



# iImagine

## Validated prototype services for three image analysis use cases

iImagine Deliverable D3.5

31/07/2025

### Abstract

Oceans, seas, coastal and inland waters are vital for our societies and the future of our planet. The EU-funded project iImagine addresses challenges in the aquatic sciences by supporting the domain scientists in better and broader use of imaging data through providing a portfolio of high-performance image analysis tools empowered with Artificial Intelligence (AI), and Best Practices documents. The deliverable presents successful development, validation, and delivery of three AI-powered prototype services for image analysis in aquatic sciences: Underwater Noise Identification, Beach Monitoring, and Freshwater Diatoms Identification. In addition, the project onboarded and helped further six external use cases in the domain, which we briefly review in this deliverable.



**Funded by  
the European Union**

iImagine receives funding from the European Union's Horizon Europe research and innovation programme under grant agreement No. 101058625.

Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union, which cannot be held responsible for them.

## Document Description

D3.5 Validated prototype services for 3 image analysis use cases			
<b>Work Package 3</b>			
Due date	31/07/2025	Actual delivery date:	31/07/2025
Nature of document	Report	Version	1.0
Dissemination level	Public		
Lead Partner	KIT		
Authors	K. Alibabaei (KIT), E. Azmi (KIT), W. Decrop (VLIZ), V. Kozlov (KIT), M. Laviale (UL-LIEC), J. Soriano-González (SOCIB)		
Reviewers	E. Martinez (UPC), D. Smyth (MI)		
Public link	<a href="https://zenodo.org/records/16638761">https://zenodo.org/records/16638761</a>		
Keywords	Aquatic science, Images, AI-based analysis, Validation, Delivery		

## Revision History

Issue	Item	Comments	Author/Reviewer
V O.1	Draft version	ToC, initial structure	V. Kozlov (KIT), K. Alibabaei (KIT)
V O.2	Draft version	Drafting first full version	K. Alibabaei (KIT), E. Azmi (KIT), W. Decrop (VLIZ), J. Soriano-González (SOCIB), M. Laviale (UL-LIEC), V. Kozlov (KIT)
V O.3	Revised version	Review	E. Martinez (UPC), D. Smyth (MI)
V O.4	Revised version	Revised after review version	V. Kozlov (KIT), J. Soriano-González (SOCIB), W. Decrop (VLIZ), M. Laviale (UL-LIEC), K. Alibabaei (KIT), E. Azmi (KIT)
V 1.0	Submitted version		

## Copyright and licence info

This material by parties of the iMagine Consortium is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

## Table of Contents

Executive summary .....	6
Introduction.....	7
Purpose of the document.....	8
Scope of the document .....	8
Structure of the document .....	8
Validation methodology.....	9
Data delivery .....	9
Model validation .....	10
Service delivery.....	10
Code Release: GitHub Open-Access.....	10
Model Sharing via iMagine Marketplace.....	11
Inference service delivery via the project's OSCAR serverless platform .....	11
AI image analysis use cases as prototype services .....	12
UC6 Underwater noise identification.....	12
Description.....	12
Achieved results.....	13
Findings and Outlook .....	17
UC7 Beach monitoring.....	18
Description.....	18
Achieved results.....	19
Findings and Outlook .....	25
UC8 Freshwater diatoms identification .....	25
Description.....	25
Achieved results.....	27
Findings and Outlook .....	31
Additional use cases .....	32
EyeOnWater .....	32
Cold water coral reefs (call 1).....	32
Satellite-derived bathymetry (call 1).....	33
Age reading from fish otoliths (call 2) .....	33
DEcentrAlised Learning (DEAL) (call 3) .....	33

Sea Wave and coastal inundation detection Methodology (SWiM) (call 3) .....	34
Conclusion.....	35

