

Deep Learning and the Oceans

Marko Radeta^{‡¶§}, Agustin Zuniga[†], Naser Hossein Motlagh[†], Mohan Liyanage*,
Ruben Freitas[‡], Moustafa Youssef**, Sasu Tarkoma[†], Huber Flores*, Petteri Nurmi[†]
[‡]Wave Labs, Faculty of Exact Sciences and Engineering, University of Madeira, Portugal
[¶] MARE – Marine and Environmental Sciences Centre, Agência Regional para o Desenvolvimento da
Investigação, Tecnologia e Inovação (ARDITI), Funchal, Madeira, Portugal
[§] Department of Astronomy, Faculty of Mathematics, University of Belgrade, Serbia
[†] Department of Computer Science, University of Helsinki, Finland
* Institute of Computer Science, University of Tartu, Estonia
** American University in Cairo and Alexandria University, Egypt

Abstract—Machine and deep learning (DL) offer significant opportunities for exploring and monitoring oceans and for tackling important problems ranging from litter and oil spill detection to marine biodiversity estimation. Reasonably priced hardware platforms, in the form of autonomous (AUV) and remote operated (ROV) underwater vehicles, are also becoming available, fuelling the growth of data and offering new types of application areas. DL not only supports emerging applications that harness this data but offers support for operating such platforms. This article presents a research vision for DL in the oceans, collating applications and use cases, identifying opportunities, constraints, and open research challenges. We conduct experiments on underwater marine litter detection to demonstrate the benefits DL can bring to underwater environments. Our results show that integrating DL in underwater explorations can automate and scale-up monitoring, and highlight practical challenges in enabling underwater operations. We also provide a research roadmap for the path forward.

Index Terms—Pervasive Computing, Underwater Sensing, Internet of Underwater Things, Marine Litter, Sea, Aquatic

©2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

I. INTRODUCTION

The oceans have long contested the role of being the last frontier for science, and this role holds also for computing research. Indeed, while other difficult to reach environments have recently gained significant research momentum, the same cannot be said for oceans. For example, increasing availability of satellite data and the emergence of diverse space applications have resulted in computing for space gaining momentum whereas estimates suggest that over 90% of the oceans and seabed still remains unexplored [1]. The difficulty of operating computing underwater also means that, even when computing support is available, data from the deep seas remains a scarce commodity [2].

Despite the difficulty in operating computing underwater, novel application areas for the oceans are steadily emerging, paving opportunities for underwater computing platforms to

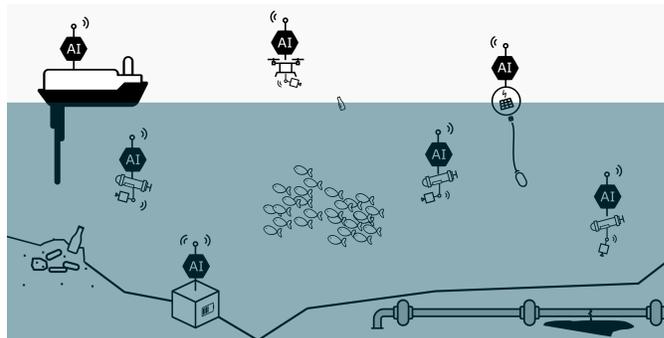


Fig. 1: Deep learning applications for oceans include, underwater pipeline, ship pollution and wildlife monitoring using underwater computing infrastructure.

support these investigations. Marine scientists and oceanographers are increasingly using machine learning (ML) techniques (mostly in the aftermath of their surveys) to analyse underwater video footage, e.g., for estimating biodiversity and the condition of the marine ecosystem [3], whereas oil companies use remote operated underwater vehicles to inspect and monitor pipeline integrity [4]. As these examples also illustrate, computer vision tends to play a central role in many of the emerging underwater applications [5]. This is mainly due to cameras being one of the few sensors that operate unhindered underwater and that are unobtrusive to marine ecosystems. The importance of vision-based data not only offers significant opportunities for harnessing deep learning (DL) but also for applied ML to support deep sea applications. However, it also presents constraints on the platforms that operate it and challenges in ensuring the technique can operate accurately [6]. Currently, the machine learning algorithms typically operate offline, by analysing footage gathered by divers or remote operated vessels [3]. With advances in underwater hardware and DL for constrained devices it should be possible to enable DL to operate directly as part of the underwater operations. For example, autonomous underwater vehicles can integrate DL techniques directly on them, enabling real-time analytics and applications. Such a development would significantly increase the scale of deep sea applications, and offer a pathway to scientific and commercial breakthroughs.

