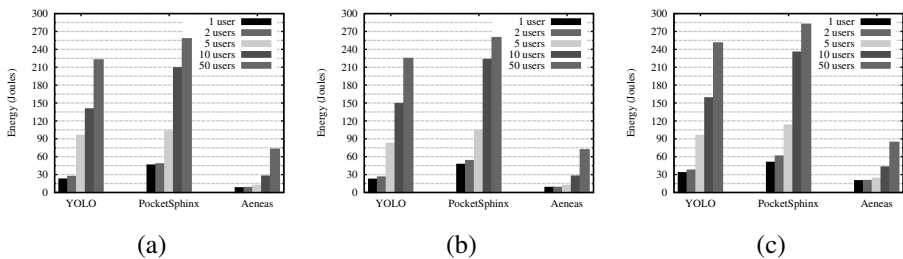


Figure 46: Results of energy consumption underwater, a) Baseline with multi-meter, b) Energy monitor (outside the water) c) Energy monitor deployed underwater.

5.6.4. Submersed Fog to Cloud Performance

Given that micro-cloud deployments are not isolated and require synchronization with the cloud, we also analyzed communication between the submerged micro-cloud and the cloud using the iPokeman application. This application introduces an additional step in task execution, whereby the results of a task processed by the micro-cloud are uploaded to a server in the cloud. The results, presented in Figure 47, indicate a higher communication latency overhead. This extra overhead creates a bottleneck in communication resources, limiting the number of concurrent users the system can accommodate. This limitation is evident when comparing Figures 47(a) and 47(d). While the baseline setup retains sufficient resources to support more than 250 users, the underwater deployment exhibits increased computation latency due to resource overutilization. Therefore, micro-clouds should be equipped with operational policies tailored to their deployment environment, whether situated on the water surface or submerged at a specific depth. The findings further indicate that submerged micro-cloud deployments are more efficient for processing data directly underwater than for transferring it to surface-based infrastructure.

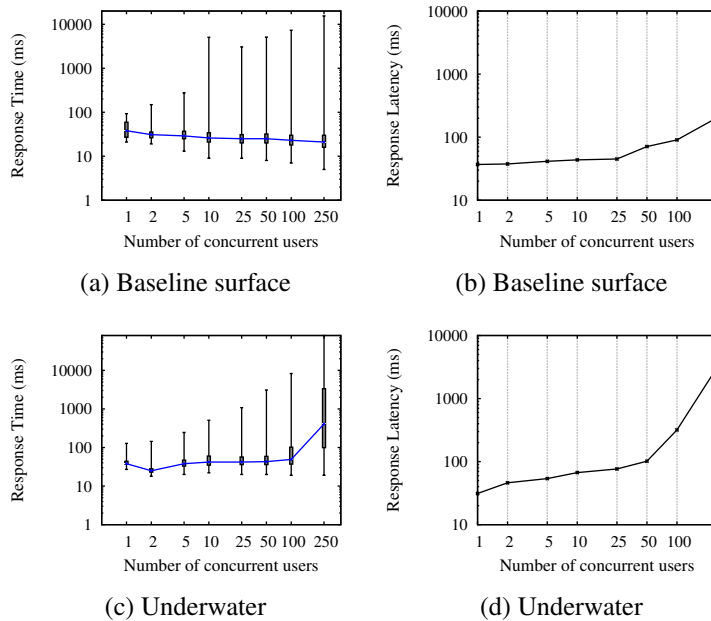


Figure 47: Impact of concurrent users using iPokeman, a) Response time, b) Response latency, c) Response time underwater, c) Response latency underwater.

5.6.5. Offloading from Underwater to the surface

Figure 48 illustrates the results of offloading computing tasks underwater. We measured the total energy consumption and response time for each application

under three conditions: (i) executed on the device, (ii) offloaded to an edge server (micro-cloud), and (iii) offloaded to a cloud server. The results show that the Stream application consumes more energy than the Code application. Additionally, offloading to the cloud results in greater overhead in both response time and energy consumption compared to offloading to the edge server. These patterns remain consistent when testing underwater, though we observed that water induces higher energy consumption and response times for both applications. When offloading occurs between underwater devices in close proximity, submerged micro-clouds function as edge servers, allowing underwater IoT devices to benefit from the same offloading advantages. Our findings further suggest that while micro-clouds are effective for underwater computing, they can also be used for synchronization with surface-based infrastructure. However, as communication remains the primary bottleneck, synchronization updates should be minimized to reduce overhead.

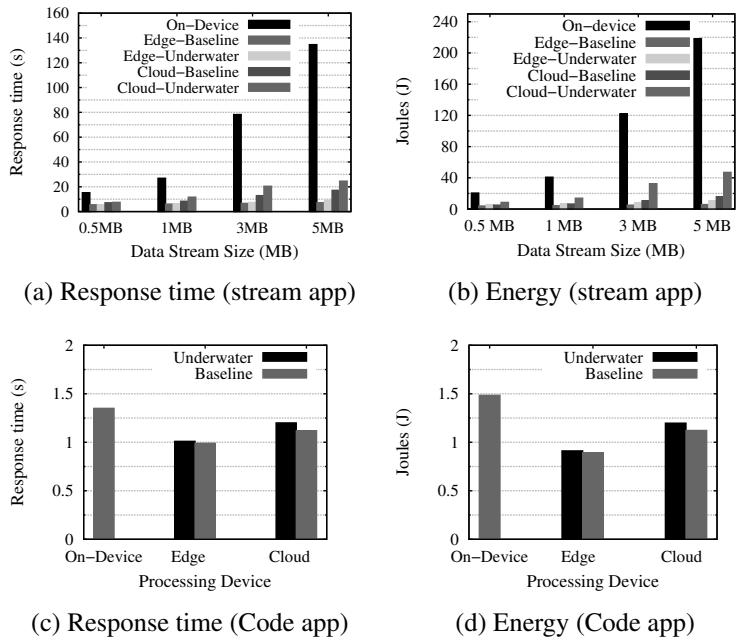


Figure 48: Response time and energy consumption for both apps (a-b) Stream app, (c-d) Code app.

5.7. Underwater Optical Communications

Given that wireless communication is unreliable and limited for long-distance use in underwater environments, we next evaluate the potential of light communication as an alternative medium. Unlike radio frequencies, the visible light spectrum experiences lower attenuation in water, making it a promising technology for supporting long-distance communications in underwater systems. Below, we




Figure 49: Testbed for analysis of underwater communication with light.

provide a brief overview of the setup and results from our optical communication experiments.

Apparatus: For the light source, we use a 650 nm, 5 mW, 3 – 5V red laser diode, and an Arduino Mega ADK microcontroller (ATmega2560) to design a transmitter that uses light for data transmission. The receiver consists of a solar panel (2.5 cm × 2.5 cm) connected to an Arduino board, which processes the received data. To generate water motion, two different sources are employed to agitate the water continuously. A pond aeration pump (Ubbink Air 100⁷) is used, which delivers air at a rate of 100 liters per hour through a 3 watt pump. Additionally, a hand mixer with a powerful motor (Model: House HB 1935, 200W) is used to create significant turbulence in the water.

Setup: Our testbed, shown in Figure 49, is constructed using a water tank with dimensions of 40 cm × 20 cm × 25 cm. The transmitter and receiver are placed on opposite sides of the tank, outside the water, and fixed in such a way that the emitted light directly hits the center of the solar panel. We developed an application that transmits data using Morse code as the encoding mechanism, chosen for its compatibility with the constrained resources of the microcontrollers. In this setup, a dot is represented by a 1 ms laser pulse, and a dash by a 3 ms pulse. While these intervals can vary, we observed that increasing the speed of data transmission leads to heavier processing demands on the devices, creating a bottleneck when decoding data. Therefore, we opted for an optical configuration that minimizes processing load.

Baseline: For the baseline setup, we positioned the laser transmitter at one end of a corridor in the university building and the receiver at the opposite end. The maximum transmission distance achieved was approximately 100 meters. At this distance, we successfully transferred a 1kB text file to the destination in 5 seconds and a 10 kB file in about 50 seconds. Notably, the light intensity from the transmitter was consistently detected by the solar panel, ensuring that there was no data loss over this distance.

Experiment: Data transmission was conducted in intervals of 6 minutes. In the

⁷www.ubbinkgarden.com

first experiment, we observed the transmission behavior in calm water. Next, we performed an experiment in which data transmission began in calm water during the first minute, followed by induced water motion for the next minute. After this, the induced motion was stopped, and the cycle was repeated for next 6 minutes.




Figure 50: Water motion (Calm water), a) Accumulated motion captured through all axes (variance), b) Experience motion (one axis only).

Results: Figure 50 illustrates the light communication results within water. To quantify water motion, we used an accelerometer floating on the water surface and estimated the overall motion by calculating the variance of the three axis over time. Figure 50(a) shows the overall motion in calm water. Figure 50(b) presents the motion data along the y-axis, along with the communication success rate, where a value of 1 represents a successful transmission and 0 represents a failure. We then assessed the performance of light communication when water motion was induced by the air pump. Figures 51(a) and 51(b) show the overall motion and motion along the accelerometer’s y-axis. The results indicate that the moderate water motion induced by the pump does not disrupt light communication. The baseline results are comparable to these two underwater scenarios: calm water and induced water motion using the pump. Finally, we examined the performance of light communication when water motion was induced by a more powerful motor (a hand mixer). Figure 51(c) shows the overall motion captured by the sensor. The results reveal that the water motion is more pronounced when using the mixer. More importantly, as depicted in Figure 51(d), the higher level of induced water motion disrupts light communication. Non-parametric Spearman correlation [254] analysis indicates a significant negative relationship between increased water motion and communication success ($\rho = -0.09, p < .05$). Interestingly, the data suggests that after stopping the motion generated by the mixer, a few seconds are required to restore the communication link.

5.8. Ocean Deployment

The previous experiments have primarily focused on controlled benchmark conditions, utilizing water containers or shallow depths. To further validate the feasibility of the solution, we extended our experiments to actual underwater environments,

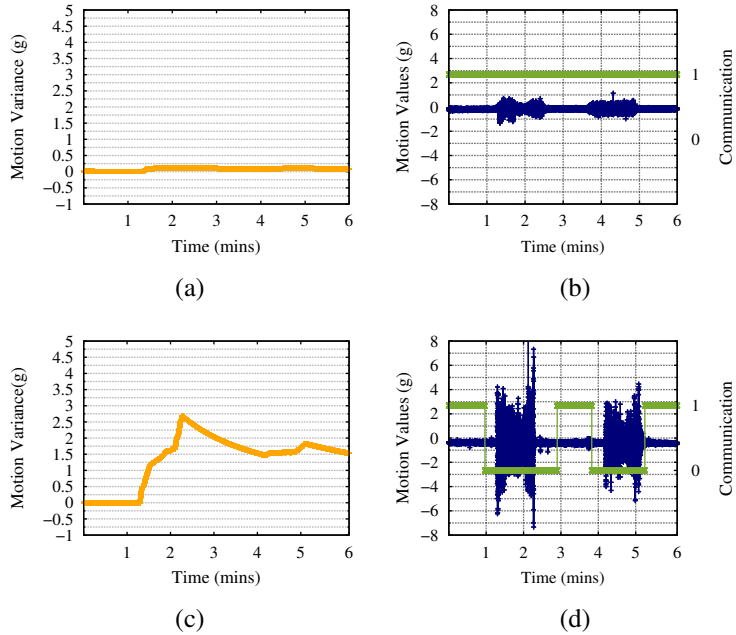


Figure 51: Light communication performance under different levels of induced water motion, a) Overall motion using AirPump (all axes) b) Motion using AirPump (one axis), c) Overall motion using a more potent motor (all axes) d) Motion using a more potent motor (one axis).

conducted during recreational scuba diving activities.

Apparatus: In this experiment, we designed a setup where the devices comprising the micro-cloud are housed inside a PVC-supported acrylic sphere. The fog nodes in our configuration are repurposed older-model smartphones, which provided support for deep learning-based object detection [224]. Two deployment scenarios were tested: one with two phones inside the same sphere (Fig. 52a), and another with two separate spheres attached to a horizontal bar, positioned adjacent to each other (Fig. 52). The spheres, containing air inside, are designed to float, with diving weights used to submerge them. A 6 kg weight is used for a single micro-cloud, while dual micro-spheres are submerged with 12 kg of weight, anchoring them to the seafloor (Fig. 52c). Additionally, a mobile scenario is tested where a diver transports the micro-cloud during a dive transect survey (Fig. 52d). For the experiments, the smartphones inside the spheres were running deep learning-based image recognition tasks. As the performance and resource utilization of image classification depend on the type of input image, a fixed set of prerecorded images were chosen for consistency [224]. Image data serves as a crucial source for underwater data science due to its ability to be collected without disturbing the environment. It is widely used in various real-world applications such as litter detection, fish school estimation, biodiversity monitoring, and pipeline leak estimation [224].

Therefore, our experimental setup directly reflects the needs of these underwater data science applications, ensuring that the performance metrics obtained in our study are relevant to actual use cases in the field.




Figure 52: Underwater Micro-Cloud Computation Units (MCCUs) consisting of individual mobile devices, communicating using off-the-shelf integrated wireless interfaces while collaborating in the execution of a task, deployed in open sea at 8m depth. From left to right, top to bottom: (a) Master and worker within same unit, (b) Master to the left and two workers to the right, (c) Seabed deployment and (d) SCUBA diver survey.

Experimental Setup: The experiments were conducted at the Carlton diving reef in Madeira Island, with both surface and underwater tests. The surface tests were carried out in 25° Celsius conditions, while the underwater tests were conducted at a depth of 8 meters, with an underwater temperature of 21° Celsius. Two types of deployments were tested:

- **Single Micro-Cloud Computation Unit (MCCU)** This setup included two mobile phones (one master and one worker) mounted in the same casing, positioned with screens on top of each other (Fig. 52a).
- **Dual Micro-Cloud Computation Units (MCCUs)** This configuration included three mobile phones (one master and two workers), with the worker devices housed in the same container (Fig. 52b).

The master device was a Sony Xperia M2, featuring a Quad-Core 1.2 GHz Cortex-A7 CPU, running Android 5.1. The worker devices were NOS NOVU II (Android 5.1) and Alcatel Go edition (Android 11). The task performed was image classifi-

cation using pretrained ImageNet models on the devices. To ensure consistency and test the performance of the fog system underwater, the same set of five images was used for all tasks, maintaining comparable memory requirements throughout the experiment.

Single MCCU Experiment: This experiment involved one master phone and one worker phone, with data collected both on the surface and at an underwater depth of 8 meters. The two phones were running separate mobile applications that allowed them to communicate via WLAN and participate in the computational task. As in previous experiments, the worker phone connected to the master's access point, where the master phone sent five images to the worker phone. The worker then performed image classification tasks on its CPU and transmitted the accuracy results back to the master phone via WLAN. The duration of the experiment was 44 minutes for both the surface and underwater setups. The data collected included environmental variables, such as accelerometer data, RSSI, RAM percentage, and CPU temperature, which were obtained from the worker phone with a duty cycle of $1/15Hz$. For the surface test, the MCCU was carried while walking, followed by the underwater test during a SCUBA dive. The underwater test was divided into three phases: pre-dive (11 minutes), dive (22 minutes), and post-dive (11 minutes), with the actual dive taking place at a depth of 8 meters.

Figure 53 shows the data from the experiments, comparing the underwater (blue) and surface (red) setups within the same micro-cloud. The graph includes accelerometer readings, CPU temperature, Received Signal Strength Indicator (RSSI), and memory usage. Vertical lines indicate the transitions between the pre-dive phase (surface tasks preparing for the dive), the actual dive, and the post-dive period. During the dive, the micro-cloud was submerged to a depth of 8 meters and stayed there until surfacing. The drop in RSSI values clearly indicates when the micro-cloud was submerged, reflecting weakened communication signals as it was taken underwater. Additionally, the temperature of the devices decreased rapidly as the dive began, with temperature effects being more prominent in shallow water near the surface. At deeper depths, heat transfer to the exterior helped stabilize the device's operations. The variations observed at the beginning and end of the experiment, during the pre-dive and post-dive phases, were likely due to the devices being moved. For instance, the early temperature differences between the two micro-clouds were a result of one device being in direct sunlight while the other was in the shade. Despite these environmental fluctuations, the communication between the master and worker devices inside the container remained stable throughout the experiment, even though the RSSI weakened during the dive.

Dual MCCU Experiment: The second experiment, which involved the dual micro-cloud MCCUs, and was designed to address the communication challenges encountered with a single MCCU setup, particularly in establishing communication between a client inside the container and one outside. In this experiment, two micro-clouds were submerged, each containing one master device in one micro-sphere and two worker devices in the second micro-sphere (as shown in Fig. 52b).

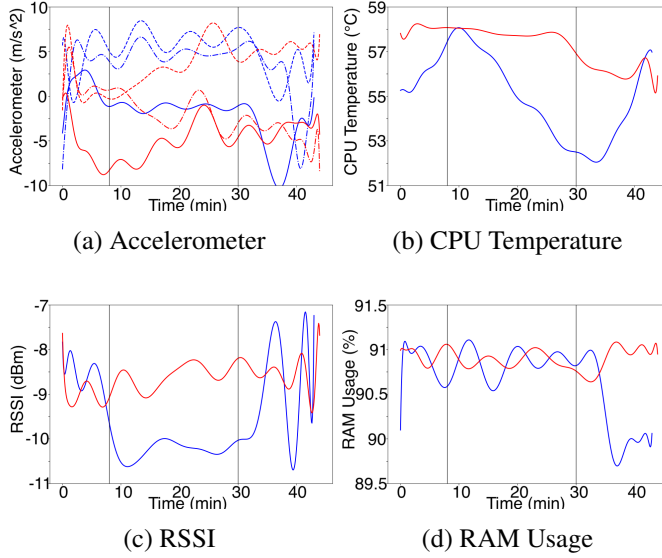


Figure 53: Obtained sensor data during ocean deployment experiment time. From left to right: (a) Triaxial accelerometer, (b) CPU temperature, (c) RSSI, and (d) RAM Percentage. Red line depicts the time at the surface, while blue line is the time underwater. Vertical lines indicate the SCUBA diving time for underwater test.

The two micro-spheres were placed 5 cm apart, adhering to the established Wi-Fi communication range for saltwater environments [278]. The RSSI measurements taken at a depth of 8 meters showed values of $-25dBm$ for the surface setting and $-41dBm$ for the underwater setting. These results indicate that stable data transmission was maintained in both environments. In Table 10, the computation times for image classification on 301 images over a 10-minute period are presented. With the single MCCU setup, 301 images were successfully classified in 19.26 minutes, while the dual MCCU setup classified the same 301 images in 12.3 minutes, achieving a speedup of nearly 7 minutes. This improvement is consistent with previous research [163], demonstrating that having multiple micro-clouds underwater can effectively accelerate computational tasks. The results highlight that robust computation can be successfully conducted in real ocean settings, and the use of more micro-clouds underwater not only ensures stability but also enhances computational efficiency.

5.9. Discussion

We have shown submersible micro-clouds to offer a potential way to support underwater data science applications and demonstrated that COTS components could be used to enable a submersible micro-cloud. Below we briefly discuss

Table 10: Comparison of underwater computation times during the image classification for 10 minutes. Increasing the number of workers expectedly speeds up computations which implies that the connectivity between devices is sufficiently stable to operate the micro-cloud despite adverse water conditions.

No. of phones	Comp. Time (s)	Mean (s)	St. Dev. (s)	No. of images	Setup
2	1156 (19.26 min)	3.84	0.3	301	Single MCCU
3	742 (12.36 min)	2.47	0.3	301	Dual MCCU

implications, possible extensions, the main limitations, and potential ways to overcome them.

Application Domains: The primary target for our research is underwater data science investigations that operate within a localized region. Examples of such applications include diverse monitoring tasks, such as pipeline integrity monitoring, reef condition monitoring, biodiversity monitoring, smart underwater navigation and litter monitoring [129, 224, 220]. At the same time, our solution can also support surface-based marine data science applications in areas that lack access to a traditional communication infrastructure (i.e., away from shore regions). Offering computational support for these kinds of applications is critical for scaling up such applications and reducing the delay between data collection and analysis. At the same time, the computing infrastructure needs to be environmentally sustainable to ensure it offers the required computational power but does not harm the underwater environment. Micro-clouds, as envisioned in our work, are well-suited for these needs, offering powerful and scalable computing support for applications that are characterized by high data velocity [163], while at the same time being easy to deploy on-demand. Indeed, our solution can be easily deployed as a temporary infrastructure instead of requiring a persistent deployment. As an example, we are currently using our solution to support scuba divers in underwater litter recognition by allowing the infrastructure to be deployed at the beginning of a dive and removed at the end of the time [224].

Implications for Fog Computing: A fundamental challenge in adopting any computational infrastructure is the need for a static and permanent physical deployment location. Indeed, available space, suitable deployment facilities, e.g., rack, cooling system, energy supplies, and sufficient computing hardware are critical to providing services to a large number of users. Our work offers a way to alleviate these issues by adopting small scale data-centres that can be submerged on-demand and potentially moved taking advantage of currents. Our results also suggest that old computing devices could be recycled for underwater settings to provide computing infrastructure near users. Naturally, it is important to ensure the micro-clouds are safely attached so that they do not get lost and end up polluting the underwater environment.