## List of Figures

- Figure UC6.1 – Overview of the data integration and processing pipeline. On the left, the AIS coordinates are matched with hydrophone recordings to compute vessel proximity. On the right, the corresponding audio segments are converted into Log-Mel spectrograms and processed via two approaches.
- Figure UC6.2 – Confusion matrices for the obtained models using the best-performing hyperparameters, evaluated in the independent test set for the approaches (a) Feature extraction and (b) Fine-tuning. Values in the confusion matrix refer to the number of instances, whereas the colouring scale shows the percentage of instances in the total test set.
- Figure UC6.3 – Comparison of the power spectrum in relation to vessel proximity from Grafton on 2022-09-13. The top section of each figure displays the power spectrum for the entire day's recording, while the bottom section shows the distances between nearby vessels and the hydrophone. AIS data is shown in black, model classifications in blue, and other AIS vessels in grey. Each classification represents a distance category of 1 km width (e.g. 0--1 km, 1--2 km), and the corresponding dot is positioned at the centre of the predicted interval (e.g. 0.5 km for the 0--1 km category). (a) shows results from the feature extraction model. (b) Fine-tuning model
- Figure UC6.4 – Comparison of the power spectrum in relation to vessel proximity from Grafton on 2022-09-13. The top section of each figure displays the power spectrum for the entire day's recording, while the bottom section shows the distances between nearby vessels and the hydrophone. AIS data is shown in black, model classifications in blue, and other AIS vessels in grey. Each classification represents a distance category of 1 km width (e.g. 0--1 km, 1--2 km), and the corresponding dot is positioned at the centre of the predicted interval (e.g. 0.5 km for the 0--1 km category). (a) shows results from the feature extraction model. (b) Fine-tuning model
- Figure UC6.5 – Time distribution of acoustic data throughout 2022, colour-coded by split type (training, validation, and testing).
- Figure UC7.1 – Left: source image. Right: detected beach wracks and masks. 'Dense wrack' indicates high-density, while 'intermediate wrack' signifies low-density or a mix of wrack and sand.
- Figure UC7.2 – Left: source image. Right: detected rip currents.
- Figure UC7.3 – Left: original crowd-sourced submission. Right: rectified image with the extracted shoreline (blue) superimposed.

- Figure UC8.1 –Examples of diatom images usually misclassified due to high inter-class similarity
- Figure UC8.2 – Image distribution across the 799 diatom species of our “Individual” dataset.
- Figure UC8.3 – Screenshots of the diatom detection and classification applications as available online

## List of Tables

- Table UC7.1 – UC7 coastal features, definition, and target groups.
- Table UC7.2 – Summary of UC7 activities and achievements. The Presentations field includes the reached audience. Acronyms are defined in the dedicated section.
- Table UC7.3 – UC7 open-source labelled datasets and reference (ref).

## Acronyms

Please see <https://confluence.egi.eu/display/IMPAIP/Glossary>

## Executive summary

This deliverable presents successful development, validation, and delivery of three AI-powered prototype services for image analysis in aquatic sciences: Underwater Noise Identification (UC6), Beach Monitoring (UC7), and Freshwater Diatoms Identification (UC8). These services were built using the iMagine AI platform, which provides tools for data preparation, model training, validation, and deployment. Each use case followed a structured methodology involving open data delivery (via Zenodo), model validation (with performance metrics and expert review), and service deployment (via GitHub and the iMagine Marketplace).

UC6 trained a CNN model to classify vessel proximity using underwater acoustic data, aiding marine conservation. UC7 applied deep learning to detect beach wracks, rip currents, and shoreline positions from images, and produced high-quality datasets. UC8 focused on automating the identification of freshwater diatoms, a key bioindicator, using a combination of synthetic and real datasets to train robust classification and detection models. All three services are now accessible to the scientific community, demonstrating the platform's effectiveness in accelerating AI adoption in aquatic research. In addition, the project onboarded and helped further six external use cases in the aquatic domain.

## Introduction

Oceans, seas, coastal and inland waters are vital for our societies and the future of our planet. The EU-funded project iImagine addresses challenges in the aquatic sciences domain by supporting the scientists in better and broader use of imaging data by providing a portfolio of high-performance image analysis tools empowered with Artificial Intelligence (AI), and Best Practices documents. In particular, the project:

- Delivers a scalable, shared IT platform for image analysis.
- Advances existing image analytical services to increase research performance in aquatic sciences.
- Develops prototype new image analytical services and datasets that can accelerate progress towards healthy oceans, seas, coastal and inland waters.
- Captures and disseminates development and operational best practices to imaging data and image analysis service providers.
- Delivers a portfolio of scientific image and image analytics services targeting researchers in marine and aquatic sciences.

To effectively achieve the objectives of the project, fourteen use cases (WP3) in various areas of aquatic sciences collaboratively engage with the iImagine AI Platform (WP4). These use cases include five mature cases started with the project:

- UC1 Marine litter assessment
- UC2 Zooscan – EcoTaxa pipeline
- UC3 Marine ecosystem monitoring at EMSO sites (OBSEA, Azores, SmartBay)
- UC4 Oil spill detection
- UC5 Flowcam plankton identification

Three prototype use cases, that also started with the project and focus on image services with high potential for uptake of AI in their analysis process:

- UC6 Underwater noise identification
- UC7 Beach monitoring
- UC8 Freshwater diatoms identification

And six additional use cases, selected and onboarded via open calls:

- EyeOnWater
- Cold water coral reefs
- Satellite-derived bathymetry
- Age reading from fish otoliths
- DEcentrAlised Learning (DEAL)
- Sea Wave and coastal inundation detection Methodology (SWiM)

This document serves as a comprehensive report that provides a detailed description of each of the three prototype use cases (UC6 – UC8) for validating their AI applications.