This article presents a research vision for DL in the oceans. The vision, illustrated in Figure 1, integrates DL directly as part of underwater operations to offer timely access to data and insights about the underwater environment. The computing can either be integrated into (autonomous or remotely operated) underwater vehicles (AUVs or ROVs) or equipment carried by people (e.g., dive computers or other portable diving equipment). Realising this vision is currently challenging due to the constraints and challenges underwater environments pose for operating computing, and especially DL. We collate different application areas and hardware and software platforms to provide an overview of the current research landscape, and to identify key constraints and open challenges. We also present proof-of-concept experiments focusing on underwater litter monitoring to demonstrate the feasibility of using DL underwater. The results highlight how DL can significantly scale up underwater monitoring and offer more timely and detailed insights about the state of underwater environments than what is currently possible. The article ends with a reflection on the current research landscape, establishing a roadmap for future explorations.

II. APPLICATIONS FOR UNDERWATER DEEP LEARNING

Underwater computing has traditionally operated as part of dedicated sensor networks [4]. Recent years have seen an increasing shift toward applications that rely on autonomous or remotely operated vehicles [2] and offline analysis of footage collected from dedicated underwater operations [3]. DL can support automated analysis of the data collected through these techniques and, as envisioned in this paper, even operate as part of the underwater vehicles that are responsible for collecting the data [2]. Besides supporting different monitoring applications, DL can also be used to optimize the operation of (autonomous or remote operated) underwater vehicles and other infrastructure such as seabed sensor networks and buoys integrating computing capabilities. Below, we briefly survey and analyse current application domains for DL in the oceans.

Marine sciences are increasingly harnessing underwater data, e.g., to estimate biodiversity or the condition of sensitive marine areas, such as reefs. Currently, the most common approach is to collect video footage through dive surveys and to analyse the data offline. There have been some initiatives to develop dedicated machine learning models, e.g., using Support Vector Machines (SVMs) and Convolutional Neural Networks (CNNs), that can support these tasks. Examples include the use of machine learning for analysing underwater audio signals to discriminate whales and dolphins [7] or the use of DL to identify different fish species [3]. Integrating such algorithms directly onto underwater vehicles (AUVs or ROVs) can scale up these surveys and pave the way to semi-autonomous operations.

Marine litter detection, especially for underwater plastics, is a key environmental sustainability challenge where computing support can bring significant benefits [2]. Learning about marine litter also allows understanding the extent of human activities and helps to limit activities to be compliant with

wildlife preservation standards. In the literature, there are efforts to develop dedicated DL models for identifying litter, but these are mostly trained with images taken on the surface and these models fail to recognise the same objects in underwater environments. The models should be further challenged by exposing them to the real underwater setting where they can be developed as methods to support discriminating individual marine litter objects [8], [9] which are subject to degradation.

Aquaculture and particularly fish farms can benefit from DL, e.g., by automatically detecting and analysing salmon behaviour and classifying fish feeding patterns [10]. Obtaining sufficient amounts of data, however, is challenging. The lack of data makes it difficult to fully realise the benefits of DL and to develop models that could be used to identify abnormal behaviours or the reasons behind it, or to understand any sort of disease in the aquaculture. Modelling abnormal fish behaviour can also reveal information about the environmental conditions, e.g., detect if the water temperature at some locations is higher or below than the desired levels.

Underwater structural monitoring in the form of inspecting underwater pipelines [4], e.g., using magnetic or thermal sensing, enables identifying cracks and potential damage that might be difficult to discern with the naked eye. The most common approach is to navigate ROVs along the pipeline and have remote operators analyse the data. DL can be used to augment and automate such analysis and to support leakage source hunting. This allows coordinating the movements of the underwater vehicles toward the source of the leakage, enabling easier and faster detection of the location of the leakage and reducing the burden on the ROV operators.

AUV and ROV operations can benefit from DL in several ways. First, DL can support enhancing the situational-awareness of AUVs by detecting obstacles and ensuring the AUVs do not collide or otherwise disturb fauna or sensitive flora, such as reef ecosystems. For example, DL can support underwater location-awareness by combining reliable anchor points with dead reckoning to estimate the path traversed by the AUVs. Second, DL can improve navigation and enable autonomous operation for AUVs. For example, AQUA [11] uses the camera of the AUV to assist navigation whereas SOAR [12] supports both collision avoidance and path optimization. Still, most AUVs use surface-based command and control systems to determine their location. This limits the movements of AUVs and the size of the area they can monitor.

Energy Awareness is critical for AUVs and ROVs as underwater movement is generally highly energy-consuming with waves and currents further increasing energy drain. Underwater vehicles rely on energy drainage prediction mechanisms that determine when the system needs to start surfacing to avoid losing the vehicle. DL can be used to build power estimation methods that can take into account a wide range of operating conditions and better determine the remaining operating time of the underwater vehicles.