Towards large-scale deployments: The present paper demonstrated how micro-clouds could be easily deployed in the open sea for short-term deployments, e.g., included as part of scuba diving missions. For longer term deployments that

could provide a broader range of underwater scenarios, such as sensors deployed onto the seabed, there are further challenges that need to be addressed. First, the waterproofing we used was designed to protect the computing units, not to offer a long-term solution. Indeed, we used acrylic spheres for the encasing of computing resources. These can only handle moderate depths (e.g., around 20 meters which is a common recreational diving depth) before pressure accumulation would break the spheres. As a result, better encasing solutions are needed for longer-term deployments and for operating deeper. Optimally, the casing should also offer adjustable buoyancy, enabling it to operate at different parts of the water column instead of being limited to the seabed. Finally, adverse environmental conditions, such as heavy turbidity, salinity or turbulence also pose challenges that require further research. Nevertheless, our research shows how it is possible to support underwater data science applications through low-cost micro cloud based fog designs, offering the first steps in developing broader computing support for underwater deployments.

Communication Interfaces: Ideally, a submersible micro-cloud has to be equipped with multiple communication interfaces, and a context-aware mechanism to decide which communication interface to use based on water conditions, e.g., high turbulence. While we have shown that submersible micro-clouds are feasible and useful (with close contact using wireless and longer range connectivity using optical connectivity), practical deployments would need to support distances of a few meters at a minimum. This can only be achieved through the use of communication mediums that are suited for underwater environments, such as optical, acoustic, laser, or electromagnetic communications. Currently, these technologies have not yet reached a level where standardized communication interfaces – let alone low-cost ones – would be available. For these reasons, we omitted their use within the benchmarks, as integration with proprietary interfaces would result in additional overheads and make it difficult to separate the computational performance of the offloaded task and applications from the overheads caused by the communication interface.

Water Conditions: The experiments were conducted both in a river and in an ocean environment with varying water motions. The results were stable across these experiments, and showed stable computational performance even in the presence of significant water motion. The main challenge, thus is not the computational aspect of the fog node, but to having robust enough connectivity. In practice, salinity, turbidity, level of pollution and extent of algae can affect the performance of both the micro-cloud and the communication interface. In particular, water characteristics affect propagation of signals as well as the thermal conductivity of the water, which is critical for cooling and thermal management of the submersible micro-clouds. Additionally, algae or particles in the water can accumulate on the surface of the casing hosting the fog and this can reduce access from the outside of the casing. Overcoming these issues requires improved material designs (that

are beyond the scope of our work) for casings besides more robust and affordable communication interfaces.

Surface-Based Micro-Cloud deployments: We demonstrated that both extents of water motion and depth of the deployment affect computational performance. Besides being highly relevant to submersed deployments, these results are also highly relevant for surface-based deployments, e.g., other computing infrastructure attached to buoys or sea vessels. Waves and other water motions can cause surface-based deployments to be momentarily submerged, which can cause disruptions in handling and processing requests. The reliability of such deployments can be improved by integrating motion-based techniques, such as accelerometers used in our work.

On-demand Fogs: While submersing and deploying permanent micro-clouds in critical areas that require continuous monitoring is important, other areas that are monitored occasionally can rely on on-demand infrastructure that is carried and deployed temporally in a location. For instance, aerial and underwater autonomous vehicles can be used for this purpose. Similarly, other transportation means can be envisioned, such as hot air balloons, airships and mobile buoys.

Multi-Modal Energy Harvesting: Tidal harvesting and solar cells are two of the most promising technologies for generating sustainable energy for underwater devices. While several works have demonstrated that energy can be harnessed using these techniques, the energy gains they offer remain small and are unlikely to suffice the needs of underwater devices facing continuous and resource intensive processing. This suggests that a multi-modal approach to preserving battery life underwater can be more effective in fostering longer explorations. We demonstrate the usage of computation offloading which complements that vision. Indeed, by using computation offloading underwater, it is possible to preserve devices underwater for longer periods of time.

Thermal Management and Casing: Our implementations of micro-clouds used sealed waterproof containers for the micro-controllers. The lack of heat exhaust can result in heat accumulation inside the container, which in turn can trigger device-internal thermal management which throttles the CPU performance. We have observed this phenomenon in earlier experiments that were conducted near the surface, but in the ocean experiments this did not occur. This is potentially a result of better thermal management in the devices that are used as fog nodes and from the deeper water being able to cool the container effectively. Further improving the performance of submersible micro-clouds requires research on casing solutions that are sufficiently lightweight to allow attaching the infrastructure into underwater devices or objects, while at the same time having sufficient cooling capacity. Another limitation of the casing that we used in the experiments is that it suffers from the fact that it houses air inside it, requiring separate weights for submerging it. Removing – or at least reducing – the air pockets is thus needed to make the overall platform easier to deploy.

Recycling opportunities for e-waste: Electronic waste from smart devices is a global concern as it pollutes natural ecosystems and fosters climate change. In our work, we demonstrate that micro-clouds can be made from aggregated smart and IoT devices. We envision that computing resources from old phones can be recycled to create portable computing racks, which then can be deployed on edge underwater to provide public services to users. For instance, a video streaming service for tourists about the sightseeing places in a city.

AUVs for fog delivery: In our work, we demonstrated the design and development of micro-clouds that can be used for underwater edge deployments. One important insight of this work is related to the weight of the micro-cloud. Indeed, micro-clouds are lighter in weight and have compact size after being encased. Thus, micro-clouds can then be transported underwater easily by AUVs. This suggests that it is possible to provide mobile fog computing services underwater.

5.10. Applications for Environmental Sustainability: Outcomes and Implications

Micro-clouds represent a transformative advancement in sustainable environmental monitoring and data-driven management for aquatic and remote environments. Built from COTS hardware and utilizing fog computing principles, micro-clouds bring efficient, localized computational power directly to areas where traditional connectivity and infrastructure are lacking. This technological innovation unlocks a diverse array of application scenarios and delivers substantial sustainability benefits, reshaping the landscape of environmental stewardship [129].

5.10.1. Application Scenarios

Micro-clouds enable a broad spectrum of underwater and remote environmental data science applications. They process high-frequency data streams from sensor arrays monitoring key parameters such as water quality, pH, salinity, temperature, hypoxia, and pollutant levels [129]. By executing analytics and anomaly detection algorithms at the edge, micro-clouds reduce the latency between data capture and analysis, facilitating rapid detection of potentially hazardous events including chemical spills, algal blooms, and sudden hypoxia. This supports adaptive ecosystem management and bolsters early-warning systems for environmental disaster prevention [91, 58].

Continual biodiversity and habitat surveillance are made feasible by onboard aggregation and machine learning, which enable the monitoring of marine life, habitat changes, and anthropogenic disturbances with unprecedented granularity. For example, coral reef monitoring benefits from continuous observation of bleaching dynamics and species richness, all processed locally on micro-cloud infrastructure [58]. Infrastructure applications include automated event detection for submerged pipelines, offshore energy installations, and aquaculture systems, en-

hancing operational safety by alerting stakeholders to faults, leaks, or unauthorized access [197].

In litter and pollution mapping, micro-clouds process sensor data—including vision, sonar, and chemical readings—to facilitate the rapid identification and geospatial mapping of macro- and microplastic debris, as well as industrial discharges. Such edge-based assessment is critical for timely remediation, especially in areas with limited connectivity [224, 246]. The affordability and portability of micro-clouds further empower field deployments and citizen science initiatives, lowering technical barriers, amplifying community engagement, and promoting high-frequency, large-scale data collection [229].

5.10.2. Sustainability Impact

Micro-clouds advance environmental sustainability across multiple dimensions. By supporting localized, edge-based analytics, they optimize energy and resource efficiency, minimizing both the energy required for data transmission and the infrastructure footprint for centralized processing [52]. Computation offloading onto micro-clouds conserves battery life for field sensors and AGVs, prolonging deployment duration and spatial coverage.

A further sustainability outcome is realized through the upcycling of decommissioned electronics—transforming e-waste into functional distributed computing platforms. Micro-cloud deployments delay device entry into the waste stream and reduce the environmental consequences of electronic waste, powerfully advancing circular economy principles [229, 326]. Their temporary or mobile deployment model lessens ecological disturbance and makes monitoring feasible in sensitive or infrastructure-poor environments.

The democratization of environmental analytics is another vital outcome. Micro-clouds facilitate high-frequency and high-resolution monitoring in previously inaccessible regions, supporting biodiversity conservation, pollution control, and climate adaptation efforts [103]. Their compatibility with renewable energy solutions—including solar, wave, and tidal power—further amplifies their impact, driving aquatic observatories toward low-carbon, sustainable operations.

As micro-cloud design continues to evolve, improvements in casing materials, buoyancy systems, and energy harvesting strategies will enhance their environmental suitability and long-term resilience. Collectively, micro-clouds offer a scalable, flexible, and environmentally responsible basis for next-generation monitoring and management, directly contributing to global goals for sustainability, equity, and climate resilience.

5.11. Beyond Aquatic Environments: Repurposing Micro-Clouds (E-Waste Smartphones) as Tiny Data Centers

While the primary evaluation of the micro-cloud framework in this chapter has focused on marine ecosystems, the underlying physical and software architectures

are inherently domain-agnostic. Having demonstrated strict operational resilience in one of the most hostile deployment environments—underwater—the micro-cloud paradigm is equally viable for providing decentralized computational support to terrestrial areas lacking traditional infrastructure. Exploring this terrestrial capability is not a departure from the core thesis, but rather a realization of its broader environmental sustainability objectives.

The exponential growth in computing demand driven by artificial intelligence, data science, and IoT applications has created unprecedented pressure on computational resources while simultaneously generating massive amounts of electronic waste. Traditional computing devices typically become obsolete within 2 – 3 years, contributing to a projected doubling of global e-waste by 2050 [299]. This paradigm presents a critical sustainability challenge where the environmental cost of manufacturing new devices conflicts with the growing need for distributed computing power.

To address this dual challenge, we expand the micro-cloud architecture into the concept of sustainable "tiny data centers." This approach aligns directly with the paradigm of "junkyard computing," first introduced by [282], which advocates for harvesting computational utility from obsolete consumer devices to reduce environmental impact. By repurposing hardware classified as e-waste—specifically aging smartphones—into functional edge or fog computing nodes [163], this approach extends device lifecycles, reduces resource consumption, and provides cost-effective distributed infrastructure. Rather than simply recycling these devices for raw materials [118]—a process that is highly energy-intensive and recovers only a fraction of the embedded value—repurposing preserves the sophisticated engineering and manufacturing investment already present in the hardware.

By grouping these discarded devices into cooperative micro-clouds, we democratize access to edge computing in far-off, infrastructure-poor regions while actively reducing the environmental footprint of the hardware itself. Figure 54 illustrates the idea, architecture, and practical deployment of this sustainable tiny data center approach, detailing the powering, assembly stages, and a use case application ranging from a single-device setup to a multi-phone prototype.

5.11.1. Technical Implementation and Architecture

The transformation of obsolete smartphones into tiny data centers requires systematic hardware modification and software adaptation. The core technical challenge involves bypassing the integrated battery system, which typically fails first in aging devices, and establishing alternative power delivery mechanisms. Through careful circuit analysis and modification, devices can be powered using external DC-DC converters, enabling stable 3.7V power delivery from various energy sources including renewable alternatives.

The implementation process involves several critical steps. First, devices undergo hardware modification where the battery circuit is carefully accessed and

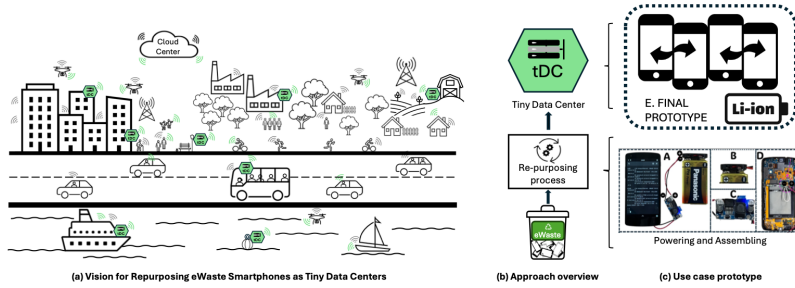


Figure 54: (a) Vision, (b) tiny data centers approach overview and (c) use case application. Powering and assembling stage: (A) Basic tiny data center: A single phone powered using a 9V battery connected to an LM2596 DC-DC, (B) External power source, (C) LM2596 DC-DC power regulator, (D) Phone rear-view. Final stage: (E) Final prototype: 4 phones and common power regulator.

connected to external power management modules such as LM2596 DC-DC converters. This modification, costing approximately 8€ per device, transforms battery-dependent smartphones into externally powered computing nodes. Second, software repurposing through custom operating system installation (such as PostmarketOS) unlocks the full computational potential of the hardware while removing manufacturer limitations and bloatware that reduce efficiency.

Performance benchmarking using High-Performance Computing (HPC) challenge tests reveals that while individual smartphones provide modest computational power compared to modern systems, their collective capability and energy efficiency make them viable for specific application domains. Testing with Nexus 5 devices showed LINPACK scores of approximately 163 MFlops, demonstrating sufficient capability for time-flexible computational tasks.

5.11.2. Applications and Deployment Scenarios

Tiny data centers constructed from repurposed smartphones demonstrate versatility across multiple application domains. In edge computing deployments, these devices can be strategically positioned in urban environments to provide localized processing for IoT sensor networks, smart city applications, and privacy-preserving analytics. Their low security risk profile—stemming from their e-waste classification—makes them suitable for public deployments where device replacement is more feasible than securing high-value equipment.

Autonomous vehicle applications represent another promising domain, where tiny data centers can serve as onboard processing units for ground-based vehicles. Their moderate weight and power requirements make them ideal for integration into vehicular platforms that require local data processing capabilities. The ability to aggregate multiple devices enables scalable processing power that can be customized based on specific vehicle requirements.

In pervasive computing scenarios, tiny data centers can be embedded into

everyday objects, toys, and interactive installations to provide intelligent behavior without requiring constant cloud connectivity. The integration potential with 3D printing technologies enables custom enclosures that protect the devices while facilitating airflow and connectivity requirements.

Renewable energy integration presents particularly compelling opportunities. Tiny data centers can be paired with energy harvesting systems—solar panels for terrestrial deployments, wave energy converters for marine applications, or wind generation for mobile platforms. This combination creates truly sustainable computing infrastructure that operates with minimal environmental impact while providing useful computational services.

5.11.3. Proof-of-Concept Demonstration and Performance Analysis

A comprehensive proof-of-concept implementation demonstrates the practical viability of tiny data centers for real-world IoT applications. The system architecture comprises sensor nodes collecting environmental data (temperature, humidity, light conditions), edge processing nodes (repurposed smartphones), and cloud visualization platforms. The edge nodes perform local data aggregation, computing rolling averages over one-minute intervals, and implementing threshold-based alerting mechanisms.

Performance monitoring reveals efficient resource utilization with CPU usage remaining below 5% during active data processing and memory consumption staying well within available limits. Power consumption analysis shows that a single repurposed device can operate continuously for over 3.5 hours on a standard 9V battery while processing approximately 2400 sensor samples. This demonstrates the feasibility of extended autonomous operation, particularly when combined with renewable energy sources or larger battery systems.

The system successfully processed over 12000 individual data points during testing, validating its capability for sustained data collection and processing scenarios. Network connectivity through standard WiFi protocols ensures compatibility with existing infrastructure while enabling remote monitoring and management capabilities.

5.11.4. Environmental Impact and Sustainability Benefits

The environmental benefits of tiny data centers extend far beyond simple waste reduction. By extending smartphone lifecycles from their typical 3.17-year lifespan to potentially decades of continued service, this approach dramatically reduces the environmental burden associated with device manufacturing. Each smartphone contains significant embedded energy and rare earth materials that are preserved through repurposing rather than lost through traditional recycling processes.

Energy efficiency comparisons show that repurposed smartphones consume significantly less power than equivalent traditional computing hardware while providing sufficient processing capability for many distributed computing tasks.

This efficiency translates directly into reduced carbon emissions, particularly when powered by renewable energy sources. The modular nature of tiny data center deployments enables precise capacity matching—using only the computational resources required for specific tasks rather than over-provisioning traditional server infrastructure.

The circular economy principles embodied in this approach demonstrate how technology can be designed for extended lifecycle rather than planned obsolescence. By proving that sophisticated computing tasks can be effectively handled by older hardware, tiny data centers challenge the assumption that constant hardware upgrades are necessary for computational progress. This paradigm shift has implications for corporate procurement policies, individual consumer behavior, and regulatory approaches to electronic waste management.

5.12. Summary

This chapter introduces the concept of *Micro-clouds*, a decentralized fog computing paradigm tailored to underwater environments. Built from COTS hardware, micro-clouds provide affordable, modular, and real-time computational support directly at the source of data. Their architecture addresses core challenges in underwater data science, including limited bandwidth, high latency, and the absence of land-based computing infrastructure, which have traditionally delayed analysis and restricted the scale of marine monitoring applications.

Micro-clouds leverage edge analytics to process high-frequency sensor streams—such as water quality, salinity, pH, temperature, and pollutant levels—near collection sites. This capability reduces the time lag in detecting hazardous events (e.g., chemical leaks, algal blooms), supports adaptive management of aquatic ecosystems, and enables continuous biodiversity assessment. The system also augments underwater drones, buoys, and sensor arrays by offloading data and executing complex analytics, such as deep learning-based object detection or anomaly identification, facilitating broader and more inclusive environmental surveillance.

Extensive experimental validation, including surface and open water deployments, demonstrates that micro-clouds—using devices like Raspberry Pi and repurposed smartphones—perform agile, energy-efficient, and scalable underwater computing. Key results show that micro-clouds maintain robust processing even when submerged, supporting concurrent workloads and remaining operational across diverse conditions, though communication and performance may be affected by factors such as water motion, submersion depth, and thermal management. The chapter highlights approaches for improving communication (including optical and Wi-Fi interfaces) and methods for handling energy and environmental constraints to prolong operational lifespans.

A notable use case presented is the repurposing of e-waste smartphones as tiny data centers. By modifying discarded devices and aggregating them into micro-cloud clusters, the framework extends hardware lifecycles, reduces electronic

waste, and enables sustainable, distributed computing infrastructure at minimal cost. These repurposed devices provide adequate processing for local sensor data, IoT analytics, and citizen science—demonstrated through proof-of-concept applications in environmental sensing and data aggregation.

Micro-cloud deployments embody circular economy principles, promote energy/resource efficiency via localized computation, reduce infrastructure footprints, and support clean, renewable operations. Their portability and scalability democratize environmental monitoring, allowing even remote or resource-limited regions to benefit from advanced data science capabilities. Ultimately, this chapter demonstrates how micro-clouds advance real-time, context-aware aquatic monitoring while significantly contributing to environmental sustainability and resilience by enabling proactive responses to ecological threats and supporting global stewardship goals.

6. CONCLUSIONS AND FUTURE DIRECTIONS

6.1. Conclusions

This thesis demonstrates how pervasive sensing and autonomous systems, combined with the creative repurposing of commodity hardware, can address urgent challenges in environmental sustainability. Through the design and validation of three core contributions—MIDAS for material characterization, LIZARD for sensing and monitoring plastics and Micro-Cloud for underwater/remote data science – this work establishes scalable, cost-effective, and practical methods for continuous environmental monitoring and resource management across diverse ecosystems.

The first contribution, the MIDAS sensing modality introduced an original approach leveraging thermal dissipation signatures to accurately and robustly characterize material types, supporting applications in waste management and recycling. The second major contribution, LIZARD, presents a novel sensing solution for autonomous plastic pollution detection, integrating thermal and optical sensing to identify macro, meso, and microplastics directly in the environment. Lastly, the fog computing framework based on micro-clouds, showcased how modular deployments of low-cost, commercial-off-the-shelf devices enable real-time data processing and actionable insights even in remote or infrastructure-poor aquatic regions.

Extensive experimental studies, including real-world case deployments, highlight that the synergy of pervasive sensing and autonomous intelligence enables fine-grained, extensible, and inclusive environmental data collection. Collectively, these innovations not only advance scientific understanding but also lay a foundation for data-driven and equitable environmental stewardship. However, practical challenges such as long-term energy autonomy, robust sensor calibration, and seamless data integration across platforms persist, warranting further investigation to expand the impact and reliability of these technologies.

6.2. Research Scope and Limitations

While this thesis breaks new ground in pervasive sensing and autonomous environmental monitoring, several limitations must be acknowledged. Recognizing these boundaries is essential for contextualizing the results and guiding future research and commercialization efforts. The limitations of this research can be categorized into general systemic constraints and the specific technical boundaries of the three proposed systems.