These use cases reached the prototype level (TRL<sup>1</sup> 6–7), making use of the iMagine framework installation and expertise of the iMagine Competence Centre experts and the synergy with the other use cases, both prototype and mature. During the course of iMagine, the UCs gained practical experience in developing AI projects and performing image analysis: from collecting images, preparing training datasets, and training AI models, to deploying models for inference, specifically in the field of aquatic sciences. In this deliverable, we go through a general overview of these processes and highlight key learnings.

### Purpose of the document

This document presents a comprehensive overview of the development, validation, and delivery of the three prototype image analysis services carried out within the iMagine project. It outlines the overall validation methodology, data delivery processes, and model validation strategies used across the three iMagine prototype use cases.

### Scope of the document

The D3.5 deliverable is our last deliverable after mature use cases progressed towards service delivery and best practices for producers and providers of image analysis applications were derived. This deliverable describes in detail what the three prototype use cases, namely, underwater noise identification, beach monitoring, and freshwater diatom identification, achieved in the course of the project. It highlights their specific goals, achievements, data workflows, model validation, key findings, and outlook. The document details how the three prototype use cases are validated based on the suggested methodology.

### Structure of the document

This deliverable is structured as follows: we begin with the validation methodology for prototype use cases containing data, model, and service delivery approaches. We continue with describing three AI image analysis use cases as prototype services and presenting their achievements and validation. In the following sections, a description of the additional use cases onboarded via open calls is provided. Finally, the experiences of both groups of use cases and the support of iMagine in achieving their goals are summarized in the conclusion.

---

<sup>1</sup> Technology Readiness Level, e.g.: <https://horizoneuropencpportal.eu/sites/default/files/2022-11/trl-assessment-tool.pdf>



## Validation methodology

The validation methodology for prototype use cases is based on our previous deliverables D3.3 “AI application upgrade, deployment, and operation plan”<sup>2</sup> and D3.4 “Best practices for producers and providers of image sets and image analysis applications in aquatic sciences”<sup>3</sup> and comprises the following items to address:

- Delivery of the training data
- AI model validation
- Delivery of the application or service to scientists via the iImagine platform

If all three items are completed, we may conclude that the image analysis use case is validated:

- Raw images are processed, and the training dataset is prepared and shared with the aquatic science community.
- The trained AI model is validated using best practices recommendations.
- The final application is made available to other scientists via the iImagine platform for further reuse. This also includes a number of software quality control tests and checks by a platform expert.

## Data delivery

Zenodo is the preferred platform for data storage and sharing due to several compelling features. First, it offers long-term storage, supported by the EU, ensuring the preservation and dissemination of research results over time. Second, Zenodo assigns each publication a Digital Object Identifier (DOI), making the datasets easy to locate and cite. Additionally, the platform provides robust version control, allowing researchers to track changes and access specific dataset versions as needed. Zenodo also accommodates generous storage needs, offering 50 GB of free storage per publication with the option to request additional space. Finally, it provides detailed usage insights, including download and access statistics, enabling researchers to gauge the reach and impact of their data effectively.

iImagine has collaborated with Zenodo as part of the EU-funded Zenodo-ZEN<sup>4</sup> project to develop a more domain-specific approach and created a dedicated template for its training image datasets, structured as a DCAT profile<sup>5</sup> and supported by aquatic vocabularies. This profile was used in dialogue with the Zenodo team, to create extensions to the generic Zenodo metadata template. As a result, more options are now available in

---

<sup>2</sup> Kozlov, V., Sipos, G., & Schaap, D. (2024). iImagine D3.3 AI application upgrade, deployment, and operation plan (Final). Zenodo. <https://doi.org/10.5281/zenodo.14893946>

<sup>3</sup> Alibabaei, K., Azmi, E., & Kozlov, V. (2024). iImagine D3.4 Best practices for producers and providers of image sets and image analysis applications in aquatic sciences (1.1 Final). Zenodo. <https://doi.org/10.5281/zenodo.14893997>

<sup>4</sup> <https://about.zenodo.org/projects/horizon-zen>

<sup>5</sup> <https://www.w3.org/TR/vocab-dcat-3>

the Zenodo upload form, allowing for capturing in more detail important information for AI training datasets.