III. SUBMERSIBLE HARDWARE

Integrating the necessary technologies into underwater vehicles is becoming increasingly possible, yet existing deployments

have predominantly relied on powerful and expensive hardware components with high processing power and on-board memory. These existing solutions also require complex infrastructure support to sustain the operations over a longer period of time. Enabling low-cost platforms to integrate the required technologies and to offer more affordable platforms requires addressing several research challenges. Indeed, accessing and using computing resources underwater is highly challenging and the level of challenge further increases with depth (due to higher pressure and water density). This section reflects on the current state of submersible hardware for marine computing and deep sea exploration whereas the challenges are covered in Section V.

Deep-sea Monitoring. Deep sea explorations traditionally rely on static cable connected deep-sea stations or ocean observatories such as NEPTUNE or the European Multidisciplinary Seafloor Observatory EMSO. These stations are fixed and connected to a central ground control stations with cables that provide high bandwidth real-time communication and power supply between the central and the deep-sea stations. These solutions have limited scale, and are costly to deploy and to operate. The alternative is to rely on a mobile platform. For example, the GEOMAR Modular Lander Systems (GML) [13] is an autonomous instrument carrier system deployed on the seafloor that can work for periods of up to 6 to 12 months without human input. However, these types of platforms are highly costly and offer limited computing capability.

AUVs and ROVs. While there are increasingly affordable off-the-shelf AUVs and ROVs for underwater explorations, their processing tends to be limited to supporting basic and simple routines so that battery life can be preserved. Examples of such platforms include the PowerVision series of underwater vehicles, and the ROVs of BlueRobotics. More advanced platforms that integrate computing support or even necessary tools for running DL are also available (e.g., HippoCampus [14], LoCO [15] or the platform proposed by Cadena et al. [16]) but the costs of these platforms are prohibitively high. Indeed, these platforms are best suited for individual explorations rather than as a large-scale solution – envisioned in this article.

Hybrid Technologies. Another possibility is to combine underwater operations with surface-based cloud access. The main benefit of these systems is the availability of powerful processing power through the cloud but this requires connectivity to a surface-based hub or gateway and limits the maximum depth at which these systems can operate. Underwater communications tend to suffer from many failures and poor bandwidth [17], and marine areas tend to have limited network coverage, which means the benefits of cloud access come at the cost of unreliable communication and high latency. Indeed, these solutions work best close to shore areas and in calm areas as this ensures both the underwater and above-surface communication can operate smoothly.

IV. CASE STUDY: UNDERWATER LITTER DETECTION

Integrating DL directly into underwater platforms has potential for significantly scaling up underwater investigations by offering access to recently gathered data and insights about the

TABLE I: Deep Litter CNN Model Performance.

Class	Accuracy	Precision	Recall	F1-Score	Support
Overall	0.80	0.82	0.80	0.80	137
Metal	0.95	0.95	0.95	0.95	24
Plastic	0.87	0.87	0.87	0.87	45
Wood	0.97	0.97	0.97	0.97	13
Other	0.80	0.83	0.80	0.81	55
Median	0.87	0.87	0.87	0.87	45

state of the underwater environment and by reducing human burden by automating (parts of) the monitoring. This section demonstrates these benefits through a case study on underwater litter monitoring. As part of the experiments we also highlight algorithm and system-level challenges in operating DL robustly in underwater environments.

A. Litter Classification

Setup. We apply DL-based object detection on the Trash-ICRA19 annotated trash dataset [8]. We expand the original labels to include four different litter categories: *plastic*, *wood*, *metal*, and *other*. MobileNetv2 with Single-Shot Detection (SSD) with default hyperparameters was used as the object recognition architecture due to its lightweight structure and good object recognition performance. Training was performed on GPU using TensorFlow 1.15, obtaining the quantized model, using a batch size of 12 and 50k iterations. No further data augmentations were considered in these experiments.

Results. Classification results were obtained by comparing the bounding boxes of the annotations with those obtained with the DL (MobileNet). We calculate the intersection of unions (IoU) for the two bounding boxes and consider the detection successful whenever IoU is at least 50%. The results are shown in Table I for the different litter categories. The median of the classification accuracy is 87% for all litter categories with wood (97%) and metal (95%) being the materials with the highest accuracy. In images containing multiple litter objects, the performance drops to 56% and the overall recall is 60%. Overall, these results show promise in using DL for detecting, and to an extent also classifying, underwater litter. This result is encouraging as it suggests that marine litter detection (and classification) could potentially be automated instead of relying on human observers or waiting for the data collection to finish. However, as we next show, there are also many challenges to extend these systems to operate robustly in the underwater environment.