6.2.1. General Limitations

Across all deployments, field operations frequently faced constraints regarding energy availability, which limited continuous, long-term operation in adverse environments. While the evaluations involved both controlled laboratory trials

and targeted field deployments, these relatively small-scale setups may not fully capture the extreme variability and stochastic events encountered in global, large-scale ecological scenarios. Additionally, environmental hardware is inherently susceptible to sensor drift, biofouling, and calibration loss over time, occasionally posing data reliability concerns. Finally, the integration of multimodal data from heterogeneous sources highlighted ongoing challenges in network synchronization. The prototype systems developed herein, although fully functional and rigorously validated, remain academic proofs-of-concept; they are several steps away from fully commercialized, standardized solutions and would significantly benefit from industrial collaboration, lifecycle analysis, and multi-year longitudinal testing.

6.2.2. Specific Limitations of the Proposed Systems

Beyond these general systemic challenges, the three core contributions of this thesis possess specific operational boundaries:

MIDAS: While MIDAS successfully demonstrates material characterization via thermal dissipation from human touch, its performance is heavily dependent on ambient thermal conditions. Extreme ambient temperatures—either very hot or very cold environments—can mask the transient heat transfer signatures, reducing classification accuracy. Furthermore, as validated in our experiments, the system’s sensing capability is currently limited to objects up to 2 mm thick. The accuracy also decreases when analyzing complex composite materials or heavily layered objects, as the thermal model primarily assumes homogenous material dissipation.

LIZARD: LIZARD’s integration of thermal and optical sensing onto autonomous ground drones is constrained by the physical limitations of the UAV/UGV platforms, particularly battery life and payload capacity, which restrict the geographic coverage of a single deployment. Optically, the system’s microplastic detection accuracy (currently up to 80%) is highly sensitive to ambient lighting conditions and surface occlusions; dense vegetation, mud, or deep water can easily obscure plastic debris from the camera’s view. Thermally, the detection of larger plastics relies on the materials absorbing and dissipating environmental heat, meaning its efficacy drops during overcast days, at night, or in highly shaded areas.

Micro-clouds: The deployment of Micro-clouds using repurposed commercial-off-the-shelf (COTS) smartphones introduces unique hardware limitations. While mitigating electronic waste, aging lithium-ion batteries in discarded smartphones present unpredictable degradation curves and potential safety risks when sealed in high-pressure underwater housings. Furthermore, the maximum operational depth is strictly dictated by the physical integrity of these custom enclosures. Computationally, while the cluster successfully processes data locally, the system is still bottlenecked by the inherent physical limitations of underwater communication. Standard radio frequency protocols (like Wi-Fi or Bluetooth) suffer from severe attenuation in aquatic environments, meaning that node-to-node synchronization and data offloading remain constrained by narrow bandwidths and high latency.

6.3. Future Directions

Building on these contributions, there are numerous paths for further research and innovation:

Energy Autonomy and Efficiency: Advances in energy harvesting—such as solar, kinetic, piezoelectric, or bio-inspired methods—are critical to enabling long-term, self-sustaining sensor operations in remote and harsh environments. Developing intelligent hybrid energy management systems that can dynamically switch between multiple energy sources and optimize power consumption through adaptive duty-cycling will significantly prolong device lifespans and reduce human intervention for maintenance and battery replacement. Further exploration of low-power electronics and energy-efficient communication protocols is essential to maximize operational efficiency across diverse deployments.

Enhanced Sensing and Calibration: The integration of emerging materials like flexible electronics, nanomaterials, and bio-compatible sensors promises to enhance sensing accuracy and durability under extreme conditions. Coupling these with AI-driven adaptive calibration techniques can compensate for sensor drift and environmental noise, ensuring consistent data quality over extended deployment periods. Such systems can learn from contextual data to self-correct, reducing the need for regular manual recalibration and increasing robustness in dynamic, real-world ecosystems.

Expanding Application Domains: Future research should broaden the scope of pervasive sensing beyond traditional monitoring to include predictive modeling for climate change impacts, disaster risk forecasting, and biodiversity trend analysis. Employing cross-disciplinary approaches that combine environmental science with social science, economics, and policy will ensure that technological solutions address real-world complexities and promote equitable access across socio-economic strata. Innovative applications may encompass precision conservation, sustainable resource management, and community-driven environmental governance.

Data Integration and AI-Enabled Insights: Establishing standardized interoperability protocols will facilitate seamless data exchange across heterogeneous sensor networks and platforms, enabling richer, multi-source environmental datasets. Advanced data fusion techniques that integrate diverse modalities—such as acoustic, visual, chemical, and thermal sensing—combined with edge AI and federated learning, will empower on-device real-time analytics while preserving privacy. These developments will create comprehensive situational awareness and enable proactive, data-driven decision-making in environmental management.

Ethics, Equity, and Privacy: As pervasive environmental sensing proliferates, so too do ethical considerations regarding data privacy, surveillance, and sovereignty, particularly in sensitive ecological regions or populated areas. Future work must develop robust frameworks ensuring transparent data governance, informed consent, and equitable participation, addressing disparities in technology access and data ownership. Ethical design principles and partnerships with local communities and

stakeholders will foster trust and social license to operate.

Circular Economy and Sustainability: Embedding circular economy principles into sensor design—emphasizing modularity, repairability, and recyclability—will reduce resource extraction and minimize e-waste generation. Lifecycle assessments should become standard for evaluating environmental footprints from production through disposal. Encouraging device repurposing and the use of eco-friendly materials will promote sustainable supply chains and foster systemic change toward sustainable technology ecosystems.

Autonomous Platforms and Edge Intelligence: Continued advances in autonomous underwater and aerial vehicle technologies, fused with distributed edge intelligence, will enhance the scope and responsiveness of environmental monitoring efforts. Enabling these platforms to perform complex analytics and autonomous decision-making in situ reduces reliance on remote cloud infrastructure and lowers communication burdens. Research should focus on collaborative multi-agent systems, adaptive mission planning, and resilience to environmental uncertainties.

Scalable and Modular Deployments: Future systems should emphasize scalability and modularity, allowing flexible assembly and dynamic reconfiguration of sensing and computing nodes tailored to diverse environmental conditions and mission requirements. This includes developing self-organizing networks capable of adaptive resource allocation and fault tolerance, ensuring robust operation despite node failures or environmental disruptions. Open standards and plug-and-play architectures will facilitate interoperability and ease of deployment.

Long-Term Field Studies and Real-World Validation: Extended longitudinal studies in diverse ecological contexts are vital to validate the efficacy, reliability, and environmental impact of pervasive sensing and autonomous systems. Collaborations with ecological experts, policymakers, and local communities will ensure relevance and uptake of technology. Such real-world deployments will illuminate operational challenges, inform design refinements, and generate rich datasets for advancing environmental science and management.

BIBLIOGRAPHY

- [1] Mohammad Aazam and Eui-Nam Huh. “Dynamic Resource Provisioning Through Fog Micro Datacenter”. In: *IEEE International Conference on Pervasive Computing and Communication Workshops*. IEEE. St. Louis, MO, USA: IEEE, 2015, pp. 105–110. DOI: 10.1109/PERCOMW.2015.7134002.
- [2] Yomna Abdelrahman et al. “Cognitive Heat: Exploring the Usage of Thermal Imaging to Unobtrusively Estimate Cognitive Load”. In: *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1.3 (2017), 33:1–33:20. ISSN: 2474-9567. DOI: 10.1145/3130898.
- [3] Yomna Abdelrahman et al. “Investigation of material properties for thermal imaging-based interaction”. In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 2015, pp. 15–18.
- [4] Yomna Abdelrahman et al. “Stay cool! understanding thermal attacks on mobile-based user authentication”. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 2017, pp. 3751–3763.
- [5] Mark A Acree. “Is there a gender difference in fingerprint ridge density?” In: *Forensic science international* 102.1 (1999), pp. 35–44.
- [6] Tommaso Addabbo, Ada Fort, Marco Mugnaini, Lorenzo Parri, Alessandro Pozzebon, and Valerio Vignoli. “Smart sensing in mobility: a lorawan architecture for pervasive environmental monitoring”. In: *2019 IEEE 5th International forum on Research and Technology for Society and Industry (RTSI)*. IEEE. 2019, pp. 421–426.
- [7] Mutiu Adesina Adegboye, Wai-Keung Fung, and Aditya Karnik. “Recent Advances in Pipeline Monitoring and Oil Leakage Detection Technologies: Principles and Approaches”. In: *Sensors* 19.11 (2019), p. 2548. DOI: <https://doi.org/10.3390/s19112548>.
- [8] Jebin Ahmed, Abhijeet Thakur, and Arun Goyal. “Industrial wastewater and its toxic effects”. In: (2021).
- [9] T Mitchell Aide, Carlos Corrada-Bravo, Marconi Campos-Cerqueira, Carlos Milan, Giovany Vega, and Rafael Alvarez. “Real-time bioacoustics monitoring and automated species identification”. In: *PeerJ* 1 (2013), e103.
- [10] *Air quality in Europe 2022*. EEA Air Quality 2022. 2022.
- [11] Mohamed Akrouf, Amal Feriani, Faouzi Bellili, Amine Mezghani, and Ekram Hossain. “Domain generalization in machine learning models for wireless communications: Concepts, state-of-the-art, and open issues”. In: *IEEE Communications Surveys & Tutorials* 25.4 (2023), pp. 3014–3037.
- [12] Ian F Akyildiz, Tommaso Melodia, and Kaushik R Chowdury. “Wireless multimedia sensor networks: A survey”. In: *IEEE Wireless Communications* 14.6 (2007), pp. 32–39.

- [13] Fawaz Alhazemi. “Sequential Clustering Phases for Environmental Noise Level Monitoring on a Mobile Crowd Sourcing/Sensing Platform”. In: *Sensors* 25.5 (2025), p. 1601.
- [14] Howraa Mehdi Mohammad Ali, Taisir EH El-Gorashi, Ahmed Q Lawey, and Jaafar MH Elmighani. “Future Energy Efficient Data Centers With Disaggregated Servers”. In: *Journal of Lightwave Technology* 35.24 (2017), pp. 5361–5380. DOI: 10.1109/JLT.2017.2767574.
- [15] Anthony L Andradý. “The plastic in microplastics: A review”. In: *Marine pollution bulletin* 119.1 (2017), pp. 12–22.
- [16] Rahmi Arda Aral, Şeref Recep Keskin, Mahmut Kaya, and Murat Hacıömerođlu. “Classification of trashnet dataset based on deep learning models”. In: *2018 IEEE International Conference on Big Data (Big Data)*. 2018, pp. 2058–2062.
- [17] Araujo et al. “Identification of microplastics using Raman spectroscopy: Latest developments and future prospects”. In: *Water research* 142 (2018), pp. 426–440.
- [18] Naveen Kumar Arora, Isha Mishra, and Pranay Arora. “SDG 14: life below water-viable oceans necessary for a sustainable planet”. In: *Environmental Sustainability* 6.4 (2023), pp. 433–439.
- [19] Ane Arregi, Oscar Vegas, Aitana Lertxundi, Ana Silva, Isabel Ferreira, Ainhoa Bereziartua, Maria Teresa Cruz, and Nerea Lertxundi. “Road traffic noise exposure and its impact on health: evidence from animal and human studies—chronic stress, inflammation, and oxidative stress as key components of the complex downstream pathway underlying noise-induced non-auditory health effects”. In: *Environmental Science and Pollution Research* 31.34 (2024), pp. 46820–46839.
- [20] Maria Avgerinou, Paolo Bertoldi, and Luca Castellazzi. “Trends in Data Centre Energy Consumption Under the European Code of Conduct for Data Centre Energy Efficiency”. In: *Energies* 10.10 (2017), p. 1470. DOI: <https://doi.org/10.3390/en10101470>.
- [21] Ali Akbar Babaei et al. “Household recycling knowledge, attitudes and practices towards solid waste management”. In: *Resources, Conservation and Recycling* 102 (2015), pp. 94–100.
- [22] Scott Bainbridge, Damien Eggeling, and Geoff Page. “Lessons From the Field—Two Years of Deploying Operational Wireless Sensor Networks on the Great Barrier Reef”. In: *Sensors* 11.7 (2011), pp. 6842–6855. DOI: <https://doi.org/10.3390/s110706842>.
- [23] Jakob E Bardram. “Clinical Proof-of-Concept—a evaluation method for pervasive healthcare systems”. In: *Proc. of the Tenth International Conference on Ubiquitous Computing*. 2008.
- [24] Ishrat Bashir, Farooq A Lone, Rouf Ahmad Bhat, Shafat A Mir, Zubair A Dar, and Shakeel Ahmad Dar. “Concerns and threats of contamination on

- aquatic ecosystems”. In: *Bioremediation and biotechnology: sustainable approaches to pollution degradation*. Springer, 2020, pp. 1–26.
- [25] Fereshteh Bashiri and Che Rosmani Che Hassan. “Light pollution and its effect on the environment”. In: *International Journal of Fundamental Physical Sciences* 4.1 (2014), pp. 8–12.
- [26] Mathias Basner, Wolfgang Babisch, Adrian Davis, Mark Brink, Charlotte Clark, Sabine Janssen, and Stephen Stansfeld. “Auditory and non-auditory effects of noise on health”. In: *The lancet* 383.9925 (2014), pp. 1325–1332.
- [27] Marc Jayson Baucas and Petros Spachos. “Using cloud and fog computing for large scale IoT-based urban sound classification”. In: *Simulation modelling practice and theory* 101 (2020), p. 102013.
- [28] Sara Meghan Beery. *Where the Wild Things Are: Computer Vision for Global-Scale Biodiversity Monitoring*. California Institute of Technology, 2023.
- [29] Trupti Mayee Behera, Sushanta Kumar Mohapatra, Umesh Chandra Samal, Mohammad S Khan, Mahmoud Daneshmand, and Amir H Gandomi. “I-SEP: An improved routing protocol for heterogeneous WSN for IoT-based environmental monitoring”. In: *IEEE Internet of Things Journal* 7.1 (2019), pp. 710–717.
- [30] Germar H Bernhard, Rachel E Neale, Paul W Barnes, PJ Neale, Richard G Zepp, Stephen R Wilson, Anthony L Andrady, AF Bais, Richard L McKenzie, Pieter J Aucamp, et al. “Environmental effects of stratospheric ozone depletion, UV radiation and interactions with climate change: UNEP Environmental Effects Assessment Panel, update 2019”. In: *Photochemical & Photobiological Sciences* 19.5 (2020), pp. 542–584.
- [31] Subhash Janardhan Bhole. “Global goals and global sustainability”. In: *International journal of environmental research and public health* 13.10 (2016), p. 991.
- [32] Mrinal Kanti Bhowmik et al. “Classification of polar-thermal eigenfaces using multilayer perceptron for human face recognition”. In: *Proceedings of the IEEE international Conference on Industrial and Information Systems*. IEEE. 2008, pp. 1–6.
- [33] Shengjie Bi et al. “Auracle: Detecting eating episodes with an ear-mounted sensor”. In: *Proceedings of ACM IMWUT 2018* 2.3 (2018), pp. 1–27.
- [34] Bianco et al. “Atmospheric transport of micro and nanoplastics and fluorescence detection of particles < 20 μm ”. In: *EGU*. 2021, EGU21–3914.
- [35] Eric Biber. “The problem of environmental monitoring”. In: *U. Colo. L. Rev.* 83 (2011), p. 1.
- [36] Frank Biermann, Norichika Kanie, and Rakhyun E Kim. “Global governance by goal-setting: the novel approach of the UN Sustainable Development Goals”. In: *Current Opinion in Environmental Sustainability* 26 (2017), pp. 26–31.

- [37] Janusz Bobulski and Mariusz Kubanek. “Waste classification system using image processing and convolutional neural networks”. In: *International work-conference on artificial neural networks*. Springer. 2019, pp. 350–361.
- [38] Frederik Zuiderveen Borgesius, Jonathan Gray, and Mireille Van Eechoud. “Open data, privacy, and fair information principles: Towards a balancing framework”. In: *Berkeley Technology Law Journal* 30.3 (2015), pp. 2073–2131.
- [39] Soren Bouma, Matthew DM Pawley, Krista Hupman, and Andrew Gilman. “Individual Common Dolphin Identification Via Metric Embedding Learning”. In: *International Conference on Image and Vision Computing New Zealand (IVCNZ)*. IEEE. New Zealand: IEEE, 2018, pp. 1–6. DOI: 10.1109/IVCNZ.2018.8634778.
- [40] Felix Büsching, Sebastian Schildt, and Lars Wolf. “DroidCluster: Towards Smartphone Cluster Computing—The Streets are Paved With Potential Computer Clusters”. In: *International Conference on Distributed Computing Systems Workshops*. IEEE. Macau, China: IEEE, 2012, pp. 114–117. DOI: 10.1109/ICDCSW.2012.59.
- [41] John Calambokidis, Jay Barlow, Kiirsten Flynn, Elana Dobson, and Gretchen H Steiger. “Update on Abundance, Trends, and Migrations of Humpback Whales Along the US West Coast”. In: *IWC Report SC A 17* (2017), pp. 18–21. URL: https://www.cascadiaresearch.org/files/RS6796_SC_A17_NP_13.pdf.
- [42] Zachary R Caldwell, Brian J Zgliczynski, Gareth J Williams, and Stuart A Sandin. “Reef Fish Survey Techniques: Assessing the Potential for Standardizing Methodologies”. In: *PLOS ONE* 11.4 (2016), pp. 1–14. URL: <https://doi.org/10.1371/journal.pone.0153066>.
- [43] Shuo Cao, Dean Zhao, Xiaoyang Liu, and Yueping Sun. “Real-Time Robust Detector for Underwater Live Crabs Based on Deep Learning”. In: *Computers and Electronics in Agriculture* 172 (2020), p. 105339. URL: <https://doi.org/10.1016/j.compag.2020.105339>.
- [44] Andrea Capponi, Claudio Fiandrino, Burak Kantarci, Luca Foschini, Dzmitry Kliazovich, and Pascal Bouvry. “A survey on mobile crowdsensing systems: Challenges, solutions, and opportunities”. In: *IEEE communications surveys & tutorials* 21.3 (2019), pp. 2419–2465.
- [45] Steven JB Carter, Ian P Bell, Jessica J Miller, and Peter P Gash. “Automated Marine Turtle Photograph Identification Using Artificial Neural Networks, With Application to Green Turtles”. In: *Journal of Experimental Marine Biology and Ecology* 452 (2014), pp. 105–110. URL: <https://doi.org/10.1016/j.jembe.2013.12.010>.

- [46] Helotonio Carvalho. “New WHO global air quality guidelines: more pressure on nations to reduce air pollution levels”. In: *The Lancet Planetary Health* 5.11 (2021), e760–e761.
- [47] Charles J Carver, Zhao Tian, Hongyong Zhang, Kofi M Odame, Alberto Quattrini Li, and Xia Zhou. “AmphiLight: Direct Air-Water Communication with Laser Light”. In: *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. USA: USENIX Association, 2020, pp. 373–388. URL: <https://www.usenix.org/conference/nsdi20/presentation/carver>.
- [48] Nuria Castell, Franck R Dauge, Philipp Schneider, Matthias Vogt, Uri Lerner, Barak Fishbain, David Broday, and Alena Bartonova. “Can commercial low-cost sensor platforms contribute to air quality monitoring and exposure estimates?” In: *Environment international* 99 (2017), pp. 293–302.
- [49] Dominique Chabot and Sarah C Marteinson. “Aerial remote sensing of aquatic microplastic pollution: the state of the science and how to move it forward”. In: *Microplastics* 3.4 (2024), pp. 685–695.
- [50] Victor Chang. “An ethical framework for big data and smart cities”. In: *Technological Forecasting and Social Change* 165 (2021), p. 120559.
- [51] Chunlei Chen, Peng Zhang, Huixiang Zhang, Jiangyan Dai, Yugen Yi, Huihui Zhang, and Yonghui Zhang. “Deep learning on computational-resource-limited platforms: A survey”. In: *Mobile Information Systems* 2020.1 (2020), p. 8454327.
- [52] Siguang Chen, Yimin Zheng, Weifeng Lu, Vijayakumar Varadarajan, and Kun Wang. “Energy-optimal dynamic computation offloading for industrial IoT in fog computing”. In: *IEEE Transactions on Green Communications and Networking* 4.2 (2019), pp. 566–576.
- [53] Mung Chiang and Tao Zhang. “Fog and IoT: An overview of research opportunities”. In: *IEEE Internet of things journal* 3.6 (2016), pp. 854–864.
- [54] Mandar Chitre, Shiraz Shahabudeen, and Milica Stojanovic. “Underwater acoustic communications and networking: Recent advances and future challenges”. In: *Marine technology society journal* 42.1 (2008), pp. 103–116.
- [55] Youngjun Cho et al. “Deep Thermal Imaging: Proximate Material Type Recognition in the Wild through Deep Learning of Spatial Surface Temperature Patterns”. In: *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, 2018, p. 2.
- [56] Dading Chong, Yuexian Zou, and Wenwu Wang. “Multi-channel convolutional neural networks with multi-level feature fusion for environmental sound classification”. In: *International conference on multimedia modeling*. Springer, 2018, pp. 157–168.