## Model validation

Generally, AI models rely heavily on data for proper training, making predictions, and drawing conclusions. A trained AI model is not immune to possible biases inherent in the data. Therefore, it is important to assess potential data biases and imply strategies for their mitigation. Performance metrics and evaluation methods are other essential components of any machine learning project, as they provide insights into the effectiveness and accuracy of the trained models. The performance metrics have to be compared not only between various AI experiments but also against the baseline model, i.e. the simplest model, which may not include AI training, and which serves as the reference point. While assessing various preprocessing techniques, AI models, and choosing the right metrics, it is also necessary to keep track of these experiments and compare them, e.g. with one of the experiments tracking frameworks or tools.

The model validation strategy thus involves:

- Assessment of data biases and their mitigation, if needed
- Definition of the baseline model
- Choosing performance metrics and evaluation methods (e.g. confusion matrix, F1-Score, ROC-AUC, MSE, mAP etc.<sup>6</sup>)
- Keeping track of AI experiments
- Validation of trained models by the domain scientists

## Service delivery

### Code Release: GitHub Open-Access

The source code of AI models of use cases is published under the ai4os-hub<sup>7</sup> GitHub organization under one of the OSI-approved licences<sup>8</sup>. The repository provides open access to the model code, documentation, and associated resources. By hosting the model within the ai4os-hub organization, the use cases benefit from increased visibility and integration with related AI initiatives.

---

<sup>6</sup> Alibabaei, K., Azmi, E., & Kozlov, V. (2024). iMagine D3.4 Best practices for producers and providers of image sets and image analysis applications in aquatic sciences (1.1 Final). Zenodo. <https://doi.org/10.5281/zenodo.14893997>

<sup>7</sup> <https://github.com/ai4os-hub>

<sup>8</sup> <https://opensource.org/licenses>

## Model Sharing via iImagine Marketplace

The simplest option is to provide the trained models via the iImagine Marketplace component of the platform. The trained and validated models are integrated with an API (recommended is DEEPaaS API, REST API for AI/ML/DL) and dockerized. This enables external users to download the AI modules as Docker images and run them on their own or third-party external compute resources. This pattern combines two complementary approaches:

1. **“Marketplace inference service delivery”** to run trained AI models for inference on the connected back-end cloud resources. The latter is also available for non-partners via short-lived (10 minutes) “Try” endpoints.
2. **“Marketplace download service delivery”**: allows users to download the AI modules as Docker images and run them on in-house/third-party external compute resources, including their own computers/laptops.

Publishing modules on the iImagine Marketplace also means that they passed a number of quality control checks, e.g. unit and security tests, and basic validation by the platform experts.

## Inference service delivery via the project’s OSCAR serverless platform

The models hosted on the Marketplace can also be made accessible for inference through the scalable OSCAR<sup>9</sup> clusters provided by the iImagine project. This allows users to choose and run trained AI models for inference on the connected back-end cloud resources of iImagine. Users will need either an EGI Check-in account or using AI4EOSC Access/Login to access the OSCAR platform. There are two OSCAR clusters available for the iImagine project, called OSCAR Inference<sup>10</sup> and OSCAR Inference-walton<sup>11</sup>. Both currently provide only CPU computing resources. The iImagine OSCAR clusters support multitenancy and offer accounting thanks to the Prometheus and GoAccess services. They allow collecting metrics like resources consumed by inference executions, the number of deployed services over a period of time, or the geolocation of the users interacting with the services. Use cases as well have the option of deploying their own OSCAR cluster, which can be set up by a project partner on computational resources provided by the iImagine consortium or on external third-party infrastructure. In such cases, support and documentation for deploying the open-source OSCAR platform are provided, but ongoing platform management remains the responsibility of the use case partner.