Challenges. Figure 2 shows selected results and highlights both successes and failures in the detection. Close-ups of individual marine litter categories are generally easy to identify as can be witnessed from the high confidence rates (see the close-ups of plastics, paper (other) and metal in Fig. 2a, 2b and 2c). At longer distances, light and water conditions decrease overall accuracy (Fig. 2d and 2e). The most difficult issues, however, are complex backgrounds (Fig. 2f) and the aggregation of sediment on marine debris (Fig. 2g). Another challenge emerges when multiple objects or object categories need to be supported (Fig. 2h) as this can cause the model to be unable to distinguish the individual objects. This may

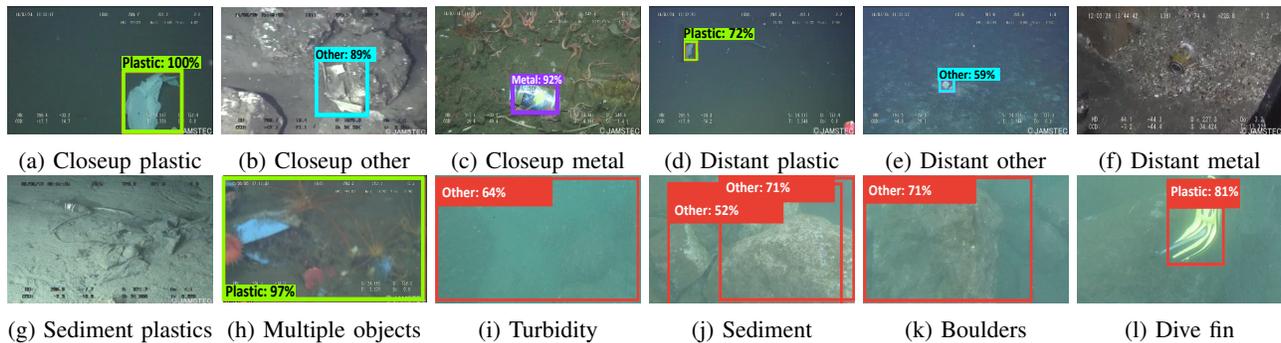


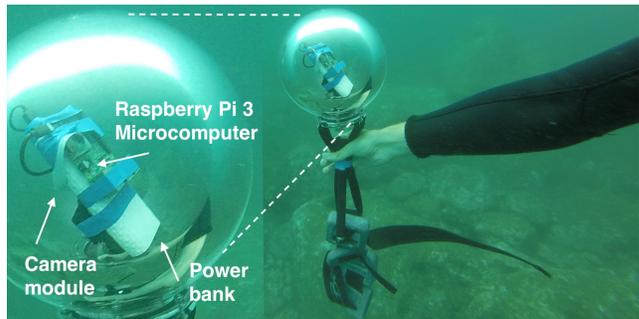
Fig. 2: Marine litter classification on existing footage (a) to (h), and during in-the-wild tests (i) to (l).

require more advanced models, e.g., recently there has been work on Region Based CNNs for marine object detection [9] but these models require higher computational time and can thus degrade AUV power autonomy.

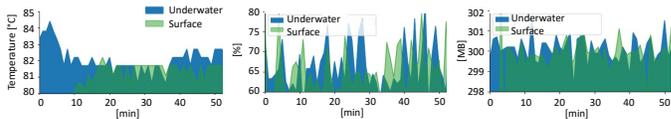
B. Deep Learning in Underwater Environments

Setup. We next consider the feasibility of using DL on underwater platforms by running object recognition on a custom-build IoT platform. Raspberry Pi 3 microcomputer with camera module and power bank were placed in a sealed container (Figure 3a) and transported during a 50 minute dive. Deep learning model was implemented using TensorFlow Lite. As the casing contains air, the deployment is positively buoyant and additional diving weights were used to submerge it. The device is then made to run continuous object recognition at 15 FPS captured by the camera. We monitor the CPU’s temperature, performance and RAM utilisation during the experiment. The same setup is then repeated on the surface including the container to see what effect (if any) the underwater environment has on the hardware performance.