- [57] Jakob Dybdal Christensen, Jens Therkelsen, Ivo Georgiev, and Henrik Sand. *Data Centre Opportunities in the Nordics: An Analysis of the Competitive Advantages*. Copenhagen, Denmark: Nordic Council of Ministers, 2018.
- [58] Salvador Climent, Antonio Sanchez, Juan Vicente Capella, Nirvana Meratnia, and Juan Jose Serrano. “Underwater acoustic wireless sensor networks: advances and future trends in physical, MAC and routing layers”. In: *Sensors* 14.1 (2014), pp. 795–833.
- [59] Alastor M Coleby and Eric PM Grist. “Prioritized Area Mapping for Multiple Stakeholders Through Geospatial Modelling: A Focus on Marine Plastics Pollution in Hong Kong”. In: *Ocean & Coastal Management* 171 (2019), pp. 131–141. DOI: <https://doi.org/10.1016/j.ocecoaman.2018.12.021>.
- [60] Jun-Hong Cui, Jiejun Kong, Mario Gerla, and Shengli Zhou. “The Challenges of Building Mobile Underwater Wireless Networks for Aquatic Applications”. In: *IEEE Network* 20.3 (2006), pp. 12–18. DOI: 10.1109/MNET.2006.1637927.
- [61] Ben Cutler, Spencer Fowers, Jeffrey Kramer, and Eric Peterson. “Dunking the Data Center”. In: *IEEE Spectrum* 54.3 (2017), pp. 26–31. DOI: 10.1109/MSPEC.2017.7864753.
- [62] Almir Davis and Hwa Chang. “Underwater Wireless Sensor Networks”. In: *2012 Oceans*. IEEE. Hampton Roads, VA, USA: IEEE, 2012, pp. 1–5. DOI: 10.1109/OCEANS.2012.6405141.
- [63] Matthew Dawkins, Linus Sherrill, Keith Fieldhouse, Anthony Hoogs, Benjamin Richards, David Zhang, Lakshman Prasad, Kresimir Williams, Nathan Lauffenburger, and Gaoang Wang. “An Open-Source Platform for Underwater Image and Video Analytics”. In: *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE. Santa Rosa, CA, USA: IEEE, 2017, pp. 898–906. DOI: 10.1109/WACV.2017.105.
- [64] Alessandra De Marco, Chiara Proietti, Alessandro Anav, Luisella Ciancarella, Ilaria D’Elia, Silvano Fares, Maria Francesca Fornasier, Lina Fusaro, Maurizio Gualtieri, Fausto Manes, et al. “Impacts of air pollution on human and ecosystem health, and implications for the National Emission Ceilings Directive: Insights from Italy”. In: *Environment international* 125 (2019), pp. 320–333.
- [65] Laurent Delauney, Chantal Compere, and Michel Lehaitre. “Biofouling protection for marine environmental sensors”. In: *Ocean Science* 6.2 (2010), pp. 503–511.
- [66] Adrián Delgado, Ciprian Briciu-Burghina, and Fiona Regan. “Antifouling strategies for sensors used in water monitoring: review and future perspectives”. In: *Sensors* 21.2 (2021), p. 389.

- [67] Felix Delgado-Ferro, Jorge Navarro-Ortiz, Natalia Chinchilla-Romero, and Juan Jose Ramos-Munoz. “A lorawan architecture for communications in areas without coverage: Design and pilot trials”. In: *Electronics* 11.5 (2022), p. 804.
- [68] Emre Can Demirors, Jiacheng Shi, Anh Duong, Neil Dave, Raffaele Guida, Bernard Herrera, Flavius Pop, Guofeng Chen, Cristian Casella, Sayed Amirhossein Tadayon, et al. “The SEANet Project: Toward a Programmable Internet of Underwater Things”. In: *Underwater Communications and Networking Conference (UComms)*. IEEE. Lerici, Italy: IEEE, 2018, pp. 1–5. DOI: 10.1109/UComms.2018.8493207.
- [69] Hylke van Der Schaaf, Jürgen Moßgraber, Sylvain Grellet, Mickaël Beaufils, Kathi Schleidt, and Thomas Usländer. “An environmental sensor data suite using the ogc sensorthings api”. In: *International Symposium on Environmental Software Systems*. Springer. 2020, pp. 228–241.
- [70] Alim Dewan, Suat U Ay, M Nazmul Karim, and Haluk Beyenal. “Alternative power sources for remote sensors: A review”. In: *Journal of power sources* 245 (2014), pp. 129–143.
- [71] Ashutosh Dhekne et al. “Liquid: A wireless liquid identifier”. In: *Proceedings of ACM MobiSys 2018*. ACM. 2018, pp. 442–454.
- [72] Mario Di Francesco, Sajal K Das, and Giuseppe Anastasi. “Data collection in wireless sensor networks with mobile elements: A survey”. In: *ACM Transactions on Sensor Networks (TOSN)* 8.1 (2011), pp. 1–31.
- [73] Mari Carmen Domingo. “An Overview of the Internet of Underwater Things”. In: *Journal of Network and Computer Applications* 35.6 (2012), pp. 1879–1890. DOI: <https://doi.org/10.1016/j.jnca.2012.07.012>.
- [74] Adam Dou, Vana Kalogeraki, Dimitrios Gunopulos, Taneli Mielikainen, and Ville H Tuulos. “Misco: A MapReduce Framework for Mobile Systems”. In: *International Conference on Pervasive Technologies related to Assistive Environments*. New York, NY, USA: ACM, 2010, p. 32. DOI: <https://doi.org/10.1145/1839294.1839332>.
- [75] Graham Dove, Shiyu Chen, Daniel Fries, Vanessa Johnson, Charlie Mydlarz, Juan Pablo Bello, and Oded Nov. “From environmental monitoring to mitigation action: considerations, challenges, and opportunities for HCI”. In: *Proceedings of the ACM on Human-Computer Interaction* 6.CSCW2 (2022), pp. 1–31.
- [76] Jhalak Dutta and Smita Das. “Enhancing disaster readiness: a comprehensive review of IoT-enabled WSN path planning for effective emergency response”. In: *Cluster Computing* 28.8 (2025), p. 537.
- [77] *Ecosystems and Air Quality*. US EPA Ecosystems Research. 2024.
- [78] *Effects of air pollution on natural ecosystems*. DAERA – Effects of Air Pollution on Ecosystems. 2022.

- [79] Yehia Elkhatib, Barry Porter, Heverson B Ribeiro, Mohamed Faten Zhani, Junaid Qadir, and Etienne Rivière. “On using micro-clouds to deliver the fog”. In: *IEEE Internet Computing* 21.2 (2017), pp. 8–15.
- [80] Emenike et al. “Characterizing Everyday Objects using Human Touch: Thermal Dissipation as a Sensing Modality”. In: *2021 IEEE PerCom*. IEEE. 2021, pp. 1–8.
- [81] Eriksen et al. “Microplastic pollution in the surface waters of the Laurentian Great Lakes”. In: *Marine pollution bulletin* 77.1-2 (2013), pp. 177–182.
- [82] Ismail Essamlali, Hasna Nhaila, and Mohamed El Khaili. “Advances in machine learning and IoT for water quality monitoring: A comprehensive review”. In: *Heliyon* 10.6 (2024).
- [83] Deborah Estrin. “Embedded networked sensing systems: motivations and challenges”. In: *Proceedings of the 2003 joint workshop on Foundations of mobile computing*. 2003, pp. 32–32.
- [84] National Council of Examiners for Engineering. *Fundamentals of Engineering: Supplied-reference Handbook*. Kaplan AEC Engineering, 2003.
- [85] Songchun Fan, Hyojeong Shin, and Romit Roy Choudhury. “Injecting life into toys”. In: *Proceedings of the 15th Workshop on Mobile Computing Systems and Applications*. 2014, pp. 1–6.
- [86] Alessio Fascista. “Toward integrated large-scale environmental monitoring using WSN/UAV/Crowdsensing: A review of applications, signal processing, and future perspectives”. In: *Sensors* 22.5 (2022), p. 1824.
- [87] Fatemeh Fazel, Maryam Fazel, and Milica Stojanovic. “Random Access Compressed Sensing for Energy-efficient Underwater Sensor Networks”. In: *IEEE Journal on Selected Areas in Communications* 29.8 (2011), pp. 1660–1670. DOI: 10.1109/JSAC.2011.110915.
- [88] Perry Neil J Fernandez. “Using a Low-Cost Green Light Emitting Diode to Measure Total Suspended Solids in a Marine Protected Area of the Philippines”. In: *Journal of Coastal Research* 38.1 (2022), pp. 99–113.
- [89] Huber Flores et al. “Evaluating energy-efficiency using thermal imaging”. In: *Proceedings of the 20th International Workshop on Mobile Computing Systems and Applications*. ACM. 2019, pp. 147–152.
- [90] Huber Flores. “The role of micromobility in environmental monitoring: Reflections and opportunities on the use of pervasive sensing”. In: *IEEE Pervasive Computing* 23.2 (2024), pp. 18–26.
- [91] Huber Flores, Naser Hossein Motlagh, Agustin Zuniga, Mohan Liyanage, Monica Passananti, Sasu Tarkoma, Moustafa Youssef, and Petteri Nurmi. “Toward large-scale autonomous marine pollution monitoring”. In: *IEEE Internet of Things Magazine* 4.1 (2021), pp. 40–45.
- [92] Huber Flores, Naser Hossein Motlagh, Agustin Zuniga, Mohan Liyanage, Monica Passananti, Sasu Tarkoma, Moustafa Youssef, and Petteri Nurmi.

- “Toward large-scale autonomous monitoring and sensing of underwater pollutants”. In: *arXiv preprint arXiv:2005.09571* (2020).
- [93] Huber Flores, Agustin Zuniga, Naser Hossein Motlagh, Mohan Liyanage, Monica Passananti, Sasu Tarkoma, Moustafa Youssef, and Petteri Nurmi. “PENGUIN: Aquatic Plastic Pollution Sensing Using AUVs.” In: *DroNet@ MobiSys*. New York, NY, USA: ACM, 2020, pp. 5–1. DOI: <https://doi.org/10.1145/3396864.3399704>.
- [94] Huber Flores, Agustin Zuniga, Leonardo Tonetto, Tristan Braud, Pan Hui, Yong Li, Sasu Tarkoma, Mostafa Ammar, and Petteri Nurmi. “Collaboration Stability: Quantifying the Success and Failure of Opportunistic Collaboration”. In: *IEEE Computer Magazine* 55 (2022), pp. 70–81. URL: [10.1109/MC.2021.3112850](https://doi.org/10.1109/MC.2021.3112850).
- [95] Trygve Olav Fossum, Øystein Sture, Petter Norgren-Aamot, Ingrid Myrnes Hansen, Bjørn Christian Kvisvik, and Anne Christine Knag. “Underwater autonomous mapping and characterization of marine debris in urban water bodies”. In: *arXiv preprint arXiv:2208.00802* (2022).
- [96] Kim Friedman, Peter Bridgewater, Vera Agostini, Tundi Agardy, Salvatore Arico, Frank Biermann, Kate Brown, Ian D Cresswell, Erle C Ellis, Pierre Failler, et al. “The CBD Post-2020 biodiversity framework: People’s place within the rest of nature”. In: *People and Nature* 4.6 (2022), pp. 1475–1484.
- [97] Michael Fulton, Jungseok Hong, Md Jahidul Islam, and Junaed Sattar. “Robotic Detection of Marine Litter Using Deep Visual Detection Models”. In: *International Conference on Robotics and Automation (ICRA)*. Montreal, QC, Canada: IEEE, 2019, pp. 5752–5758. DOI: [10.1109/ICRA.2019.8793975](https://doi.org/10.1109/ICRA.2019.8793975).
- [98] Fabio Galatioto, Margaret Carol Bell, and Graeme Hill. “Understanding the characteristics of the microenvironments in urban street canyons through analysis of pollution measured using a novel pervasive sensor array”. In: *Environmental monitoring and assessment* 186.11 (2014), pp. 7443–7460.
- [99] Kevin J Gaston, Jonathan Bennie, Thomas W Davies, and John Hopkins. “The ecological impacts of nighttime light pollution: a mechanistic appraisal”. In: *Biological reviews* 88.4 (2013), pp. 912–927.
- [100] Bakhram Gaynullin, Christine Hummelgård, Henrik Rödjegård, Claes Mattsson, and Goran Thungstrom. “Environmental monitoring of methane utilizing multispectral NDIR gas sensing for compensation of spectral impact from water vapor in air”. In: *2022 IEEE Sensors*. IEEE, 2022, pp. 1–4.
- [101] S Geetha and SJSW Gouthami. “Internet of things enabled real time water quality monitoring system”. In: *Smart Water* 2.1 (2016), p. 1.
- [102] Julien Gigault, Hind El Hadri, Brian Nguyen, Bruno Grassl, Laura Rowencyk, Nathalie Tufenkji, Siyuan Feng, and Mark Wiesner. “Nanoplastics are

- neither microplastics nor engineered nanoparticles”. In: *Nature nanotechnology* 16.5 (2021), pp. 501–507.
- [103] Francesca Glaviano, Roberta Esposito, Anna Di Cosmo, Francesco Esposito, Luca Gerevini, Andrea Ria, Mario Molinara, Paolo Bruschi, Maria Costantini, and Valerio Zupo. “Management and sustainable exploitation of marine environments through smart monitoring and automation”. In: *Journal of Marine Science and Engineering* 10.2 (2022), p. 297.
- [104] Lonneke Goddijn-Murphy et al. “On thermal infrared remote sensing of plastic pollution in natural waters”. In: *Remote Sensing* 11.18 (2019), p. 2159.
- [105] Michael F Goodchild. “Citizens as sensors: the world of volunteered geography”. In: *GeoJournal* 69.4 (2007), pp. 211–221.
- [106] Andrew S Goudie. *Human impact on the natural environment: Past, present and future*. John Wiley & Sons, 2018.
- [107] Jayavardhana Gubbi, Rajkumar Buyya, Slaven Marusic, and Marimuthu Palaniswami. “Internet of Things (IoT): A vision, architectural elements, and future directions”. In: *Future generation computer systems* 29.7 (2013), pp. 1645–1660.
- [108] Neil Gunningham. “The new collaborative environmental governance: The localization of regulation”. In: *Journal of Law and Society* 36.1 (2009), pp. 145–166.
- [109] Pengtian Guo, Kai Xiao, Xiaohui Wang, and Daoxing Li. “Multi-source heterogeneous data access management framework and key technologies for electric power Internet of Things”. In: *Global energy interconnection* 7.1 (2024), pp. 94–105.
- [110] Unsoo Ha, Yunfei Ma, Zexuan Zhong, Tzu-Ming Hsu, and Fadel Adib. “Learning Food Quality and Safety from Wireless Stickers.” In: *HotNets*. 2018, pp. 106–112.
- [111] Karim Habak, Mostafa Ammar, Khaled A Harras, and Ellen Zegura. “Femto Clouds: Leveraging Mobile Devices to Provide Cloud Service at the Edge”. In: *IEEE International Conference on Cloud Computing*. New York, NY, USA: IEEE, 2015, pp. 9–16. DOI: 10.1109/CLOUD.2015.12.
- [112] Soheyl Faghir Hagh, Parmida Amngostar, Adam Zylka, Maddy Zimmerman, Lauren Cresanti, Spencer Karins, Jarlath Patrick O’Neil-Dunne, Koa Ritz, Clayton J Williams, Ana M Morales-Williams, et al. “Autonomous UAV-mounted LoRaWAN system for real-time monitoring of harmful algal blooms (HABs) and water quality”. In: *IEEE Sensors Journal* 24.7 (2024), pp. 11414–11424.
- [113] Balázs Harangi, Tibor Csordás, and András Hajdu. “Detecting the excessive activation of the ciliaris muscle on thermal images”. In: *2011 IEEE 9th International Symposium on Applied Machine Intelligence and Informatics (SAMI)*. IEEE, 2011, pp. 329–331.

- [114] Jane K Hart and Kirk Martinez. *Environmental sensor networks*. 2005.
- [115] Muhammad Zharfan Hasbullah, Hafizal Mohamad, Ainul Hakim Fizal Sabillah, Ukasyah Mahamod, Khairul Nabilah Zainul Ariffin, and Shah-nurriman Abdul Rahman. “IoT Based Indoor Air and Water Quality Monitoring System Using Node-RED”. In: *2023 9th International Conference on Computer and Communication Engineering (ICCCE)*. IEEE. 2023, pp. 161–166.
- [116] David Hasenfratz, Olga Saukh, Silvan Sturzenegger, Lothar Thiele, et al. “Participatory air pollution monitoring using smartphones”. In: *Mobile Sensing 1* (2012), pp. 1–5.
- [117] Jetmir Haxhibeqiri, Eli De Poorter, Ingrid Moerman, and Jeroen Hoebeke. “A survey of LoRaWAN for IoT: From technology to application”. In: *Sensors* 18.11 (2018), p. 3995.
- [118] Pengwei He, Haibo Feng, Gyan Chhipi-Shrestha, Kasun Hewage, and Rehan Sadiq. “Life Cycle Assessment of e-Waste–Waste Cellphone Recycling”. In: *Electronic Waste: Recycling and Reprocessing for a Sustainable Future* (2022), pp. 231–253.
- [119] John Heidemann, Milica Stojanovic, and Michele Zorzi. “Underwater sensor networks: applications, advances and challenges”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 370.1958 (2012), pp. 158–175. DOI: <https://doi.org/10.1098/rsta.2011.0214>.
- [120] José A Hidalgo-López, Jorge Romero-Sánchez, Raquel Fernández-Ramos, José F Martín-Canales, and José F Ríos-Gómez. “A low-cost, high-accuracy temperature sensor array”. In: *Measurement* 125 (2018), pp. 425–431.
- [121] Thomas J Hofler and Steven L Garrett. “Thermal noise in a fiber optic sensor”. In: *The Journal of the Acoustical Society of America* 84.2 (1988), pp. 471–475.
- [122] Pablo Horstrand, Raúl Guerra, Aythami Rodriguez, Maria Diaz, Sebastián López, and José Fco López. “A UAV platform based on a hyperspectral sensor for image capturing and on-board processing”. In: *IEEE Access* 7 (2019), pp. 66919–66938.
- [123] Xiangwang Hou, Zhiyuan Ren, Jingjing Wang, Shuya Zheng, Wenchi Cheng, and Hailin Zhang. “Distributed fog computing for latency and reliability guaranteed swarm of drones”. In: *IEEE access* 8 (2020), pp. 7117–7130.
- [124] Huang et al. “Lili: liquor quality monitoring based on light signals”. In: *Proceedings of the 27th Annual Mobicom*. 2021, pp. 256–268.
- [125] Mikhail Hushchyn and Andrey Ustyuzhanin. “Generalization of change-point detection in time series data based on direct density ratio estimation”. In: *Journal of Computational Science* 53 (2021), p. 101385.