---

<sup>9</sup> <https://docs.oscar.grycap.net/latest/>

<sup>10</sup> <https://inference.cloud.imagine-ai.eu/>

<sup>11</sup> <https://inference-walton.cloud.imagine-ai.eu/ui/#/login>

## AI image analysis use cases as prototype services

### UC6 Underwater noise identification

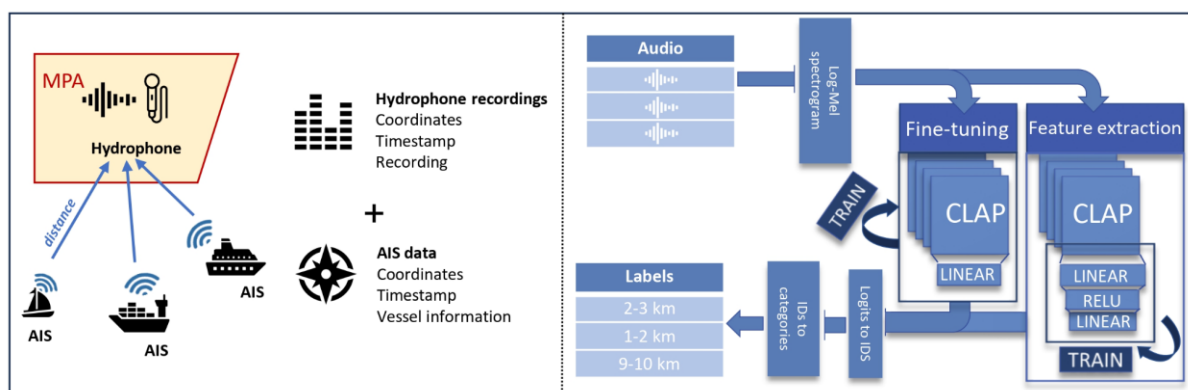
#### Description

Marine environments are increasingly affected by human activities, which generate underwater noise as a by-product. Acoustic data from these environments can offer valuable insights for tracking human activity and improving the monitoring of sensitive areas, such as Marine Protected Areas (MPAs) and offshore wind farms.

This UC presents a convolutional neural network (CNN) trained to classify vessel distances from passive acoustic recordings. During pre-processing, each 10-second audio file produced one log-mel spectrogram with a size of (1001, 64).

We constructed an open-source, diverse dataset by integrating 116 days of acoustic data from two stations in the Belgian North Sea with Automatic Identification System (AIS) data. The CNN was trained to classify acoustic clips into discrete distance bins, representing the proximity of the nearest vessel.

Two approaches were taken, in the feature extraction approach, the extracted features were passed through three custom-made layers, whereas in the fine-tuning approach, the entire model was re-trained and followed by just one linear layer, as shown in **Figure UC6.1**.



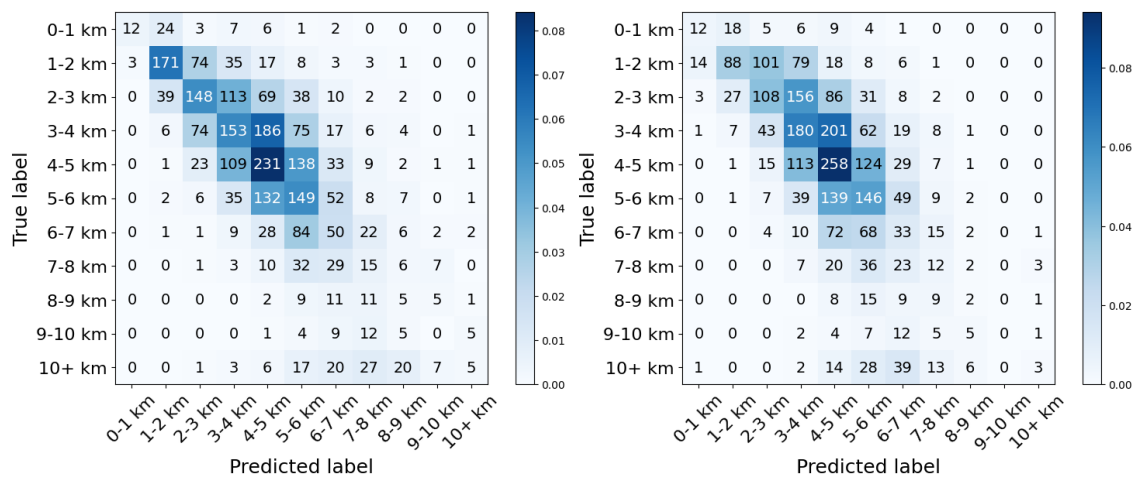
**Figure UC6.1** – Overview of the data integration and processing pipeline. On the left, the AIS coordinates are matched with hydrophone recordings to compute vessel proximity. On the right, the corresponding audio segments are converted into Log-Mel spectrograms and processed via two approaches.

Our results demonstrate that the model can effectively distinguish between distance categories using underwater sound alone, confirming the feasibility of passive acoustic monitoring for vessel activity. This technology provides an innovative approach to enhance MPA oversight and represents a first step in a promising pathway for conservation efforts.

## Achieved results

The fine-tuned model achieved an RMSE of 1.587, whereas the feature extraction model reached 1.709. The obtained confusion matrices for the best-performing fine-tuning and feature extraction models, respectively, can be seen in **Figure UC6.2**. Most values are concentrated near the diagonal in both matrices, indicating that the models correctly predict the majority of classes. However, some misclassifications are evident.