Results. We observe that CPU temperature has an early spike during the underwater operation (Figure 3b), which is expected due to the lack of air outlet inside the container. Once the dive submerges deeper, the colder water starts to offer cooling for the container. In terms of CPU (Figure 3c), the results show that the performance is similar between the underwater environment and the surface. The surface-based experiments present slightly higher CPU rates than the underwater experiments, suggesting that the Raspberry Pi needs to occasionally throttle the CPU to avoid the risk of overheating. A similar behaviour is observed in RAM usage (Figure 3d). Overall the results show that even relatively simple computing platforms, such as Raspberry Pis or other low-cost micro-controllers, could be integrated with off-the-shelf AUVs and ROVs to offer affordable underwater computing platforms. However, these require better casing and cooling as well as advanced energy management solutions to ensure the computations do not hamper the navigation and movement functionality of the underwater vehicles. Additionally, while the computing power of embedded hardware is increasing, it remains better suited for running the models rather than learning them. Thus, there is also a need for interfaces and solutions that allow easily deploying DL models on the underwater platforms.



(a) Deployed microsphere for real-time marine litter detection.



(b) CPU temperature (c) CPU usage (d) RAM usage

Fig. 3: In-the-wild test for litter classification with prototype (a), CPU temperature (Celsius) (b), CPU usage (percentage) (c) and RAM usage (MB) (d).

Challenges. Figure 2 shows selected object recognition results from the in-the-wild tests. The image quality is significantly decreased compared to the standalone tests, but overall the results of object recognition largely mirror those presented in the previous section. Indeed, the results confirm that background complexity derived from the diversity of underwater conditions affects classification accuracy (e.g., turbidity and sediment in Fig. 2i and Fig. 2j). The experiments also show the need to use advance models that differentiate multiple objects in marine environments (e.g., boulders and dive fin in Fig. 2k and Fig. 2l). Besides improving the DL models, overcoming these issues may demand having a higher amount of samples during the training and testing of the model.

V. THE ROAD AHEAD AND CHALLENGES

The results of the experiments demonstrated that underwater DL can automate operations that currently require significant manual effort, thus helping to scale up underwater monitoring. By operating the DL directly underwater, it becomes possible to perform the monitoring in real-time rather than having to wait for the data to be collected and analysed.

Operating the DL directly on the underwater platforms reduces communication bandwidth to relay the data, and can reduce human effort in curating and analysing the data. What makes the integration of DL into underwater platforms challenging, however, is the lack of DL models that can be easily adopted to operate in different underwater environments. Underwater operations are prone to extreme and changing conditions, including strong turbulence, different visibility, and differing water temperatures with particularly those around the freezing point posing challenges for monitoring. Similarly, turbulence (e.g., currents), turbidity, sunlight availability, the presence of marine wildlife, water salinity and other factors all influence the quality of the images. Also, the platforms themselves are subject to challenges as they can be attacked by marine wildlife and are subjected to corrosion and fouling. Below, we outline key research challenges to establish a path for the research roadmap ahead. A summary is also presented in Table II.

A. Technical and Hardware Constraints

To operate in extreme conditions, the effective encasing of resources is critical. The encasing must protect the resources from water damage while at the same time it should be lightweight enough to be easily transported between locations, e.g., using AUVs. We describe below key requirements to preserve the correct operations of different resources underwater.

Underwater sensing. Currently, data sampling is performed mostly using cameras. However, image-based solutions are not sufficient to provide accurate information underwater objects and a richer set of sensors is needed to capture more information about the monitored objects. For instance, a common way to detect material is to rely on spectroscopy, such as Fourier transform infrared (FTIR) spectroscopy and Raman spectroscopy (FTR) [18]. Spectroscopic instruments are difficult to migrate to underwater environments due to being costly and bulky and (typically) requiring the monitored objects to be placed between a sensor and a receptor. As a result, easily portable and lightweight solutions with low energy profile are required for underwater operations.

Available processing resources. Affordable underwater vehicles only come equipped with limited computing power. While there have been efforts to build affordable underwater vehicles that could do heavier processing, these vehicles remain limited. They do not have sufficient processing capability to manipulate streams of sensor data and video footage in real-time nor to run DL – or if they do they are costly and difficult to deploy at scale. As a result, distributed computing underwater is required to have high performance while at the same avoiding an increased processing load for the vehicles. Besides aiding DL execution with distributed computing power, it could be possible to train DL models incrementally underwater in a federated manner. Achieving this requires flexible and transparent cooperation mechanisms for the underwater vehicles.

Fault tolerance design and operational time. Existing underwater vehicles have limited operational time, surviving at most 24 hours between recharges. Once computing resources are submerged in the deep, it is impossible to replace failing,

exhausted or inoperative components without extracting the resources first. As a result, computing tools need to be designed to be robust and capable of taking advantage of different opportunities. For instance, power supplies that feed the computing resources can rely on tidal energy harvesting mechanisms, super-capacitors, underwater battery-free sensor network that use back-scatter technology or offloading computing to surface infrastructure via laser [19] to extend battery lifetime.