- [126] Zeba Idrees, Zhuo Zou, and Lirong Zheng. “Edge computing based IoT architecture for low cost air pollution monitoring systems: a comprehensive system analysis, design considerations & development”. In: *Sensors* 18.9 (2018), p. 3021.
- [127] *Impacts of air pollution on ecosystems in Europe*. EEA Report 2022. 2022.
- [128] Eduardo Iscar, Corina Barbalata, Nicholas Goumas, and Matthew Johnson-Roberson. “Towards Low Cost, Deep Water AUV Optical Mapping”. In: *OCEANS MTS/IEEE Charleston*. Charleston, SC, USA: IEEE, 2018, pp. 1–6. DOI: 10.1109/OCEANS.2018.8604772.
- [129] Mohammad Jahanbakht, Wei Xiang, Lajos Hanzo, and Mostafa Rahimi Azghadi. “Internet of underwater things and big marine data analytics—a comprehensive survey”. In: *IEEE Communications Surveys & Tutorials* 23.2 (2021), pp. 904–956.
- [130] Jenna R Jambeck, Roland Geyer, Chris Wilcox, Theodore R Siegler, Miriam Perryman, Anthony Andrady, Ramani Narayan, and Kara Lavender Law. “Plastic waste inputs from land into the ocean”. In: *science* 347.6223 (2015), pp. 768–771.
- [131] Imad Jawhar, Nader Mohamed, Jameela Al-Jaroodi, and Sheng Zhang. “An Efficient Framework for Autonomous Underwater Vehicle Extended Sensor Networks for Pipeline Monitoring”. In: *IEEE International Symposium on Robotic and sensors Environments (ROSE)*. Washington, DC, USA: IEEE, 2013, pp. 124–129. DOI: 10.1109/ROSE.2013.6698430.
- [132] Jia-Qian Jiang. “Occurrence of microplastics and its pollution in the environment: A review”. In: *Sustainable production and consumption* 13 (2018), pp. 16–23.
- [133] Jinfang Jiang, Guangjie Han, Chunsheng Zhu, Sammy Chan, and Joel JPC Rodrigues. “A trust cloud model for underwater wireless sensor networks”. In: *IEEE Communications Magazine* 55.3 (2017), pp. 110–116.
- [134] Weiwei Jiang et al. “Probing Sucrose Contents in Everyday Drinks Using Miniaturized Near-Infrared Spectroscopy Scanners”. In: *Proceedings of ACM IMWUT 2019* 3.4 (2019), pp. 1–25.
- [135] Leilei Jin and Hong Liang. “Deep Learning for Underwater Image Recognition in Small Sample Size Situations”. In: *OCEANS 2017-Aberdeen*. Aberdeen, UK: IEEE, 2017, pp. 1–4. DOI: 10.1109/OCEANSE.2017.8084645.
- [136] Lucas N Joppa. “Technology for nature conservation: An industry perspective”. In: *Ambio* 44.Suppl 4 (2015), pp. 522–526.
- [137] Michael I Jordan and Tom M Mitchell. “Machine learning: Trends, perspectives, and prospects”. In: *Science* 349.6245 (2015), pp. 255–260.
- [138] Muneza Kagzi, Vishal Dagar, Nadia Doytch, Deepika Krishnan, and Manisha Raj. “Curbing environmental degradation to balance sustainable de-

- velopment: Evidence from China”. In: *Environmental and Sustainability Indicators* 24 (2024), p. 100465.
- [139] Matthew E Kahn. “The climate change adaptation literature”. In: *Review of Environmental Economics and Policy* (2016).
- [140] Aman Kansal, Jason Hsu, Sadaf Zahedi, and Mani B Srivastava. “Power management in energy harvesting sensor networks”. In: *ACM Transactions on Embedded Computing Systems (TECS)* 6.4 (2007), 32–es.
- [141] Chien-Chi Kao, Yi-Shan Lin, Geng-De Wu, and Chun-Ju Huang. “A Comprehensive Study on the Internet of Underwater Things: Applications, Challenges, and Channel Models”. In: *Sensors* 17.7 (2017), p. 1477. DOI: <https://doi.org/10.3390/s17071477>.
- [142] Federico Karagulian, Maurizio Barbieri, Alexander Kotsev, Laurent Spinelle, Michel Gerboles, Friedrich Lagler, Nathalie Redon, Sabine Crunaire, and Annette Borowiak. “Review of the performance of low-cost sensors for air quality monitoring”. In: *Atmosphere* 10.9 (2019), p. 506.
- [143] Md Rezaul Karim, Md Shajalal, Alexander Graß, Till Döhmen, Sisay Adugna Chala, Alexander Boden, Christian Beecks, and Stefan Decker. “Interpreting black-box machine learning models for high dimensional datasets”. In: *2023 IEEE 10th International Conference on Data Science and Advanced Analytics (DSAA)*. IEEE. 2023, pp. 1–10.
- [144] Mohammed Al-Khafajiy, Thar Baker, Hilal Al-Libawy, Zakaria Maamar, Moayad Aloqaily, and Yaser Jararweh. “Improving Fog Computing Performance Via Fog-2-Fog Collaboration”. In: *Future Generation Computer Systems* 100 (2019), pp. 266–280. DOI: <https://doi.org/10.1016/j.future.2019.05.015>.
- [145] Mohammad-Ali Khalighi, Chadi Gabriel, Tasnim Hamza, Salah Bourenane, Pierre Leon, and Vincent Rigaud. “Underwater Wireless Optical Communication; Recent Advances and Remaining Challenges”. In: *International Conference on Transparent Optical Networks (ICTON)*. Graz, Austria: IEEE, 2014, pp. 1–4. DOI: 10.1109/ICTON.2014.6876673.
- [146] Ruhul Khalil, Mohammad Babar, Tariqullah Jan, and Nasir Saeed. “Towards the Internet of Underwater Things: Recent Developments and Future Challenges”. In: *IEEE Consumer Electronics Magazine* 10 (2020), pp. 32–37. DOI: 10.1109/MCE.2020.2988441.
- [147] M Imran Khan, Tabassam Yasmeen, Mushtaq Khan, Noor Ul Hadi, Muhammad Asif, Muhammad Farooq, and Sami G Al-Ghamdi. “Integrating industry 4.0 for enhanced sustainability: Pathways and prospects”. In: *Sustainable Production and Consumption* 54 (2025), pp. 149–189.
- [148] Simon Klakegg et al. “Assisted Medication Management in Elderly Care Using Miniaturised Near-Infrared Spectroscopy”. In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2.2 (2018), p. 69.

- [149] Roberta L Klatzky, Susan J Lederman, and Victoria A Metzger. “Identifying objects by touch: An “expert system””. In: *Perception & psychophysics* 37.4 (1985), pp. 299–302.
- [150] Benjamin Koger, Adwait Deshpande, Jeffrey T Kerby, Jacob M Graving, Blair R Costelloe, and Iain D Couzin. “Quantifying the movement, behaviour and environmental context of group-living animals using drones and computer vision”. In: *Journal of Animal Ecology* 92.7 (2023), pp. 1357–1371.
- [151] Lian Pin Koh and Serge A Wich. “Dawn of drone ecology: low-cost autonomous aerial vehicles for conservation”. In: *Tropical conservation science* 5.2 (2012), pp. 121–132.
- [152] Andrey N Kokoulin et al. “Convolutional neural networks application in plastic waste recognition and sorting”. In: *IEEE EIConRus*. IEEE. 2018, pp. 1094–1098.
- [153] Christopher M Korpela, Todd W Danko, and Paul Y Oh. “MM-UAV: Mobile manipulating unmanned aerial vehicle”. In: *Journal of Intelligent & Robotic Systems* 65.1 (2012), pp. 93–101.
- [154] Emmanouil Koukoumidis, Dimitrios Lymberopoulos, Karin Strauss, Jie Liu, and Doug Burger. “Pocket Cloudlets”. In: *ACM SIGARCH Computer Architecture News*. Vol. 39. New York, NY, USA: ACM, 2011, pp. 171–184. DOI: <https://doi.org/10.1145/1961296.1950387>.
- [155] Bridget M Kuehn. “WHO: More than 7 million air pollution deaths each year”. In: *Jama* 311.15 (2014), pp. 1486–1486.
- [156] Prashant Kumar, Andrew Garmory, Matthias Ketzel, Ruwim Berkowicz, and Rex Britter. “Comparative study of measured and modelled number concentrations of nanoparticles in an urban street canyon”. In: *Atmospheric Environment* 43.4 (2009), pp. 949–958.
- [157] Prashant Kumar, Lidia Morawska, Claudio Martani, George Biskos, Marina Neophytou, Silvana Di Sabatino, Margaret Bell, Leslie Norford, and Rex Britter. “The rise of low-cost sensing for managing air pollution in cities”. In: *Environment international* 75 (2015), pp. 199–205.
- [158] Sachin Kumar, Prayag Tiwari, and Mikhail Zymbler. “Internet of Things is a revolutionary approach for future technology enhancement: a review”. In: *Journal of Big data* 6.1 (2019), pp. 1–21.
- [159] Sudeshna Kumar, Sumitro Bhaumik, and Haimanti Banerji. “Methodology for framing indicators for assessing economic-socio-cultural sustainability of the neighbourhood level urban communities in Indian megacities: evidence from Kolkata”. In: *Social Indicators Research* 154.2 (2021), pp. 511–544.
- [160] V Vinoth Kumar and G Sasikala. “Internet of Things (IoT) Enabled Air Quality Monitoring System for Conventional and UAV Application.” In: *Nature Environment & Pollution Technology* 21.1 (2022).

- [161] Rameez Raja Kureshi, Dhavalkumar Thakker, Bhupesh Kumar Mishra, and Jo Barnes. “From raising awareness to a Behavioural change: a case study of indoor air quality improvement using IoT and COM-B model”. In: *Sensors* 23.7 (2023), p. 3613.
- [162] Eemil Lagerspetz, Huber Flores, Niko Mäkitalo, Pan Hui, Petteri Nurmi, Sasu Tarkoma, Andrea Passarella, Jörg Ott, Peter Reichl, Marco Conti, et al. “Pervasive Communities in the Internet of People”. In: *International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. Vol. 1. Athens, Greece: IEEE, 2018, pp. 40–45. DOI: 10.1109/PERCOMW.2018.8480273.
- [163] Eemil Lagerspetz, Xin Li, Jonatan Hamberg, Huber Flores, Petteri Nurmi, Nigel Davis, and Sumi Helal. “Pervasive Data Science on the Edge”. In: *IEEE Pervasive Computing* 18 (2019), pp. 40–49. DOI: 10.1109/MPRV.2019.2925600.
- [164] Liu Lanbo, Zhou Shengli, and Cui Jun-Hong. “Prospects and problems of wireless communication for underwater sensor networks”. In: *Wireless Communications and Mobile Computing* 8.8 (2008), pp. 977–994.
- [165] Eric Larson et al. “HeatWave: thermal imaging for surface user interaction”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2011, pp. 2565–2574.
- [166] Gregory Lee, Ralf Gommers, Filip Waselewski, Kai Wohlfahrt, and Aaron O’Leary. “PyWavelets: A Python package for wavelet analysis”. In: *Journal of Open Source Software* 4.36 (2019), p. 1237.
- [167] Hoesung Lee, Katherine Calvin, Dipak Dasgupta, Gerhard Krinner, Aditi Mukherji, Peter Thorne, Christopher Trisos, José Romero, Paulina Aldunce, Ko Barret, et al. “IPCC, 2023: Climate change 2023: Synthesis report, summary for policymakers. Contribution of working groups I, II and III to the sixth assessment report of the intergovernmental panel on climate change [core writing team, h. Lee and j. Romero (eds.)]. IPCC, Geneva, Switzerland.” In: (2023).
- [168] József Lehel and Sathbh Murphy. “Microplastics in the food chain: food safety and environmental aspects”. In: *Reviews of Environmental Contamination and Toxicology Volume 259* (2021), pp. 1–49.
- [169] Alfin Leo, Gino Rinaldi, Ion Stiharu, and Rama Bhat. “Wireless sensing using acoustic signals for measurement of dynamic pressure and temperature in harsh environment”. In: *Sensor Review* 32.2 (2012), pp. 142–148.
- [170] Heather A Leslie, Martin Jm Van Velzen, Sicco H Brandsma, A Dick Vethaak, Juan J Garcia-Vallejo, and Marja H Lamoree. “Discovery and quantification of plastic particle pollution in human blood”. In: *Environment international* 163 (2022), p. 107199.
- [171] Guozhen Li et al. “Skin-inspired quadruple tactile sensors integrated on a robot hand enable object recognition”. In: *Science Robotics* 5.49 (2020).

- [172] He Li, Kaoru Ota, and Mianxiong Dong. “Learning IoT in Edge: Deep Learning for the Internet of Things with Edge Computing”. In: *IEEE Network* 32.1 (2018), pp. 96–101. DOI: 10.1109/MNET.2018.1700202.
- [173] Yinian Liang, Yan Wang, Fangjiong Chen, Hua Yu, Fei Ji, and Yankun Chen. “Automatic detection and unsupervised clustering-based classification of cetacean vocal signals”. In: *Applied Sciences* 15.7 (2025), p. 3585.
- [174] Kaiqiang Lin, Tong Hao, Wuan Zheng, and Wenchao He. “Analysis of LoRa Link Quality for Underwater Wireless Sensor Networks: A Semi-empirical Study”. In: *IEEE Asia-Pacific Microwave Conference (APMC)*. Singapore: IEEE, 2019, pp. 120–122. DOI: 10.1109/APMC46564.2019.9038666.
- [175] Shaoshan Liu, Liangkai Liu, Jie Tang, Bo Yu, Yifan Wang, and Weisong Shi. “Edge computing for autonomous driving: Opportunities and challenges”. In: *Proceedings of the IEEE* 107.8 (2019), pp. 1697–1716.
- [176] Mohan Liyanage, Farooq Dar, Rajesh Sharma, and Huber Flores. “GEESE: Edge computing enabled by UAVs”. In: *Pervasive and Mobile Computing* 72 (2021), p. 101340.
- [177] Hoi-Shing Lo et al. “Field test of beach litter assessment by commercial aerial drone”. In: *Marine Pollution Bulletin* 151 (2020), p. 110823.
- [178] Javier Lorenzo-Navarro et al. “SMACC: a system for microplastics automatic counting and classification”. In: *IEEE Access* 8 (2020), pp. 25249–25261.
- [179] Sheng-Lan Ma, Shao-Fei Jiang, and Jun Li. “Structural damage detection considering sensor performance degradation and measurement noise effect”. In: *Measurement* 131 (2019), pp. 431–442.
- [180] Dario Madeo, Alessandro Pozzebon, Chiara Mocenni, and Duccio Bertoni. “A low-cost unmanned surface vehicle for pervasive water quality monitoring”. In: *IEEE Transactions on Instrumentation and Measurement* 69.4 (2020), pp. 1433–1444.
- [181] Rosalia Maglietta, Rocco Caccioppoli, Emanuele Seller, Stefano Bellomo, Francesca Cornelia Santacesaria, Roberto Colella, Giulia Cipriano, Ettore Stella, Karin Hartman, Carmelo Fanizza, et al. “Convolutional Neural Networks for Risso’s Dolphins Identification”. In: *IEEE Access* 8 (2020), pp. 80195–80206. DOI: 10.1109/ACCESS.2020.2990427.
- [182] TMI Mahlia, BN Taufiq, HH Masjuki, et al. “Correlation between thermal conductivity and the thickness of selected insulation materials for building wall”. In: *Energy and Buildings* 39.2 (2007), pp. 182–187.
- [183] Nicolas Maisonneuve, Matthias Stevens, Maria E Niessen, Peter Hanappe, and Luc Steels. “Citizen noise pollution monitoring”. In: (2009).
- [184] Titti Malmivirta, Jonatan Hamberg, Eemil Lagerspetz, Xin Li, Ella Peltonen, Huber Flores, and Petteri Nurmi. “Hot or not? robust and accurate continuous thermal imaging on flir cameras”. In: *2019 IEEE International*

- Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 2019, pp. 1–9.
- [185] Gonçalo Marques and Rui Pitarma. “A real-time noise monitoring system based on internet of things for enhanced acoustic comfort and occupational health”. In: *IEEE Access* 8 (2020), pp. 139741–139755.
- [186] Cecilia Martin, Stephen Parkes, Qiannan Zhang, Xiangliang Zhang, Matthew F McCabe, and Carlos M Duarte. “Use of unmanned aerial vehicles for efficient beach litter monitoring”. In: *Marine pollution bulletin* 131 (2018), pp. 662–673.
- [187] Ricardo Martins, Joao Borges de Sousa, Renato Caldas, Chiara Petrioli, and John Potter. “SUNRISE project: Porto University Testbed”. In: *Underwater Communications and Networking (UComms)*. Sestri Levante, Italy: IEEE, 2014, pp. 1–5. DOI: 10.1109/UComms.2014.7017143.
- [188] Jonathan McChesney, Nan Wang, Ashish Tanwer, Eyal de Lara, and Blesson Varghese. “DeFog: Fog Computing Benchmarks”. In: *ACM/IEEE Symposium on Edge Computing*. New York, NY, USA: ACM, 2019, pp. 47–58. DOI: <https://doi.org/10.1145/3318216.3363299>.
- [189] Michael J McGrath and Cliodhna Ní Scanail. “Regulations and standards: Considerations for sensor technologies”. In: *Sensor Technologies: Healthcare, Wellness, and Environmental Applications*. Springer, 2014, pp. 115–135.
- [190] Stephen R McNutt. “Seismic monitoring and eruption forecasting of volcanoes: a review of the state-of-the-art and case histories”. In: *Monitoring and mitigation of volcano hazards* (1996), pp. 99–146.
- [191] Assefa M Melesse, Qihao Weng, Prasad S Thenkabail, and Gabriel B Senay. “Remote sensing sensors and applications in environmental resources mapping and modelling”. In: *Sensors* 7.12 (2007), pp. 3209–3241.
- [192] David K Mellinger and Christopher W Clark. “MobySound: A Reference Archive for Studying Automatic Recognition of Marine Mammal Sounds”. In: *Applied Acoustics* 67.11-12 (2006), pp. 1226–1242. DOI: <https://doi.org/10.1016/j.apacoust.2006.06.002>.
- [193] Agustín Merino, Bruno Chávez-Vergara, Josefa Salgado, María T Fonturbel, Felipe Garcia-Oliva, and Jose A Vega. “Variability in the composition of charred litter generated by wildfire in different ecosystems”. In: *Catena* 133 (2015), pp. 52–63.
- [194] Luis Merino, Fernando Caballero, J Ramiro Martínez-de-Dios, Iván Maza, and Aníbal Ollero. “An unmanned aircraft system for automatic forest fire monitoring and measurement”. In: *Journal of Intelligent & Robotic Systems* 65.1 (2012), pp. 533–548.
- [195] Tymoteusz Miller, Irmina Durlík, Ewelina Kostecka, Polina Kozłowska, Adrianna Łobodzińska, Sylwia Sokołowska, and Agnieszka Nowy. “Integrating artificial intelligence agents with the internet of things for enhanced

- environmental monitoring: applications in water quality and climate data”. In: *Electronics* 14.4 (2025), p. 696.
- [196] Georgia Miskell, Jennifer A Salmond, and David E Williams. “Solution to the problem of calibration of low-cost air quality measurement sensors in networks”. In: *ACS sensors* 3.4 (2018), pp. 832–843.
- [197] Nader Mohamed, Imad Jawhar, Jameela Al-Jaroodi, and Liren Zhang. “Sensor network architectures for monitoring underwater pipelines”. In: *Sensors* 11.11 (2011), pp. 10738–10764.
- [198] Roberto Morabito. “A Performance Evaluation of Container Technologies on Internet of Things Devices”. In: *IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. San Francisco, CA, USA: IEEE, 2016, pp. 999–1000. DOI: 10.1109/INFOCOMW.2016.7562228.
- [199] Nils Morozs, Paul D Mitchell, and Roe Diamant. “Scalable Adaptive Networking for the Internet of Underwater Things”. In: *IEEE Internet of Things Journal* 7.10 (2020), pp. 10023–10037. DOI: 10.1109/JIOT.2020.2988621.
- [200] Abdallah Moubayed, Mohamed Sharif, Marco Luccini, Serguei Primak, and Abdallah Shami. “Water leak detection survey: Challenges & research opportunities using data fusion & federated learning”. In: *IEEE Access* 9 (2021), pp. 40595–40611.
- [201] Enda Murphy and Eoin A King. *Environmental noise pollution: Noise mapping, public health, and policy*. Elsevier, 2022.
- [202] Naga Venkata Sudha Rani Nalukurthi, Ismaila Abimbola, Tasneem Ahmed, Iulia Anton, Khurram Riaz, Qusai Ibrahim, Arghadyuti Banerjee, Ananya Tiwari, and Salem Gharbia. “Challenges and opportunities in calibrating low-cost environmental sensors”. In: *Sensors* 24.11 (2024), p. 3650.
- [203] NDGT Nanayakkara, GASY Nanayakkara, KMCM Karunanayaka, CPTK Kumara, and N Karunanayake. “Smart Cities, Secure Data: Navigating the Evolving Landscape of Data Protection Challenges”. In: *Data Protection: The Wake of AI and Machine Learning*. Springer, 2024, pp. 251–285.
- [204] Brian O’Connor, Stephan Bojinski, Claudia Röösl, and Michael E Schaepman. “Monitoring global changes in biodiversity and climate essential as ecological crisis intensifies”. In: *Ecological Informatics* 55 (2020), p. 101033.
- [205] Edel O’Connor, Alan F Smeaton, Noel E O’Connor, and Fiona Regan. “A neural network approach to smarter sensor networks for water quality monitoring”. In: *Sensors* 12.4 (2012), pp. 4605–4632.
- [206] Ziad Obermeyer, Jasmeet K Samra, and Sendhil Mullainathan. “Individual differences in normal body temperature: longitudinal big data analysis of patient records”. In: *Bmj* 359 (2017).