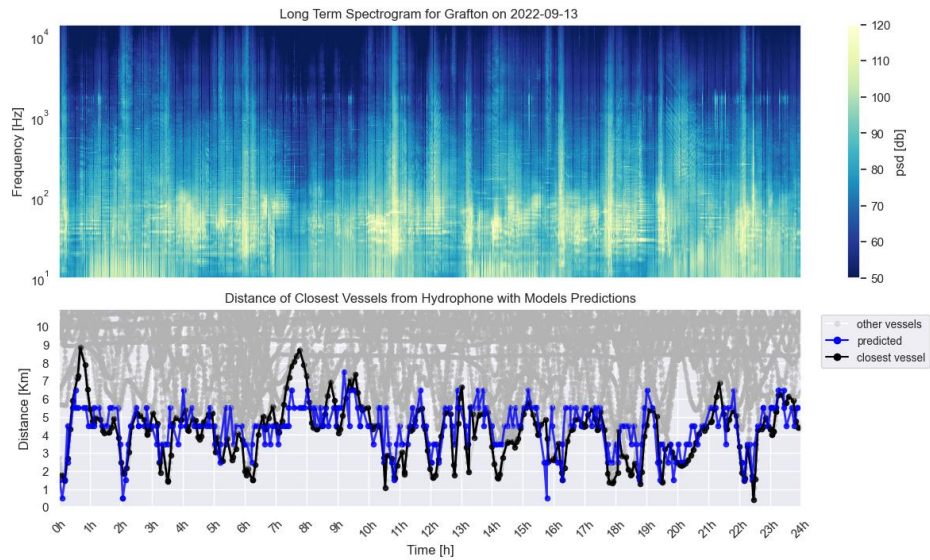
Underestimations—where vessels are classified as being closer than they actually are—are more common than overestimations. This is especially noticeable in the middle-distance bins (e.g. 3–6 km), where vessel sounds may resemble closer-range acoustic patterns.



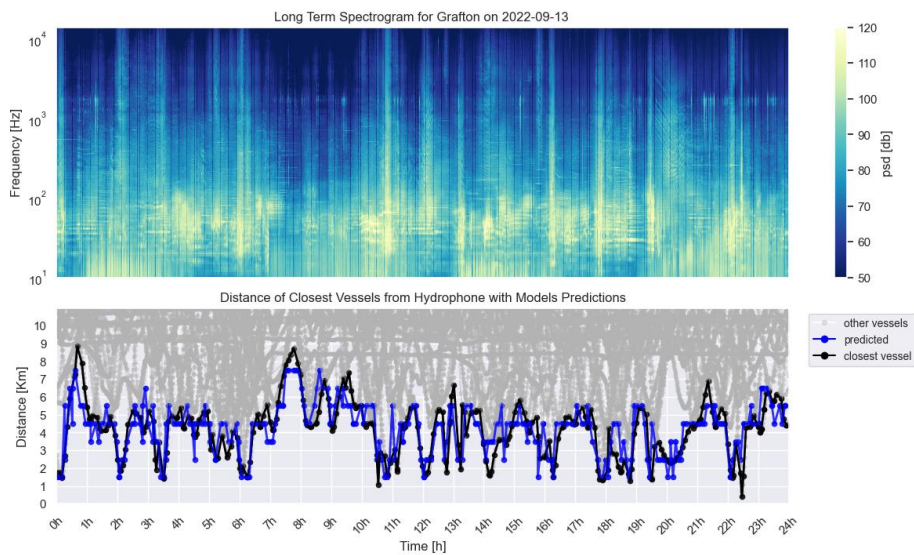
**Figure UC6.2** – Confusion matrices for the obtained models using the best-performing hyperparameters, evaluated in the independent test set for the approaches (a) Feature extraction and (b) Fine-tuning. Values in the confusion matrix refer to the number of instances, whereas the colouring scale shows the percentage of instances in the total test set.

**Figure UC6.3** and **Figure UC6.4** display the classifications of the best-performing model for an entire day's recording at two stations, comparing the fine-tuning and feature extraction approaches. In both cases, the model predictions (blue lines) align more closely with the AIS ground truth (black lines) when using the fine-tuning approach. This outcome is expected, as the fine-tuned model consistently outperformed the feature extraction model across evaluation metrics, reflecting its improved ability to learn and generalise from the training data. Specifically, it demonstrates greater sensitivity to nuanced acoustic patterns, allowing it to better detect distant vessel presence. This improved alignment is particularly evident in the outer distance bins (e.g. 9–10 km and

10+ km), supporting the earlier observations from the confusion matrices and RMSE plots.