Communication and cooperation. Underwater communication technologies have limited range, coverage, and bandwidth [17]. As a result, multiple communications interfaces are preferable instead of relying on a single one, e.g., combining acoustic, electromagnetic, and optical. Communication technologies are sensitive to environmental characteristics such as salinity, water temperature or currents. In addition, lack of standards and high resource fragmentation are critical challenges that limit cooperative operations underwater. Thus, submerged vehicles or other computing technology mostly work in isolation. Seamless integration of protocols, communication interfaces, services and processing resources are necessary to foster distributed cooperation in the underwater environment.

B. Software Platforms and Marine Datasets

Currently most data collected underwater are analysed passively on the surface rather than proactively in the underwater environment. Existing software platforms mostly have been designed for offline analysis, e.g., there are software platforms for aiding ecologists such as CurvRank for whales, finFindR for dolphins, MYDAS for turtles, as well as more generic platforms such as NNPool [20] or PhotoID Ninja. Despite effort to perform heavy computations underwater, unfortunately none of the available systems can provide real-time analysis.

Resource intensive processing in deep sea. While the surrounding cold water can provide natural cooling for the encased computing resources, the capacity of the material for thermal absorption may be not enough to cope with heating caused due to heavy DL processing. This can result in reduced processing performance for DL. Distributed processing can alleviate this issue and offer increased processing capacity. However, several other challenges have to be addressed first to enable it.

Advanced autonomy. Currently human effort is required for underwater deployments. Efficient operational modes of infrastructure, e.g., energy-efficient underwater communication, and separate idle, moderate and high processing modes are key factors for achieving autonomous, continuous and active underwater presence through AUVs. This requires mechanisms for self-management, self-configuration, self-healing, self-optimisation, and self-protection of the AUVs.

Open SDKs and extendible APIs. Existing DL frameworks, e.g., TensorFlow, can easily be ported into smart devices. However, using these frameworks underwater is still in its infancy. A key limitation is that specialised underwater equipment does not have available open source firmware, e.g., Blue Robotics

TABLE II: Current AUV/ROV underwater technologies, challenges and directions for enabling deep learning deployment.

	State-of-the-Art	Key Research Challenges	Emerging Challenges
Underwater sensing	Data sampling approaches are mostly image-based	Migration of existing sensors for underwater operations	New portable and lightweight sensing solutions with low energy footprint
Available processing resources	Underwater solutions have limited processing resources	Augmenting resources with additional infrastructure	Advanced augmentation of resources with distributed and collaborative processing
Fault tolerance and operational time	Components operating underwater depend on operational time and correct functionality	Recovery and replacement of components without extraction from underwater	Multi-modal techniques to provide robust recovery and continuous operations
Communication and co-operation	Communications suffer from limited coverage and bandwidth performance. Deployments with dedicated wires	Adoption of different communication technologies, including electromagnetic, acoustic and optical	Emerging mature interfaces and better integration with new communication paradigms, including 5G and 6G
Resource intensive processing	Cooling functionality piggybacking environmental characteristics of the deployment	Design of better encasing to improve thermal absorption of heavy processing	Emerging approaches reduce thermal overhead based on distributed processing
Advance autonomy	Underwater operations require human support from experts, e.g., divers	New autonomous functionalities that reduce human intervention partially, e.g., back to home routines of UAVs and ROVs	Total autonomy for underwater solutions, e.g., self-healing, optimisation, protection and configuration
Open SDKs and extendable APIs	Existing frameworks are easily ported to smart and IoT devices	Open firmware to build a wider ecosystems of solutions	Adoption of a common and reusable platform to build underwater solutions
Lack of data diversity and massive datasets	Analysis of data is performed over passive collected measurements	Adoption of static monitoring solutions in different aquatic environments, e.g., lakes and rivers	Emerging integration of dynamic monitoring and collection of data using UAVs and ROVs

is one of the few options that provide open SDK for AUVs and ROVs. Thus, solutions are tailored for specific cases and not integrated into broader applications.

Lack of data diversity and massive datasets. The difficulty of conducting underwater explorations constrains the generation of datasets for research and development purposes. This is not limited only to oceans, but other aquatic areas, such as lakes and rivers. Data availability and diversity are key to generalise DL models that can be applied to different cases. For instance, DL models to classify water pollution can work seamlessly in open sea and rivers. Massive dataset generation can be envisioned by relying on a combination of data collection techniques such as AUVs, crowdsourcing, and even aerial imaging. Data diversity is also essential for ensuring robust DL performance and for developing solutions that can enhance performance. For example, data augmentation techniques can improve robustness of DL models in practical deployments, but they need to be able to capture variations in the data resulting from the actual deployment environment.