- [207] Arzu Özkara and Dilek Akyıl. “Environmental pollution and the effects of the pollutants on the ecosystem”. In: *Türk bilimsel derlemeler dergisi* 11.2 (2018), pp. 11–17.
- [208] Ivleva Natalia P et al. “Microplastic in aquatic ecosystems”. In: *Ange wandte Chemie International Edition* 56.7 (2017), pp. 1720–1739.
- [209] Mohsen Padervand et al. “Removal of microplastics from the environment. A review”. In: *Environmental Chemistry Letters* 18.3 (2020), pp. 807–828.
- [210] Hossein Parineh, Majid Sarvi, and Saeed Asadi Bagloee. “Acoustic Sensors and Audio Signal Processing in Intelligent Transportation Systems: A Survey”. In: *IEEE Transactions on Intelligent Vehicles* (2024).
- [211] Soohyun Park, Soyi Jung, Haemin Lee, Joongheon Kim, and Jae-Hyun Kim. “Large-scale water quality prediction using federated sensing and learning: A case study with real-world sensing big-data”. In: *Sensors* 21.4 (2021), p. 1462.
- [212] Yiming Peng, Peipei Wu, Amina T Schartup, and Yanxu Zhang. “Plastic Waste Release Caused by COVID-19 and its Fate in the Global Ocean”. In: *Proceedings of the National Academy of Sciences* 118.47 (2021), pp. 1–6. DOI: <https://doi.org/10.1073/pnas.2111530118>.
- [213] Henrique Miguel Pereira, Simon Ferrier, Michele Walters, Gary N Geller, Rob HG Jongman, Robert J Scholes, Michael William Bruford, Neil Brummitt, Stuart HM Butchart, AC Cardoso, et al. “Essential biodiversity variables”. In: *Science* 339.6117 (2013), pp. 277–278.
- [214] AA Periola, AA Alonge, and KA Ogudo. “Architecture and System Design for Marine Cloud Computing Assets”. In: *The Computer Journal* 63.6 (2020), pp. 927–941. DOI: 10.1093/comjnl/bxz169.
- [215] Linn Persson, Anders Arvidson, Mats Lannerstad, Hanna Lindskog, Tim Morrissey, Linda Nilsson, Stacey Noel, and Jacqueline Senyagwa. *Impacts of pollution on ecosystem services for the millennium development goals*. Stockholm Environment Institute., 2022.
- [216] C Arden Pope III and Douglas W Dockery. “Health effects of fine particulate air pollution: lines that connect”. In: *Journal of the air & waste management association* 56.6 (2006), pp. 709–742.
- [217] Kateryna Popova and Tomáš Prošek. “Corrosion monitoring in atmospheric conditions: a review”. In: *Metals* 12.2 (2022), p. 171.
- [218] Joana C Prata et al. “COVID-19 pandemic repercussions on the use and management of plastics”. In: *Environmental Science & Technology* 54.13 (2020), pp. 7760–7765.
- [219] Teemu Pulkkinen et al. “Understanding wifi cross-technology interference detection in the real world”. In: *2020 IEEE 40th ICDCS*. IEEE. 2020, pp. 954–964.
- [220] Tie Qiu, Zhao Zhao, Tong Zhang, Chen Chen, and CL Philip Chen. “Underwater Internet of Things in Smart Ocean: System Architecture and

- Open Issues”. In: *IEEE transactions on industrial informatics* 16.7 (2019), pp. 4297–4307. DOI: 10.1109/TII.2019.2946618.
- [221] Marko Radeta, Nuno J Nunes, Dinarte Vasconcelos, and Valentina Nisi. “Poseidon-Passive-Acoustic Ocean Sensor for Entertainment and Interactive Data-Gathering in Opportunistic Nautical-Activities”. In: *Designing Interactive Systems Conference*. New York, NY, USA: ACM, 2018, pp. 999–1011. DOI: <https://doi.org/10.1145/3196709.3196752>.
- [222] Marko Radeta, Miguel Ribeiro, Dinarte Vasconcelos, Jorge Lopes, Michael Sousa, João Monteiro, and Nuno Jardim Nunes. “SeaMote-Interactive Remotely Operated Apparatus for Aquatic Expeditions”. In: *IFIP Conference on Human-Computer Interaction*. New York, NY, USA: Springer, 2019, pp. 237–248. DOI: https://doi.org/10.1007/978-3-030-29387-1_14.
- [223] Marko Radeta, Miguel Ribeiro, Dinarte Vasconcelos, Hildegardo Noronha, and Nuno Jardim Nunes. “LoRaquatica: Studying Range and Location Estimation Using LoRa and IoT in Aquatic Sensing”. In: *International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. Austin, TX, USA: IEEE, 2020, pp. 1–6. DOI: 10.1109/PerComWorkshops48775.2020.9156088.
- [224] Marko Radeta, Agustin Zuniga, Naser Hossein Motlagh, Mohan Liyanage, Ruben Freitas, Moustafa Youssef, Sasu Tarkoma, Huber Flores, and Petteri Nurmi. “Deep learning and the oceans”. In: *Computer* 55.5 (2022), pp. 39–50.
- [225] Mohamed Shabeer Mohamed Rafi, Mehran Behjati, and Ahmad Sahban Rafsanjani. “for Smart Agriculture: A Comparative Study of LPWAN, 5G, and Hybrid Connectivity Models”. In: *Smart Computing Paradigms: Intelligence and Network Applications: Proceedings of Seventh International Conference on Smart Computing and Informatics (SCI 2025), Volume 1*. Springer Nature, p. 389.
- [226] Md Adnan Rahman, Md Mehedi Hasan Emon, Tahsina Khan, Abu Bakar Abdul Hamid, and Noor Inayah Yaakub. “Eco-Intelligence: AI’s Contribution to a Sustainable World”. In: *AI and Green Technology Applications in Society*. IGI Global Scientific Publishing, 2025, pp. 31–62.
- [227] Md Wahidur Rahman et al. “Intelligent waste management system using deep learning with IoT”. In: *JKSUCI* (2020).
- [228] Avinash Kumar Rai, Sita Ram Bana, Dhruvendra Singh Sachan, and Barinderjit Singh. “Advancing sustainable agriculture: a comprehensive review for optimizing food production and environmental conservation”. In: *Int. J. Plant Soil Sci* 35.16 (2023), pp. 417–425.
- [229] Rajkumar Ramasamy, Rahul Sanmugam Gopi, and Rohini Palanisamy. “E-Waste Management in the digital era: A sustainable computing approach”. In: *IEEE Potentials* (2024).

- [230] Bharti Rana, Yashwant Singh, and Harvinder Singh. “Metaheuristic routing: a taxonomy and energy-efficient framework for internet of things”. In: *IEEE access* 9 (2021), pp. 155673–155698.
- [231] Christopher Rasmussen, Jiayi Zhao, Danielle Ferraro, and Arthur Trembanis. “Deep Census: AUV-based Scallop Population Monitoring”. In: *International Conference on Computer Vision Workshops*. Venice, Italy: IEEE, 2017, pp. 2865–2873. DOI: 10.1109/ICCVW.2017.338.
- [232] Priyanka Rawat, Kamal Deep Singh, Hakima Chaouchi, and Jean Marie Bonnin. “Wireless sensor networks: a survey on recent developments and potential synergies”. In: *The Journal of supercomputing* 68.1 (2014), pp. 1–48.
- [233] Sasank Reddy, Deborah Estrin, and Mani Srivastava. “Recruitment framework for participatory sensing data collections”. In: *International Conference on Pervasive Computing*. Springer. 2010, pp. 138–155.
- [234] M Anline Rejula, S Jasmine Minija, S Gnana Sophia, and J Angel Barakka. “Decentralized water quality classification using federated learning with recurrent neural networks”. In: *Water Quality Research Journal* 60.1 (2025), pp. 135–150.
- [235] Ad JHM Reniers, JH MacMahan, EB Thornton, TP Stanton, M Henriquez, JW Brown, JA Brown, and E Gallagher. “Surf zone surface retention on a rip-channeled beach”. In: *Journal of Geophysical Research: Oceans* 114.C10 (2009).
- [236] Bernd Resch, Manfred Mittlboeck, and Michael Lippautz. “Pervasive monitoring—An intelligent sensor pod approach for standardised measurement infrastructures”. In: *Sensors* 10.12 (2010), pp. 11440–11467.
- [237] SM2781844 Rhind. “Anthropogenic pollutants: a threat to ecosystem sustainability?” In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 364.1534 (2009), pp. 3391–3401.
- [238] Francisca Ribeiro et al. “Accumulation and fate of nano- and micro-plastics and associated contaminants in organisms”. In: *TrAC* 111 (2019), pp. 139–147.
- [239] Mikko Rinta-Homi et al. “How Low Can You Go? Performance Trade-offs in Low-Resolution Thermal Sensors for Occupancy Detection: A Systematic Evaluation”. In: *IMWUT* 5.3 (2021), pp. 1–22.
- [240] Mikko Rinta-Homi, Naser Hossein Motlagh, Agustin Zuniga, Huber Flores, and Petteri Nurmi. “How Low Can You Go? Performance Trade-offs in Low-Resolution Thermal Sensors for Occupancy Detection: A Systematic Evaluation”. In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT)* 5.3 (2021), p. 126.
- [241] Johan Rockström, Will Steffen, Kevin Noone, Åsa Persson, F Stuart Chapin III, Eric Lambin, Timothy M Lenton, Marten Scheffer, Carl Folke, Hans

- Joachim Schellnhuber, et al. “Planetary boundaries: exploring the safe operating space for humanity”. In: *Ecology and society* 14.2 (2009).
- [242] Jon Paul Rodríguez, David A Keith, Kathryn M Rodríguez-Clark, Nicholas J Murray, Emily Nicholson, Tracey J Regan, Rebecca M Miller, Edmund G Barrow, Lucie M Bland, Kaia Boe, et al. “A practical guide to the application of the IUCN Red List of Ecosystems criteria”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 370.1662 (2015), p. 20140003.
- [243] Victoria Ruiz et al. “Automatic image-based waste classification”. In: *International Work-Conference on the Interplay Between Natural and Artificial Computation*. Springer. 2019, pp. 422–431.
- [244] Peter G Ryan. “Entanglement of birds in plastics and other synthetic materials”. In: *Marine Pollution Bulletin* 135 (2018), pp. 159–164.
- [245] Peter G Ryan. “Ingestion of plastics by marine organisms”. In: *Hazardous chemicals associated with plastics in the marine environment*. Springer, 2016, pp. 235–266.
- [246] Peter G Ryan, Charles J Moore, Jan A Van Franeker, and Coleen L Moloney. “Monitoring the abundance of plastic debris in the marine environment”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 364.1526 (2009), pp. 1999–2012.
- [247] Nasir Saeed, Abdulkadir Celik, Tareq Y Al-Naffouri, and Mohamed-Slim Alouini. “Underwater Optical Wireless Communications, Networking, and Localization: A Survey”. In: *Ad Hoc Networks* 94 (2019), p. 101935. DOI: <https://doi.org/10.1016/j.adhoc.2019.101935>.
- [248] Silvia Santini, Benedikt Ostermaier, and Andrea Vitaletti. “First experiences using wireless sensor networks for noise pollution monitoring”. In: *Proceedings of the workshop on Real-world wireless sensor networks*. 2008, pp. 61–65.
- [249] Arjuna Sathiaselan, Adisorn Lertsinsrubtavee, Adarsh Jagan, Prakash Baskaran, and Jon Crowcroft. “Cloudrone: Micro Clouds in the Sky”. In: *Workshop on Micro Aerial Vehicle Networks, Systems, and Applications for Civilian Use*. New York, NY, USA: ACM, 2016, pp. 41–44. DOI: <https://doi.org/10.1145/2935620.2935625>.
- [250] Mahadev Satyanarayanan. “The emergence of edge computing”. In: *Computer* 50.1 (2017), pp. 30–39.
- [251] Mahadev Satyanarayanan, Paramvir Bahl, Ramón Caceres, and Nigel Davies. “The Case for VM-Based Cloudlets in Mobile Computing”. In: *IEEE Pervasive Computing* 8.4 (2009), pp. 14–23. DOI: 10.1109/MPRV.2009.82.
- [252] Stefano Savazzi, Monica Nicoli, Mehdi Bennis, Sanaz Kianoush, and Luca Barbieri. “Opportunities of federated learning in connected, cooperative,

- and automated industrial systems”. In: *IEEE Communications Magazine* 59.2 (2021), pp. 16–21.
- [253] Jeffrey Sayer, Douglas Sheil, Glenn Galloway, Rebecca A Riggs, Gavyn Mewett, Kenneth G MacDicken, Bas Arts, Agni K Boedihartono, James Langston, David P Edwards, et al. “SDG 15: life on land–The Central role of forests in sustainable development”. In: *Sustainable development goals: their impacts on forests and people* (2019), pp. 482–509.
- [254] Patrick Schober, Christa Boer, and Lothar A Schwarte. “Correlation coefficients: appropriate use and interpretation”. In: *Anesthesia & Analgesia* 126.5 (2018), pp. 1763–1768.
- [255] Gabriel Schwartz and Ko Nishino. “Recognizing material properties from images”. In: *IEEE transactions on pattern analysis and machine intelligence* (2019).
- [256] René P Schwarzenbach, Thomas Egli, Thomas B Hofstetter, Urs Von Gunten, and Bernhard Wehrli. “Global water pollution and human health”. In: *Annual review of environment and resources* 35 (2010), pp. 109–136.
- [257] Yongni Shao et al. “Near-infrared spectroscopy for classification of oranges and prediction of the sugar content”. In: *IJFP* 12.3 (2009), pp. 644–658.
- [258] Kanhaiya Sharma, Bindu G Madhavi, Ankur Goyal, Deepak Parashar, Lalit Shrotriya, and Anurag Gupta. “Real-Time Data Analysis, Significance, Architectures and Applications for Informed Decision-Making”. In: *2023 International Conference on Power Energy, Environment & Intelligent Control (PEEIC)*. IEEE. 2023, pp. 298–302.
- [259] Sandhya Sharma, Kazuhiko Sato, and Bishnu Prasad Gautam. “A methodological literature review of acoustic wildlife monitoring using artificial intelligence tools and techniques”. In: *Sustainability* 15.9 (2023), p. 7128.
- [260] Ashay Devidas Shende, Shrikant A Tekade, Arpan Arunrao Deshmukh, Sandeep Prabhudas Tembhurkar, and P Selvakumar. “Water Quality Monitoring and Control With AI”. In: *Smart Water Technology for Sustainable Management in Modern Cities*. IGI Global Scientific Publishing, 2025, pp. 51–80.
- [261] Chenhua Shi, Zhiyuan Ren, Kun Yang, Chen Chen, Hailin Zhang, Yao Xiao, and Xiangwang Hou. “Ultra-low latency cloud-fog computing for industrial internet of things”. In: *2018 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE. 2018, pp. 1–6.
- [262] Weisong Shi, Jie Cao, Quan Zhang, Youhuizi Li, and Lanyu Xu. “Edge computing: Vision and challenges”. In: *IEEE internet of things journal* 3.5 (2016), pp. 637–646.
- [263] Basem Shihada, Osama Amin, Christopher Bainbridge, Seifallah Jar-dak, Omar Alkhazragi, Tien Khee Ng, Boon Ooi, Michael Berumen, and Mohamed-Slim Alouini. “Aqua-Fi: Delivering Internet Underwater Using

- Wireless Optical Networks”. In: *IEEE Communications Magazine* 58.5 (2020), pp. 84–89. DOI: 10.1109/MCOM.001.2000009.
- [264] Michelle Sigler. “The Effects of Plastic Pollution on Aquatic Wildlife: Current Situations and Future Solutions”. In: *Water, Air, & Soil Pollution* 225.11 (2014), pp. 1–9. DOI: <https://doi.org/10.1007/s11270-014-2184-6>.
- [265] Abhimanyu Singh, Pankhuri Aggarwal, and Rahul Arora. “IoT based waste collection system using infrared sensors”. In: *Proceedings of International Conference on Reliability, Infocom Technologies and Optimization*. IEEE, 2016, pp. 505–509.
- [266] Baihaqi Siregar, Ahmad Badril Azmi Nasution, and Fahmi Fahmi. “Integrated pollution monitoring system for smart city”. In: *2016 International conference on ICT For smart society (ICISS)*. IEEE, 2016, pp. 49–52.
- [267] Dan A Smale, Gary A Kendrick, Euan S Harvey, Timothy J Langlois, Renae K Hovey, Kimberly P Van Niel, Kris I Waddington, Lynda M Bellchambers, Matthew B Pember, Russ C Babcock, et al. “Regional-Scale Benthic Monitoring for Ecosystem-based Fisheries Management (EBFM) Using an Autonomous Underwater Vehicle (AUV)”. In: *ICES Journal of Marine Science* 69.6 (2012), pp. 1108–1118. DOI: <https://doi.org/10.1093/icesjms/fss082>.
- [268] John Smith. “IoT-Enabled Wireless Sensor Networks for Environmental Sensing”. In: *American Journal Of Sensor Networks And Wireless Communications* 5.5 (2024), pp. 18–28.
- [269] Masashi Soga and Kevin J Gaston. “Extinction of experience: the loss of human–nature interactions”. In: *Frontiers in Ecology and the Environment* 14.2 (2016), pp. 94–101.
- [270] de Souza Machado et al. “Microplastics as an emerging threat to terrestrial ecosystems”. In: *Global change biology* 24.4 (2018), pp. 1405–1416.
- [271] Swarnalatha Srinivas, P Ranjitha, R Ramya, and G Kumar Narendra. “Investigation of Oceanic Environment Using Large-Scale UWSN and UANETs”. In: *International Conference on Wireless Communications, Networking and Mobile Computing*. Shanghai, China: IEEE, 2012, pp. 1–5. DOI: 10.1109/WiCOM.2012.6478552.
- [272] Jatin Srivastava, Shilpa Kayastha, Sarah Jamil, and Vandana Srivastava. “Environmental perspectives of *Vetiveria zizanioides* (L.) Nash”. In: *Acta Physiologiae Plantarum* 30.4 (2008), pp. 413–417.
- [273] Will Steffen, Katherine Richardson, Johan Rockström, Hans Joachim Schellnhuber, Opha Pauline Dube, Sébastien Dutreuil, Timothy M Lenton, and Jane Lubchenco. “The emergence and evolution of Earth System Science”. In: *Nature Reviews Earth & Environment* 1.1 (2020), pp. 54–63.

- [274] Bruce A Stein. “Tools and Techniques of Biodiversity Inventory”. In: *Precious heritage: The status of biodiversity in the United States* (2000), p. 19.
- [275] Milica Stojanovic and James Preisig. “Underwater Acoustic Communication Channels: Propagation Models and Statistical Characterization”. In: *IEEE Communications Magazine* 47.1 (2009), pp. 84–89. DOI: 10.1109/MCOM.2009.4752682.
- [276] Dan Stowell, Michael D Wood, Hanna Pamuła, Yannis Stylianou, and Hervé Glotin. “Automatic acoustic detection of birds through deep learning: the first bird audio detection challenge”. In: *Methods in Ecology and Evolution* 10.3 (2019), pp. 368–380.
- [277] Zornitsa Stoyanova and Hristina Harizanova. “Impact of agriculture on water pollution”. In: *AGROFOR International Journal* 4.1 (2019), pp. 111–118.
- [278] Kay Strama, Daniel Weber, and Helge Renkewitz. “Evaluation of Wifi Data Transmission Algorithms for Short Distance Underwater Communication”. In: *OCEANS 2021: San Diego–Porto*. San Diego, CA, USA: IEEE, 2021, pp. 1–6. DOI: 10.23919/OCEANS44145.2021.9705847.
- [279] Jian Su et al. “Design of Litter Collection Robot for Urban Environment”. In: *Intl. Mechanical Engineering Congress and Exposition*. Vol. 85611. ASME. 2021, V07AT07A018.
- [280] Jérôme Sueur, Almo Farina, Amandine Gasc, Nadia Pieretti, and Sandrine Pavoine. “Acoustic indices for biodiversity assessment and landscape investigation”. In: *Acta Acustica united with Acustica* 100.4 (2014), pp. 772–781.
- [281] Guanghao Sun et al. “A portable screening system for onboard entry screening at international airports using a microwave radar, reflective photo sensor and thermography”. In: *Proceedings of International Conference on Instrumentation, Communications, Information Technology, and Biomedical Engineering*. IEEE. 2011, pp. 107–110.
- [282] Jennifer Switzer, Gabriel Marcano, Ryan Kastner, and Pat Pannuto. “Junkyard computing: Repurposing discarded smartphones to minimize carbon”. In: *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*. 2023, pp. 400–412.
- [283] *TACO - Open image dataset of waste in the wild*. <http://tacodataset.org/>. 2022.
- [284] Maen Takruri, Khalid Aboura, and Subhash Challa. “Distributed recursive algorithm for auto calibration in drift aware wireless sensor networks”. In: *Innovations and advanced techniques in systems, computing sciences and software engineering*. Springer, 2008, pp. 21–25.