(a) Feature extraction model



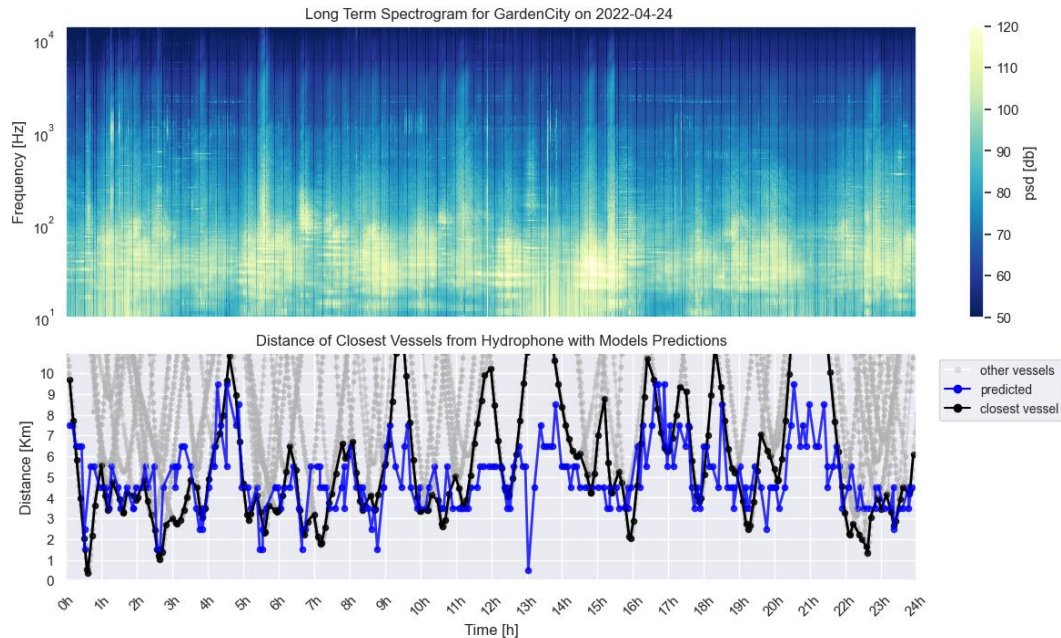
(b) Fine-tuning model

**Figure UC6.3** – Comparison of the power spectrum in relation to vessel proximity from Grafton on 2022-09-13. The top section of each figure displays the power spectrum for the entire day's recording, while the bottom section shows the distances between nearby vessels and the hydrophone. AIS data is shown in black, model classifications in blue, and other AIS vessels in grey. Each classification represents a distance category of 1 km width (e.g. 0--1 km, 1--2 km), and the corresponding dot is positioned at the centre of the predicted interval (e.g. 0.5 km for the 0--1 km category). (a) shows results from the feature extraction model. (b) Fine-tuning model

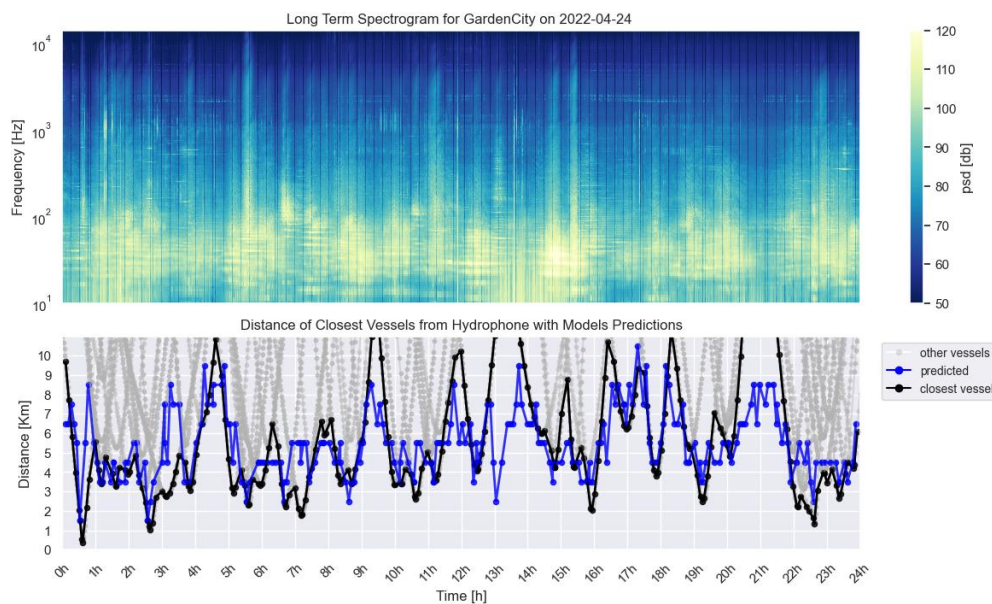
Notably, in **Figure UC6.4**, both model's classifications diverge from the AIS data around 13h. However, upon reviewing the audio and power spectrum, evidence of a vessel's presence emerges that is not captured by the AIS data. This indicates that the model