VI. SUMMARY AND CONCLUSION

Deep learning can bring significant benefits to underwater computing, supporting automated analysis of data in underwater monitoring and offering a mechanisms to support the operations of vehicles operating underwater. As hardware platforms are becoming increasingly affordable and available, the interest in underwater computing is likely to gather momentum. These applications are of significant interest to computing researchers as they offer unique algorithm and system challenges. Using plastic litter detection as an example application domain, we presented how embedded DL can increase the scale of monitoring operation by supporting automated in-situ analysis while also highlighting algorithmic and system level challenges in running the detection robustly.

ACKNOWLEDGMENTS

Authors are thankful for the suggestions by the anonymous reviewers. The research is partially supported by

the grants LARGESCALE (PTDC/CCI-CIF/32474/2017), LARSyS (UIDB/50009/2020), INTERWHALE (PTDC/CCI-COM/0450/2020), INTERTAGUA (MAC2/1.1.a/385) by FCT, MCTES, PIDDAC, MAC-INTERREG, the European Social Fund via the IT Academy Programme, and the Academy of Finland project 339614.

REFERENCES

- [1] L. Mayer, M. Jakobsson, G. Allen, B. Dorschel, R. Falconer, V. Ferrini, G. Lamarche, H. Snaith, and P. Weatherall, "The Nippon Foundation—GEBCO seabed 2030 project: The quest to see the world's oceans completely mapped by 2030," *Geosciences*, vol. 8, no. 2, p. 63, 2018.
- [2] H. Flores, N. H. Motlagh, A. Zuniga, M. Liyanage, M. Passananti, S. Tarkoma, M. Youssef, and P. Nurmi, "Toward Large-Scale Autonomous Monitoring and Sensing of Underwater Pollutants," *IEEE Internet of Things Magazine*, 2021.
- [3] L. Jin and H. Liang, "Deep learning for underwater image recognition in small sample size situations," in *OCEANS 2017-Aberdeen*. IEEE, 2017, pp. 1–4.
- [4] M. Ho, S. El-Borgi, D. Patil, and G. Song, "Inspection and monitoring systems subsea pipelines: A review paper," *Structural Health Monitoring*, vol. 19, no. 2, pp. 606–645, 2020.
- [5] M. Moniruzzaman, S. M. S. Islam, M. Bennamoun, and P. Lavery, "Deep learning on underwater marine object detection: A survey," in *International Conference on Advanced Concepts for Intelligent Vision Systems*. Springer, 2017, pp. 150–160.
- [6] M. Yang, J. Hu, C. Li, G. Rohde, Y. Du, and K. Hu, "An in-depth survey of underwater image enhancement and restoration," *IEEE Access*, vol. 7, pp. 123 638–123 657, 2019.
- [7] M. Radeta, N. J. Nunes, D. Vasconcelos, and V. Nisi, "Poseidon-passive-acoustic ocean sensor for entertainment and interactive data-gathering in opportunistic nautical-activities," in *Proceedings of the 2018 Designing Interactive Systems Conference*, 2018, pp. 999–1011.
- [8] M. Fulton, J. Hong, M. J. Islam, and J. Sattar, "Robotic detection of marine litter using deep visual detection models," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 5752–5758.
- [9] D. V. Politikos, E. Fakiris, A. Davvetas, I. A. Klampanos, and G. Papatheodorou, "Automatic detection of seafloor marine litter using towed camera images and deep learning," *Marine Pollution Bulletin*, vol. 164, p. 111974, 2021.
- [10] M. Sun, X. Yang, and Y. Xie, "Deep learning in aquaculture: A review," *Journal of Computers*, vol. 31, no. 1, pp. 294–319, 2020.
- [11] G. Dudek, P. Giguere, C. Prahacs, S. Saunderson, J. Sattar, L.-A. Torres-Mendez, M. Jenkin, A. German, A. Hogue, A. Ripsman *et al.*, "Aqua: An amphibious autonomous robot," *Computer*, vol. 40, no. 1, pp. 46–53, 2007.