- [285] Liu Tang et al. “Non-invasive blood glucose monitoring technology: a review”. In: *Sensors* 20.23 (2020), p. 6925.
- [286] Jaime W Thompson, Victoria H Zero, Lori H Schwacke, Todd R Speakman, Brian M Quigley, Jeanine SM Morey, and Trent L McDonald. “finFindR: Computer-assisted Recognition and Identification of Bottlenose Dolphin Photos in R”. In: *bioRxiv* 598 (2019), pp. 1–18. DOI: <https://doi.org/10.1101/825661>.
- [287] Richard C Thompson, Charles J Moore, Frederick S Vom Saal, and Shanna H Swan. “Plastics, the environment and human health: current consensus and future trends”. In: *Philosophical transactions of the royal society B: biological sciences* 364.1526 (2009), pp. 2153–2166.
- [288] Christa Thomsen. “Sustainability (world commission on environment and development definition)”. In: *Encyclopedia of corporate social responsibility*. Springer, 2013, pp. 2358–2363.
- [289] Graham A Tobin et al. “The effectiveness of street sweepers in removing pollutants from road surfaces in Florida”. In: *J Environ Sci.Health* 37.9 (2002), pp. 1687–1700.
- [290] Nguyen H Tran, Wei Bao, Albert Zomaya, Nguyen Minh NH, and Choong Seon Hong. “Federated Learning Over Wireless Networks: Optimization Model Design and Analysis”. In: *IEEE INFOCOM Conference on Computer Communications*. Paris, France: IEEE, 2019, pp. 1387–1395. DOI: [10.1109/INFOCOM.2019.8737464](https://doi.org/10.1109/INFOCOM.2019.8737464).
- [291] Maulana Triananda et al. “Analysis of effect of center of gravity change towards type M1 vehicle stability”. In: *AIP Conference Proceedings*. Vol. 2376. 1. AIP Publishing LLC, 2021, p. 070013.
- [292] Hjalmar K Turesson, Sidarta Ribeiro, Danillo R Pereira, Joao P Papa, and Victor Hugo C de Albuquerque. “Machine learning algorithms for automatic classification of marmoset vocalizations”. In: *PloS one* 11.9 (2016), e0163041.
- [293] Prince O Ukaogo, Ugochukwu Ewuzie, and Chibuzo V Onwuka. “Environmental pollution: causes, effects, and the remedies”. In: *Microorganisms for sustainable environment and health*. Elsevier, 2020, pp. 419–429.
- [294] Thiago Lara Vasques, Pedro Moura, and Aníbal de Almeida. “A Review on Energy Efficiency and Demand Response with Focus on Small and Medium Data Centers”. In: *Energy Efficiency* 12 (2019), pp. 1–30. DOI: <https://doi.org/10.1007/s12053-018-9753-2>.
- [295] J Vellingiri, K Kalaivanan, MP Gopinath, C Gobinath, et al. “Strategies for classifying water quality in the Cauvery River using a federated learning technique”. In: *International Journal of Cognitive Computing in Engineering* 4 (2023), pp. 187–193.
- [296] R Venkatesan, VR Shamji, G Latha, Simi Mathew, RR Rao, Arul Muthiah, and MA Atmanand. “In Situ Ocean Subsurface Time-Series Measurements

- from OMNI Buoy Network in the Bay of Bengal”. In: *Current Science* 1.1 (2013), pp. 1166–1177. URL: <https://www.jstor.org/stable/24092396>.
- [297] Sébastien Villon, David Mouillot, Marc Chaumont, Emily S Darling, Gérard Subsol, Thomas Claverie, and Sébastien Villéger. “A Deep Learning Method for Accurate and Fast Identification of Coral Reef Fishes in Underwater Images”. In: *Ecological informatics* 48 (2018), pp. 238–244. DOI: <https://doi.org/10.1016/j.ecoinf.2018.09.007>.
- [298] Jacob Virtue, Darren Turner, Guy Williams, Stephanie Zeliadt, Matthew McCabe, and Arko Lucieer. “Thermal sensor calibration for unmanned aerial systems using an external heated shutter”. In: *Drones* 5.4 (2021), p. 119.
- [299] Anusha Vishwakarma, Kavita Kanaujia, and Subrata Hait. “Global scenario of E-waste generation: trends and future predictions”. In: *Global E-Waste Management Strategies and Future Implications*. Elsevier, 2023, pp. 13–30.
- [300] Jana Wäldchen and Patrick Mäder. “Machine learning for image based species identification”. In: *Methods in Ecology and Evolution* 9.11 (2018), pp. 2216–2225.
- [301] Jianxiong Wan, Xiang Gui, Shoji Kasahara, Yuanyu Zhang, and Ran Zhang. “Air Flow Measurement and Management for Improving Cooling and Energy Efficiency in Raised-Floor Data Centers: A Survey”. In: *IEEE Access*. 6 (2018), pp. 48867–48901. DOI: <https://doi.org/10.1109/ACCESS.2018.2866840>.
- [302] Jingyang Wang, Xiaolei Li, Lukai Jin, Jiazheng Li, Qiuhong Sun, and Haiyao Wang. “An air quality index prediction model based on CNN-ILSTM”. In: *Scientific Reports* 12.1 (2022), p. 8373.
- [303] Ju Wang, Jie Xiong, Xiaojiang Chen, Hongbo Jiang, Rajesh Krishna Balan, and Dingyi Fang. “TagScan: Simultaneous target imaging and material identification with commodity RFID devices”. In: *Proceedings of ACM MobiCom 2017*. ACM. 2017, pp. 288–300.
- [304] Limin Wang et al. “Birds and plastic pollution: recent advances”. In: *Avian Research* 12.1 (2021), pp. 1–9.
- [305] Lin Wang, Qianyu Li, Quanyi Qiu, Lipeng Hou, Jingyi Ouyang, Ruihan Zeng, Sha Huang, Jing Li, Lina Tang, and Yang Liu. “Assessing the ecological risk induced by PM_{2.5} pollution in a fast developing urban agglomeration of southeastern China”. In: *Journal of Environmental Management* 324 (2022), p. 116284.
- [306] Shiqiang Wang, Tiffany Tuor, Theodoros Salonidis, Kin K Leung, Christian Makaya, Ting He, and Kevin Chan. “Adaptive Federated Learning in Resource Constrained Edge Computing Systems”. In: *IEEE Journal on*

- Selected Areas in Communications* 37.6 (2019), pp. 1205–1221. DOI: 10.1109/JSAC.2019.2904348.
- [307] X Wang and IF Akilydiz. “A survey on sensor networks”. In: *IEEE communication magazine* 40.8 (2002), pp. 102–114.
- [308] Xiaoming Wang, Jianzhong Shang, Zirong Luo, Li Tang, Xiangpo Zhang, and Juan Li. “Reviews of power systems and environmental energy conversion for unmanned underwater vehicles”. In: *Renewable and Sustainable Energy Reviews* 16.4 (2012), pp. 1958–1970.
- [309] Jun-Ichiro Watanabe, Yang Shao, and Naoto Miura. “Underwater and Airborne Monitoring of Marine Ecosystems and Debris”. In: *Journal of Applied Remote Sensing* 13.4 (2019), p. 044509. URL: <https://doi.org/10.1117/1.JRS.13.044509>.
- [310] Mengxi Wei et al. “The study of liquid surface waves with a smartphone camera and an image recognition algorithm”. In: *European Journal of Physics* 36.6 (2015), p. 065026.
- [311] Weigang Wei, Timothy Van Renterghem, Bert De Coensel, and Dick Botteldooren. “Dynamic noise mapping: A map-based interpolation between noise measurements with high temporal resolution”. In: *Applied Acoustics* 101 (2016), pp. 127–140.
- [312] Philip Weinstein. “Human health is harmed by ecosystem degradation, but does intervention improve it? A research challenge from the Millennium Ecosystem Assessment”. In: *EcoHealth* 2.3 (2005), pp. 228–230.
- [313] Gary White et al. “WasteNet: Waste Classification at the Edge for Smart Bins”. In: *arXiv preprint arXiv:2006.05873* (2020).
- [314] Erik Wilhelm, Sandra Siby, Yuren Zhou, Xavier Jayaraj Siddarth Ashok, Melani Jayasuriya, Shaohui Foong, Jacksheeng Kee, Kristin L Wood, and Nils Ole Tippenhauer. “Wearable environmental sensors and infrastructure for mobile large-scale urban deployment”. In: *IEEE Sensors Journal* 16.22 (2016), pp. 8111–8123.
- [315] Hannah J Williams, Lucy A Taylor, Simon Benhamou, Allert I Bijleveld, Thomas A Clay, Sophie de Grissac, Urška Demšar, Holly M English, Novella Franconi, Agustina Gómez-Laich, et al. “Optimizing the Use of Biologgers for Movement Ecology Research”. In: *Journal of Animal Ecology* 89.1 (2020), pp. 186–206. DOI: <https://doi.org/10.1111/1365-2656.13094>.
- [316] Mattis Wolf et al. “Machine learning for aquatic plastic litter detection, classification and quantification (APLASTIC-Q)”. In: *Environmental Research Letters* 15.11 (2020), p. 114042.
- [317] Wayne Wolf. “Cyber-physical systems”. In: *Computer* 42.03 (2009), pp. 88–89.
- [318] Cuebong Wong, Erfu Yang, Xiu-Tian Yan, and Dongbing Gu. “Autonomous robots for harsh environments: a holistic overview of current solutions

- and ongoing challenges”. In: *Systems Science & Control Engineering* 6.1 (2018), pp. 213–219.
- [319] Chenshu Wu et al. “mSense: Towards mobile material sensing with a single millimeter-wave radio”. In: *IMWUT* 4.3 (2020), pp. 1–20.
- [320] Mingyu Wu, Che Fai Yeong, Eileen Lee Ming Su, William Holderbaum, and Chenguang Yang. “A review on energy efficiency in autonomous mobile robots”. In: *Robotic Intelligence and Automation* 43.6 (2023), pp. 648–668.
- [321] Ji-Yan Wu et al. “MRIM: Enabling Mixed-Resolution Imaging for Low-Power Pervasive Vision Tasks”. In: *2022 Percom*. IEEE. 2022, pp. 44–53.
- [322] Q. Xie, M. J. Dousti, and M. Pedram. “Therminator: A thermal simulator for smartphones producing accurate chip and skin temperature maps”. In: *IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*. Aug. 2014, pp. 117–122. DOI: 10 . 1145 / 2627369 . 2627641.
- [323] Jie Xu, Lixing Chen, and Pan Zhou. “Joint Service Caching and Task Offloading for Mobile Edge Computing in Dense Networks”. In: *IEEE Conference on Computer Communications- INFOCOM*. 5. IEEE. Honolulu, HI, USA: IEEE, 2018, pp. 207–215. DOI: 10 . 1109 / INFOCOM . 2018 . 8485977.
- [324] Orhan Yaman, Turker Tuncer, and Beyda Tasar. “DES-Pat: A Novel DES Pattern-based Propeller Recognition Method Using Underwater Acoustical Sounds”. In: *Applied Acoustics* 175 (2021), p. 107859. URL: <https://doi.org/10.1016/j.apacoust.2020.107859>.
- [325] Geng Yang, Li Xie, Matti Mäntysalo, Xiaolin Zhou, Zhibo Pang, Li Da Xu, Sharon Kao-Walter, Qiang Chen, and Li-Rong Zheng. “A health-IoT platform based on the integration of intelligent packaging, unobtrusive biosensor, and intelligent medicine box”. In: *IEEE transactions on industrial informatics* 10.4 (2014), pp. 2180–2191.
- [326] Haichao Yang, Sheng Zhang, Weifeng Ye, Yuzhe Qin, Meng Xu, and Lidong Han. “Emission reduction benefits and efficiency of e-waste recycling in China”. In: *Waste Management* 102 (2020), pp. 541–549.
- [327] Weidong Yang et al. “Wi-Wheat: Contact-free wheat moisture detection with commodity WiFi”. In: *2018 IEEE ICC*. IEEE. 2018, pp. 1–6.
- [328] Yanni Yang et al. “HearLiquid: Non-intrusive Liquid Fraud Detection Using Commodity Acoustic Devices”. In: *IEEE Internet of Things Journal* (2022).
- [329] Siyi Yao, Yongheng Zhang, Peng Wang, Zhipeng Xu, Yongmei Wang, and Youhua Zhang. “Long-term water quality prediction using integrated water quality indices and advanced deep learning models: a case study of Chaohu Lake, China, 2019–2022”. In: *Applied Sciences* 12.22 (2022), p. 11329.

- [330] Hui-Shyong Yeo, Juyoung Lee, Andrea Bianchi, David Harris-Birtill, and Aaron Quigley. “SpeCam: sensing surface color and material with the front-facing camera of a mobile device”. In: *Proceedings of MobileHCI 2017*. ACM. 2017, p. 25.
- [331] Kim Hak Yeong et al. “Application of functional near-infrared spectroscopy to the study of brain function in humans and animal models”. In: *Molecules and cells* 40.8 (2017), p. 523.
- [332] Shanhe Yi, Cheng Li, and Qun Li. “A Survey of Fog Computing: Concepts, Applications and Issues”. In: *Proceedings of the Workshop on Mobile Big Data*. New York, NY, United States: ACM, 2015, pp. 37–42. DOI: <https://doi.org/10.1145/2757384.2757397>.
- [333] Zhigang Yin, Mayowa Olapade, Mohan Liyanage, Farooq Dar, Agustin Zuniga, Naser Hossein Motlagh, Xiang Su, Sasu Tarkoma, Pan Hui, Petteri Nurmi, et al. “Toward city-scale litter monitoring using autonomous ground vehicles”. In: *IEEE Pervasive Computing* 21.3 (2022), pp. 74–83.
- [334] Zhigang Yin, Marko Radeta, Mohan Liyanage, Mayowa Olapade, Abdul-Rasheed Ottun, Agustin Zuniga, Pan Hui, Petteri Nurmi, and Huber Flores. “SNAKE: Harnessing Human Touch for Produce Quality Estimation to Foster Sustainable Retail Practices”. In: *ACM Transactions on Sensor Networks* (2025).
- [335] Huang Yongzhi et al. “A Portable and Convenient System for Unknown Liquid Identification with Smartphone Vibration”. In: *IEEE TMC* (2021).
- [336] Shichao Yue and Dina Katabi. “Liquid Testing with Your Smartphone”. In: *Proceedings of ACM MobiSys 2019*. ACM. 2019, pp. 275–286.
- [337] Zhaoquan Zeng, Shu Fu, Huihui Zhang, Yuhan Dong, and Julian Cheng. “A Survey of Underwater Optical Wireless Communications”. In: *IEEE Communications Surveys & Tutorials* 19.1 (2016), pp. 204–238. DOI: <https://doi.org/10.1109/COMST.2016.2618841>.
- [338] *Zero pollution monitoring assessment*. EEA Zero Pollution Web Report. 2022.
- [339] Chunhua Zhang and John M Kovacs. “The application of small unmanned aerial systems for precision agriculture: a review”. In: *Precision agriculture* 13.6 (2012), pp. 693–712.
- [340] Xie Zhang and Chenshu Wu. “TADAR: Thermal Array-based Detection and Ranging for Privacy-Preserving Human Sensing”. In: *Proceedings of the Twenty-fifth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing*. 2024, pp. 11–20.
- [341] Yu Zhang, Jinfang Zeng, Youming Li, and Da Chen. “Convolutional neural network-gated recurrent unit neural network with feature fusion for environmental sound classification”. In: *Automatic Control and Computer Sciences* 55.4 (2021), pp. 311–318.

- [342] Kan Zheng, Qiang Zheng, Periklis Chatzimisios, Wei Xiang, and Yiqing Zhou. “Heterogeneous vehicular networking: A survey on architecture, challenges, and solutions”. In: *IEEE communications surveys & tutorials* 17.4 (2015), pp. 2377–2396.
- [343] Yu Zheng, Furui Liu, and Hsun-Ping Hsieh. “U-air: When urban air quality inference meets big data”. In: *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. 2013, pp. 1436–1444.
- [344] Zhangbing Zhou, Beibei Yao, Riliang Xing, Lei Shu, and Shengrong Bu. “E-CARP: An Energy Efficient Routing Protocol for UWSNs in the Internet of Underwater Things”. In: *IEEE Sensors Journal* 16.11 (2015), pp. 4072–4082. DOI: 10.1109/JSEN.2015.2437904.
- [345] Pengfei Zhu, Longyin Wen, Xiao Bian, Haibin Ling, and Qinghua Hu. “Vision meets drones: A challenge”. In: *arXiv preprint arXiv:1804.07437* (2018).
- [346] Ziang Zhu et al. “PlasticNet: Deep learning for automatic microplastic recognition via FT-IR spectroscopy”. In: *JCVIS* 6.1 (2020), pp. 1–3.
- [347] Oliver Zielinski, Julia A Busch, Allan D Cembella, Kendra L Daly, Johan Engelbrektsson, Angelos K Hannides, and Heinar Schmidt. “Detecting marine hazardous substances and organisms: sensors for pollutants, toxins, and pathogens”. In: *Ocean Science* 5.3 (2009), pp. 329–349.
- [348] Paula Zuccotti. *Every Thing We Touch: A 24-hour Inventory of Our Lives*. Penguin UK, 2015.
- [349] Agustin Zuniga, Huber Flores, and Petteri Nurmi. “Ripe or rotten? Low-cost produce quality estimation using reflective green light sensing”. In: *IEEE Pervasive Computing* 20.3 (2021), pp. 60–67.
- [350] Agustin Zuniga, Mayowa Olapade, Naser Hossein Motlagh, Mohan Liyanage, Zhigang Yin, Farooq Dar, Ngoc Thi Nguyen, Adeyinka Akintola, Marko Radeta, Sasu Tarkoma, et al. “Low-cost sensing for environmental sustainability”. In: *IEEE Pervasive Computing* (2024).

ACKNOWLEDGEMENTS

First of all, I would like to express my deepest gratitude to *Almighty Allah* for His countless blessings and for granting me the strength and perseverance to complete this journey.

I am profoundly indebted to my **family**. To my *parents*, thank you for your unconditional love, endless support, and for instilling in me the moral values of education. To my *siblings*, I appreciate your constant encouragement. A very special thanks goes to my beloved *wife*, **Dr. Sadaf**, whose patience, sacrifice, and unwavering belief in me have been my rock. To my dear *son*, **Musa**, you are my greatest inspiration.

I would like to express my sincere gratitude to my supervisor, Associate Professor **Huber Flores**. His guidance and insightful feedback were truly instrumental throughout my PhD journey. I am also grateful to Professor **Petteri Nurmi** for his valuable suggestions and continuous support during my doctoral research. Special thanks are due to Professor **Amnir Hadachi**, whose insightful tips and mentorship were invaluable to the successful completion of this thesis.

My sincere thanks go to my lab mates at the **Institute of Computer Science, University of Tartu**, especially *Mohan Liyanage*. Your camaraderie, collaborative spirit, and the stimulating discussions we shared have been a constant source of motivation and have enriched my research experience immeasurably. I am also deeply thankful to the Institute for providing an exceptional environment for learning. To all my colleagues—both those I worked with directly and the familiar faces in the halls and labs—thank you. Your collective presence and contributions made the Institute a truly welcoming and inspiring place.

To all the amazing people I have had the pleasure of meeting in Tartu, thank you for making this city feel like a home away from home. Your friendship and the memories we have created will always be cherished. I want to give a massive thank you to my **friends** who have accompanied me on this journey. I am especially grateful to *Qurat* and *Hamza* for their steadfast friendship and support. Finally, a special mention goes to the **Tartu Wolves Cricket Club** for being a fantastic source of recreation and for making my experience in Tartu truly memorable.

This research has been financed by the **European Social Fund** via the “*ICT programme*” measure.

All content in the thesis is original, and generative AI tools were used solely to help polish the text for improved readability.

SISUKOKKUVÕTE

Autonoomne lausandmetel põhinev sensortehnoloogia keskkonna jätkusuutlikkuse ennetavaks tagamiseks

Keskkonna jätkusuutlikkuse tagamine on tänapäeval üks olulisemaid ülemaailmseid ülesandeid, mis mõjutab nii looduslikke ökosüsteeme kui ka inimühiskonda tervikuna. Ressursside ületarbimine, saastamine, prügi kuhjumine ning elurikkuse vähenemine nõuavad viivitamatult tõhusate tehnoloogiliste lahenduste kasutuselevõttu, et toetada keskkonnaseiret ja ressursside säästlikumat kasutamist. See doktoritöö keskendub uue põlvkonna anduritehnoloogiate ja autonoomsete süsteemide integreerimisele keskkonnaseires. Töö peamine eesmärk on tutvustada ja arendada uusi meetodeid, mis võimaldavad reaalsajas koguda, töödelda ja analüüsida ülimalt detailsusega keskkonnaandmeid, et täiustada keskkonnajuhtimist ningja toetada jätkusuutlikkust.

Töös on välja töötatud kolm olulist süsteemi või kontseptsiooni.

Esiteks: MIDAS – uuenduslik kontaktivaba andurisüsteem erinevate esemete materjali määramiseks soojushajumise põhimõttel. Jälgides inimeste ja esemete vahelist soojusvahetust, suudab MIDAS iseloomustada igapäevaseid esemeid neid kahjustamata ja füüsiliselt puudutamata. Katsed on kinnitanud, et süsteem suudab täpselt tuvastada mitmesuguseid materjale, toimib erinevate kasutajate puhul ja suudab analüüsida mitut eset korraga. Seda tehnoloogiat on võimalik kasutada näiteks jäätmete sortimisel ja taaskasutuse tõhustamisel ning toetada seeläbi toodete keskkonnasäästlikku elutsüklit ja ressursside taaskasutust.

Teiseks: LIZARD – autonoomne prügiseire- ja plastituvastussüsteem, mis ühendab termopildi ja valguspeegelduvuse andurid. LIZARD on loodud plastprügi (sh mikro- ja mesoosakeste) avastamiseks keskkonnas, keskendudes eelkõige traditsiooniliste meetoditega raskesti leitavatele väikestele osakestele. Süsteem töötab energiasäästlikult ning on hõlpsasti integreeritav maapealsete droonidega, võimaldades automaatset, ulatuslikku ja pidevat prügi kaardistamist ilma inimese sekkumiseta. See pakub uusi võimalusi prügisaaste varajaseks avastamiseks ja tõhusaks likvideerimiseks nii linnades kui ka loodusmaastikel, aidates võidelda plasti kuhjumise probleemiga.

Kolmandaks: Micro-cloud ehk mikropilv – veealune hajusate arvutiarvutusvõrgustike (ingl fog computing) kontseptsioon, mis põhineb kommertskasutuseks mõeldud valmisseedmetel (COTS) ja võimaldab andmeid töödelda otse vee all või selle lähedal. See lahendus lubab teha reaalsajas keerukaid analüüse otse andmete kogumiskohas – olgu selleks veealused sensorvõrgud, bioloogilise mitmekesisuse hindamine, veekvaliteedi seire või näiteks torustike turvalisuse jälgimine. MikropilvMicro-cloud on modulaarne, seda on hõlbus transportida (isegi droonide abil) ning see suurendab märkimisväärselt veealuste andmerakenduste töökindlust, skaleeritavust ja energiatõhusust. Selline lähenemine võimaldab laiendada seiret ka piirkondadesse, kus puudub püsiv sidevõrgustik või taristu (nt sisevee-

kogud, süvameri), mis on oluline mitmekesiste keskkonnaprobleemide varajaseks avastamiseks.

Kõigi kolme süsteemi töökindlust on katsetatud nii laboris kui ka reaalses keskkonnas ning tulemused näitavad võrreldes olemasolevate lahendustega märkimisväärselt suuremat täpsust, töökindlust ja energiatõhusust. Töös rõhutatakse ka keskkonnatehnoloogiate jätkusuutlikkust: süsteemide ülesehitus võimaldab vanade seadmete taaskasutust, vähendab elektroonikajäätmete teket ja optimeerib energia-kasutust. Hajusarvutus- ja anduritehnoloogiate kombineerimine võimaldab luua terviklikke ökosüsteeme, milles andurid, servtöötlus, võrgud, pilvandmetöötlus ja kasutajarakendused (koondpaneelid, teavitussüsteemid) loovad andmevoo objektilt lõpplahenduseni.

Kokkuvõttes pakub doktoritöö keskkonnaseireks uue põlvkonna digitaalse alusraamistiku, mis võimaldab kiiret andmetel põhinevat otsustamist ja terviklikku ülevaadet ökosüsteemist ning toetab tõhusat ja jätkusuutlikku keskkonnajuhtimist Eestis ja mujal maailmas.

CURRICULUM VITAE

Personal data

Name: Farooq Ayoub Dar
Date of birth: 15/12/1993
Citizenship: India
E-mail: aaqib4466@gmail.com

Education

2020–2026 University of Tartu, Doctoral studies,
Specialty: Computer Science
2016–2018 Central University of Jammu, Master studies,
Specialty: Computer Science and Information Technology
2011–2015 School of Technology, Islamic University of Science and
Technology, Bachelor studies,
Specialty: Computer Science and Engineering

Experience

2020–2024 University of Tartu, Junior Research Fellow

Teaching experience

Spring 2021 Teaching assistant in: LTAT.06.007 Distributed Systems
Spring 2019 Guest Lecturer: Government Polytechnic College

Academic community service

Year 2023 Student Volunteer (ACM MobiSys)
Year 2018 Organizing Committee (ICRIC 2018)

ELULOOKIRJELDUS

Isikuandmed

Nimi: Farooq Ayoub Dar
Sünniaeg: 15.12.1993
Kodakondsus: India
Meil: aaqib4466@gmail.com

Haridus

2020–2026 Tartu Ülikool, doktoriõpe,
Eriala: informaatika
2016–2018 Jammu Keskülikool, magistriõpe,
Eriala: informaatika ja infotehnoloogia
2011–2015 Islami Teaduse ja Tehnoloogia Ülikooli tehnoloogiakool,
bakalaureuseõpe
Eriala: informaatika ja inseneriteadus

Töökogemus

2020–2024 Tartu Ülikool, nooremteadur

Õpetamiskogemus

Kevad 2021 Õppeassistent aines Hajussüsteemid (LTAT.06.007)
Kevad 2019 Külalislektor Valitsuse Polütehnilises Kolledžis

Akadeemiline kogukonnateenistus

2023 ACM MobiSys, vabatahtlik
2018 ICRIC 2018, korraldustoimkonna liige

**DISSERTATIONES INFORMATICAЕ
PREVIOUSLY PUBLISHED IN
DISSERTATIONES MATHEMATICAE
UNIVERSITATIS TARTUENSIS**

19. **Helger Lipmaa.** Secure and efficient time-stamping systems. Tartu, 1999, 56 p.
22. **Kaili Müürisep.** Eesti keele arvutigrammatika: süntaks. Tartu, 2000, 107 lk.
23. **Varmo Vene.** Categorical programming with inductive and coinductive types. Tartu, 2000, 116 p.
24. **Olga Sokratova.** Ω -rings, their flat and projective acts with some applications. Tartu, 2000, 120 p.
27. **Tiina Puolakainen.** Eesti keele arvutigrammatika: morfoloogiline ühestamine. Tartu, 2001, 138 lk.
29. **Jan Villemson.** Size-efficient interval time stamps. Tartu, 2002, 82 p.
45. **Kristo Heero.** Path planning and learning strategies for mobile robots in dynamic partially unknown environments. Tartu 2006, 123 p.
49. **Härmel Nestra.** Iteratively defined transfinite trace semantics and program slicing with respect to them. Tartu 2006, 116 p.
53. **Marina Issakova.** Solving of linear equations, linear inequalities and systems of linear equations in interactive learning environment. Tartu 2007, 170 p.
55. **Kaarel Kaljurand.** Attempto controlled English as a Semantic Web language. Tartu 2007, 162 p.
56. **Mart Anton.** Mechanical modeling of IPMC actuators at large deformations. Tartu 2008, 123 p.
59. **Reimo Palm.** Numerical Comparison of Regularization Algorithms for Solving Ill-Posed Problems. Tartu 2010, 105 p.
61. **Jüri Reimand.** Functional analysis of gene lists, networks and regulatory systems. Tartu 2010, 153 p.
62. **Ahti Peder.** Superpositional Graphs and Finding the Description of Structure by Counting Method. Tartu 2010, 87 p.
64. **Vesal Vojdani.** Static Data Race Analysis of Heap-Manipulating C Programs. Tartu 2010, 137 p.
66. **Mark Fišel.** Optimizing Statistical Machine Translation via Input Modification. Tartu 2011, 104 p.
67. **Margus Niitsoo.** Black-box Oracle Separation Techniques with Applications in Time-stamping. Tartu 2011, 174 p.
71. **Siim Karus.** Maintainability of XML Transformations. Tartu 2011, 142 p.
72. **Margus Treumuth.** A Framework for Asynchronous Dialogue Systems: Concepts, Issues and Design Aspects. Tartu 2011, 95 p.
73. **Dmitri Lepp.** Solving simplification problems in the domain of exponents, monomials and polynomials in interactive learning environment T-algebra. Tartu 2011, 202 p.

74. **Meelis Kull.** Statistical enrichment analysis in algorithms for studying gene regulation. Tartu 2011, 151 p.
77. **Bingsheng Zhang.** Efficient cryptographic protocols for secure and private remote databases. Tartu 2011, 206 p.
78. **Reina Uba.** Merging business process models. Tartu 2011, 166 p.
79. **Uuno Puus.** Structural performance as a success factor in software development projects – Estonian experience. Tartu 2012, 106 p.
81. **Georg Singer.** Web search engines and complex information needs. Tartu 2012, 218 p.
83. **Dan Bogdanov.** Sharemind: programmable secure computations with practical applications. Tartu 2013, 191 p.
84. **Jevgeni Kabanov.** Towards a more productive Java EE ecosystem. Tartu 2013, 151 p.
87. **Margus Freudenthal.** Simpl: A toolkit for Domain-Specific Language development in enterprise information systems. Tartu, 2013, 151 p.
90. **Raivo Kolde.** Methods for re-using public gene expression data. Tartu, 2014, 121 p.
91. **Vladimir Sor.** Statistical Approach for Memory Leak Detection in Java Applications. Tartu, 2014, 155 p.
92. **Naved Ahmed.** Deriving Security Requirements from Business Process Models. Tartu, 2014, 171 p.
94. **Liina Kamm.** Privacy-preserving statistical analysis using secure multi-party computation. Tartu, 2015, 201 p.
100. **Abel Armas Cervantes.** Diagnosing Behavioral Differences between Business Process Models. Tartu, 2015, 193 p.
101. **Fredrik Milani.** On Sub-Processes, Process Variation and their Interplay: An Integrated Divide-and-Conquer Method for Modeling Business Processes with Variation. Tartu, 2015, 164 p.
102. **Huber Raul Flores Macario.** Service-Oriented and Evidence-aware Mobile Cloud Computing. Tartu, 2015, 163 p.
103. **Tauno Metsalu.** Statistical analysis of multivariate data in bioinformatics. Tartu, 2016, 197 p.
104. **Riivo Talviste.** Applying Secure Multi-party Computation in Practice. Tartu, 2016, 144 p.
108. **Siim Orasmaa.** Explorations of the Problem of Broad-coverage and General Domain Event Analysis: The Estonian Experience. Tartu, 2016, 186 p.
109. **Prastudy Mungkas Fauzi.** Efficient Non-interactive Zero-knowledge Protocols in the CRS Model. Tartu, 2017, 193 p.
110. **Pelle Jakovits.** Adapting Scientific Computing Algorithms to Distributed Computing Frameworks. Tartu, 2017, 168 p.
111. **Anna Leontjeva.** Using Generative Models to Combine Static and Sequential Features for Classification. Tartu, 2017, 167 p.
112. **Mozhgan Pourmoradnasseri.** Some Problems Related to Extensions of Polytopes. Tartu, 2017, 168 p.

113. **Jaak Randmets.** Programming Languages for Secure Multi-party Computation Application Development. Tartu, 2017, 172 p.
114. **Alisa Pankova.** Efficient Multiparty Computation Secure against Covert and Active Adversaries. Tartu, 2017, 316 p.
116. **Toomas Saarsen.** On the Structure and Use of Process Models and Their Interplay. Tartu, 2017, 123 p.
121. **Kristjan Korjus.** Analyzing EEG Data and Improving Data Partitioning for Machine Learning Algorithms. Tartu, 2017, 106 p.
122. **Eno Tõnisson.** Differences between Expected Answers and the Answers Offered by Computer Algebra Systems to School Mathematics Equations. Tartu, 2017, 195 p.

DISSERTATIONES INFORMATICAЕ UNIVERSITATIS TARTUENSIS

1. **Abdullah Makkeh.** Applications of Optimization in Some Complex Systems. Tartu 2018, 179 p.
2. **Riivo Kikas.** Analysis of Issue and Dependency Management in Open-Source Software Projects. Tartu 2018, 115 p.
3. **Ehsan Ebrahimi.** Post-Quantum Security in the Presence of Superposition Queries. Tartu 2018, 200 p.
4. **Ilya Verenich.** Explainable Predictive Monitoring of Temporal Measures of Business Processes. Tartu 2019, 151 p.
5. **Yauhen Yakimenka.** Failure Structures of Message-Passing Algorithms in Erasure Decoding and Compressed Sensing. Tartu 2019, 134 p.
6. **Irene Teinmaa.** Predictive and Prescriptive Monitoring of Business Process Outcomes. Tartu 2019, 196 p.
7. **Mohan Liyanage.** A Framework for Mobile Web of Things. Tartu 2019, 131 p.
8. **Toomas Krips.** Improving performance of secure real-number operations. Tartu 2019, 146 p.
9. **Vijayachitra Modhukur.** Profiling of DNA methylation patterns as biomarkers of human disease. Tartu 2019, 134 p.
10. **Elena Sügis.** Integration Methods for Heterogeneous Biological Data. Tartu 2019, 250 p.
11. **Tõnis Tasa.** Bioinformatics Approaches in Personalised Pharmacotherapy. Tartu 2019, 150 p.
12. **Sulev Reisberg.** Developing Computational Solutions for Personalized Medicine. Tartu 2019, 126 p.
13. **Huishi Yin.** Using a Kano-like Model to Facilitate Open Innovation in Requirements Engineering. Tartu 2019, 129 p.
14. **Faiz Ali Shah.** Extracting Information from App Reviews to Facilitate Software Development Activities. Tartu 2020, 149 p.
15. **Adriano Augusto.** Accurate and Efficient Discovery of Process Models from Event Logs. Tartu 2020, 194 p.
16. **Karim Baghery.** Reducing Trust and Improving Security in zk-SNARKs and Commitments. Tartu 2020, 245 p.
17. **Behzad Abdolmaleki.** On Succinct Non-Interactive Zero-Knowledge Protocols Under Weaker Trust Assumptions. Tartu 2020, 209 p.
18. **Janno Siim.** Non-Interactive Shuffle Arguments. Tartu 2020, 154 p.
19. **Ilya Kuzovkin.** Understanding Information Processing in Human Brain by Interpreting Machine Learning Models. Tartu 2020, 149 p.
20. **Orlenys López Pintado.** Collaborative Business Process Execution on the Blockchain: The Caterpillar System. Tartu 2020, 170 p.
21. **Ardi Tampuu.** Neural Networks for Analyzing Biological Data. Tartu 2020, 152 p.

22. **Madis Vasser.** Testing a Computational Theory of Brain Functioning with Virtual Reality. Tartu 2020, 106 p.
23. **Ljubov Jaanuska.** Haar Wavelet Method for Vibration Analysis of Beams and Parameter Quantification. Tartu 2021, 192 p.
24. **Arnis Parsovs.** Estonian Electronic Identity Card and its Security Challenges. Tartu 2021, 214 p.
25. **Kaido Lepik.** Inferring causality between transcriptome and complex traits. Tartu 2021, 224 p.
26. **Tauno Palts.** A Model for Assessing Computational Thinking Skills. Tartu 2021, 134 p.
27. **Liis Kolberg.** Developing and applying bioinformatics tools for gene expression data interpretation. Tartu 2021, 195 p.
28. **Dmytro Fishman.** Developing a data analysis pipeline for automated protein profiling in immunology. Tartu 2021, 155 p.
29. **Ivo Kubjas.** Algebraic Approaches to Problems Arising in Decentralized Systems. Tartu 2021, 120 p.
30. **Hina Anwar.** Towards Greener Software Engineering Using Software Analytics. Tartu 2021, 186 p.
31. **Veronika Plotnikova.** FIN-DM: A Data Mining Process for the Financial Services. Tartu 2021, 197 p.
32. **Manuel Camargo.** Automated Discovery of Business Process Simulation Models From Event Logs: A Hybrid Process Mining and Deep Learning Approach. Tartu 2021, 130 p.
33. **Volodymyr Leno.** Robotic Process Mining: Accelerating the Adoption of Robotic Process Automation. Tartu 2021, 119 p.
34. **Kristjan Krips.** Privacy and Coercion-Resistance in Voting. Tartu 2022, 173 p.
35. **Elizaveta Yankovskaya.** Quality Estimation through Attention. Tartu 2022, 115 p.
36. **Mubashar Iqbal.** Reference Framework for Managing Security Risks Using Blockchain. Tartu 2022, 203 p.
37. **Jakob Mass.** Process Management for Internet of Mobile Things. Tartu 2022, 151 p.
38. **Gamal Elkoumy.** Privacy-Enhancing Technologies for Business Process Mining. Tartu 2022, 135 p.
39. **Lidia Feklistova.** Learners of an Introductory Programming MOOC: Background Variables, Engagement Patterns and Performance. Tartu 2022, 151 p.
40. **Mohamed Ragab.** Bench-Ranking: A Prescriptive Analysis Approach for Large Knowledge Graphs Query Workloads. Tartu 2022, 158 p.
41. **Mohammad Anagreh.** Privacy-Preserving Parallel Computations for Graph Problems. Tartu 2023, 181 p.
42. **Rahul Goel.** Mining Social Well-being Using Mobile Data. Tartu 2023, 104 p.

43. **Anti Ingel.** Algorithms using information theory: classification in brain-computer interfaces and characterising reinforcement-learning agents. Tartu 2023, 142 p.
44. **Shakshi Sharma.** Fighting Misinformation in the Digital Age: A Comprehensive Strategy for Characterizing, Identifying, and Mitigating Misinformation on Online Social Media Platforms. Tartu 2023, 158 p.
45. **Kristiina Rahkema.** Quality Analysis of iOS Applications with Focus on Maintainability and Security Aspects. Tartu 2023, 182 p.
46. **Ivan Slobozhan.** Studying Online Social Media Engagement in CIS Countries during Protests, Mass Demonstrations and War. Tartu 2023, 81 p.
47. **Nurlan Kerimov.** Building a catalogue of molecular quantitative trait loci to interpret complex trait associations. Tartu 2023, 248 p.
48. **Pavlo Tertychnyi.** Machine Learning Methods for Anti-Money Laundering Monitoring. Tartu 2023, 117 p.
49. **Abasi-amefon Obot Affia.** A Framework and Teaching Approach for IoT Security Risk Management. Tartu 2023, 180 p.
50. **Raimond-Hendrik Tunnel.** Video Game Design and Development Bachelor's Curriculum for Estonia. Tartu 2024, 137 p.
51. **Ahto Salumets.** Bioinformatics analysis of various aspects in immunology. Tartu 2024, 198 p.
52. **Mohammed Abdulhameed Shaif Ali.** Deep Learning Methods for Cell Microscopy Image Analysis. Tartu 2024, 143 p.
53. **Pille Pullonen-Raudvere.** Foundations of Efficient and Secure Algorithm Development for Secure Multiparty Computation. Tartu 2024, 265 p.
54. **Marili Rõõm.** Multiple approaches to learners' success and factors affecting it in computer programming MOOCs. Tartu 2024, 170 p.
55. **Shivananda Rangappa Poojara.** Design and Orchestration of Scalable, Event-Driven Serverless Data Pipelines for Internet of Things (IoT) Applications. Tartu 2024, 172 p.
56. **Hassan Abdulgaleel Hassan Salim Eldeeb.** Empowering Machine Learning Pipelines with Automated Feature Engineering. Tartu 2024, 121 p.
57. **Muhammad Uzair.** Soft decision making for agri-food 4.0. Tartu 2024, 158 p.
58. **Kirill Milintsevich.** Estimation of Depression Level from Text: Symptom-Based Approach, External Knowledge, Dataset Validity. Tartu 2024, 130 p.
59. **Maksym Del.** Multilingual and Multi-Domain Representational Patterns Across Trpansformer-Based Models. Tartu 2024, 131 p.
60. **Kristo Raun.** Adaptive Out-of-order Handling in Streaming Conformance Checking. Tartu 2024, 118 p.
61. **Toivo Vajakas.** Towards integration of mobile network data into analyzing human mobility. Tartu 2024, 103 p.
62. **Katsiaryna Lashkevich.** Data-Driven Analysis and Optimization of Waiting Times in Business Processes. Tartu 2024, 169 p.
63. **Alejandra Duque-Torres.** Classifying, Constraining and Ranking Metamorphic Relations. Tartu 2025, 159 p.

64. **Mariia Bakhtina.** A Method for Information Security and Privacy Management in Smart Solutions. Tartu 2025, 199 p.
65. **Andre Tättar.** Multilingual Machine Translation for Under-Resourced Languages. Tartu 2025, 170 p.
66. **Mahmoud Shoush.** Prescriptive Process Monitoring Under Uncertainty and Resource Constraints. Tartu 2025, 178 p.
67. **Alireza Akhavi Zadegan.** A Multimodal approach for refining Mapping and Localization by Integrating Generative AI and Pedestrian-Centric Data. Tartu 2025, 147 p.
68. **Eerik Muuli.** Automating the assessment and feedback processes in IT teaching – improving creation and maintenance from the teaching staff perspective. Tartu 2025, 196 p.
69. **Kateryna Kubrak.** Towards User-Centered Prescriptive Process Monitoring Systems. Tartu 2025, 151 p.
70. **Zhigang Yin.** Computing and Sensing in a Smart Ring. Tartu 2025, 251 p.
71. **Abdul-Rasheed Olatunji Ottun.** Practical Trustworthy Artificial Intelligence with Human Oversight. Tartu 2025, 239 p.
72. **Sander Mikelsaar.** Analysis and Optimization of Iteratively Decodable Codes. Tartu 2025, 146 p.
73. **Marharyta Domnich.** Advancing Human-Centric Counterfactual Explanations in Explainable AI. Tartu 2025, 210 p.
74. **Viacheslav Komisarenko.** Aligning Training Loss to Evaluation Metrics in Deep Learning. Tartu 2026, 165 p.
75. **Heidi Taveter.** Using Programming-Process Data of Introductory Programming Courses: Finding Solver Types, Giving Feedback, and Detecting Plagiarism. Tartu 2026, 184 p.
76. **Daniel Majoral Lopez.** Deep neural networks for microscopy images. Tartu 2026, 81 p.
77. **Mahir Gulzar.** Addressing Real-world Scenarios via Motion Prediction in Autonomous Driving. Tartu 2026, 141 p.
78. **Hele-Andra Kuulmets.** Cross-Lingual Transfer Learning and Evaluation in Low-Resource Settings. Tartu 2026, 180 p.
79. **Simmo Saan.** Correctness Witnesses for Thread-Modular Program Analysis. Tartu 2026, 251 p.
80. **Tarun Khajuria.** Scene understanding in human and computer vision. Tartu 2026, 151 p.
81. **Modar Sulaiman.** From Data to Fair Decisions: On Ensuring Fairness in Machine Learning Models. Tartu 2026, 209 p.