- [12] J. Hong, K. de Langis, C. Wyethy, C. Walaszek, and J. Sattar, "Semantically-aware strategies for stereo-visual robotic obstacle avoidance," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 2450–2456.
- [13] O. Pfannkuche and P. Linke, "GEOMAR landers as long-term deep-sea observatories," *Sea Technology*, vol. 44, no. 9, pp. 50–55, 2003.
- [14] D. A. Duecker, N. Bauschmann, T. Hansen, E. Kreuzer, and R. Seifried, "HippoCampusX—A hydrobatic open-source micro AUV for confined environments," in *2020 IEEE/OES Autonomous Underwater Vehicles Symposium (AUV)*(50043). IEEE, 2020, pp. 1–6.
- [15] C. Edge, S. S. Enan, M. Fulton, J. Hong, J. Mo, K. Barthelemy, H. Bashaw, B. Kallevig, C. Knutson, K. Orpen *et al.*, "Design and Experiments with LoCO AUV: A Low Cost Open-Source Autonomous Underwater Vehicle," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 1761–1768.
- [16] A. Cadena, "A modular hardware-software architecture of an autonomous underwater vehicle for deep sea exploration," in *Proceedings of the World Congress on Engineering*, vol. 1, 2017.
- [17] S. Li, W. Qu, C. Liu, T. Qiu, and Z. Zhao, "Survey on high reliability wireless communication for underwater sensor networks," *Journal of Network and Computer Applications*, vol. 148, p. 102446, 2019.
- [18] H.-L. Wang, E.-M. You, R. Panneerselvam, S.-Y. Ding, and Z.-Q. Tian, "Advances of surface-enhanced Raman and IR spectroscopies: from nano/microstructures to macro-optical design," *Light: Science & Applications*, vol. 10, no. 1, pp. 1–19, 2021.
- [19] A. Trichili, C. Issaid, B. Ooi, and M. Alouini, "A cnn based structured light communication scheme for internet of underwater things applications," *IEEE Internet of Things Journal*, 2020.
- [20] R. Maglietta, V. Renò, R. Caccioppoli, E. Seller, S. Bellomo, F. Santacesaria, R. Colella, G. Cipriano, E. Stella, K. Hartman *et al.*, "Convolutional neural networks for risso's dolphins identification," *IEEE Access*, 2020.

Marko Radeta is an Assistant Professor at the Department of Faculty of Exact Sciences and Engineering at the University of Madeira, 9000 Funchal, Portugal; Coordinator of Wave Labs; and researcher at the Marine Environmental Sciences Centre (MARE) and ARDITI. His research interests include the Internet of Things, artificial intelligence, marine sensing, Human-Computer Interaction and astrophysics. Contact him at marko@wave-labs.org.

Agustin Zuniga is a doctoral researcher with the Department of Computer Science, University of Helsinki, 00014 Helsinki, Finland. His research interests include pervasive sensing, machine learning, and data science. Contact him at agustin.zuniga@helsinki.fi.

Naser Hossein Motlagh is a postdoctoral researcher with the Department of Computer Science, University of Helsinki, 00014 Helsinki, Finland. His research interests include the Internet of Things, wireless sensor networks, unmanned aerial vehicles, and autonomous underwater vehicles. Contact him at naser.motlagh@helsinki.fi.

Mohan Liyanage is a lecturer at the University of Tartu, 51009 Tartu, Estonia. He obtained his Ph.D. in computer science from the University of Tartu in 2019. His research interests include mobile edge computing, the Internet of Things, and pervasive systems. Contact him at mohan.liyanage@ut.ee.

Ruben Freitas is a doctoral researcher at the Department of Faculty of Exact Sciences and Engineering at the University of Madeira, 9000 Funchal, Portugal; and researcher at the Wave Labs, Marine Environmental Sciences Centre (MARE) and ARDITI. His research interests include artificial intelligence and marine monitoring. Contact him at ruben@wave-labs.org.

Moustafa Youssef is a Professor of computer science and engineering at American University in Cairo, 11835 New Cairo, Egypt and Alexandria University, 21544 Cairo, Egypt. He completed his Ph.D. in Computer Science, University of Maryland in 2004. His research interests include mobile wireless networks, mobile computing, location determination technologies, pervasive computing, and network security. Contact him at moustafa@alexu.edu.eg.

Sasu Tarkoma is a professor in the Department of Computer Science, University of Helsinki, 00014 Helsinki, Finland. His research interests include mobile computing, Internet technologies, and artificial intelligence. Contact him at sasu.tarkoma@helsinki.fi.

Huber Flores is an Associate Professor with the Institute of Computer Science, University of Tartu, 51009 Tartu, Estonia; and a docent with the Department of Computer Science, University of Helsinki, Finland. He received his PhD in computer science in 2015 at the University of Tartu, Estonia. His research interests include mobile and pervasive computing, distributed systems and mobile cloud computing. Contact him at huber.flores@ut.ee.

Petteri Nurmi is an Associate Professor at the University of Helsinki and Science, Finland. He received his PhD in computer science in 2009 at the University of Helsinki, 00014 Helsinki, Finland. His research is in the area of pervasive computing and Internet of Things, and is characterized by a focus on research questions surrounding the collection, analysis and use of sensor data. Contact him at petteri.nurmi@helsinki.fi.