successfully detected a *dark* vessel around 13h, which the AIS failed to register. Such detections are particularly valuable, as they highlight the model's capacity to identify real-world vessel activity that eludes traditional AIS-based monitoring. These results demonstrate the potential of passive acoustic models to complement AIS data by filling observational gaps. While this is just one illustrative case, similar occurrences were noted at other time points and locations; additional examples are provided in the supplementary material.



(a) Feature extraction model



(b) Fine-tuning model

**Figure UC6.4** – Comparison of the power spectrum in relation to vessel proximity from Grafton on 2022-09-13. The top section of each figure displays the power spectrum for the entire day's recording, while the bottom section shows the distances between nearby vessels and the hydrophone. AIS data is shown in black, model classifications in blue, and other AIS vessels in grey. Each classification represents a distance category of 1 km width (e.g. 0--1 km, 1--2 km), and

*the corresponding dot is positioned at the centre of the predicted interval (e.g. 0.5 km for the 0–1 km category). (a) shows results from the feature extraction model. (b) Fine-tuning model*

## Data delivery

To support reproducibility and future research in underwater noise monitoring, we constructed a labelled dataset by integrating 116 days of continuous passive acoustic recordings from two hydrophone stations in the Belgian North Sea with vessel position data from the Automatic Identification System (AIS). The acoustic data was segmented into 10-second clips and annotated with the distance to the nearest vessel using a window-based approach that accounts for intermittent AIS transmission. Each clip is paired with metadata including timestamp, station ID, and a discrete distance bin label.

The resulting dataset contains several hundred thousand annotated audio samples spanning a wide range of acoustic conditions and vessel proximities. It serves as a valuable resource for developing and evaluating machine learning models for passive acoustic monitoring.

The dataset is publicly available via the VLIZ data portal and can be accessed through the following DOI<sup>12</sup>. It will be publicly available after the publication of its corresponding paper (accepted).

## Model validation

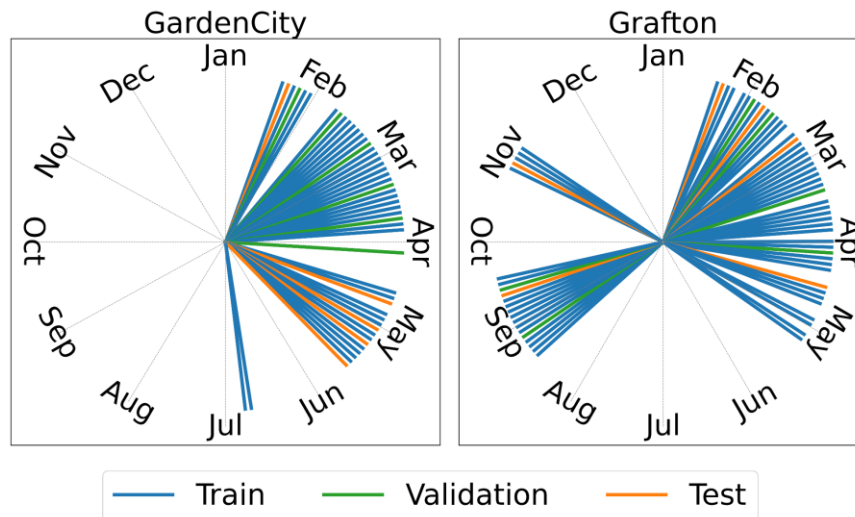
The distribution from the **Figure UC6.5** was used for data training (79.4%), testing (9.9%), and validation (10.6%). As a baseline, we used the AIS-ground truth. So we could compare the predictions of our model with the ones from the dataset. A pre-trained version of CLAP-LAION<sup>13</sup> partially trained on underwater bioacoustics data, was selected for transfer learning and finetuning. The loss function can be compared to the same functionality as a Mean Squared Error (MSE).

---

<sup>12</sup> <https://doi.org/10.14284/723>

<sup>13</sup> [doi:10.1109/ICASSP48485.2024.10447250](https://doi.org/10.1109/ICASSP48485.2024.10447250)





**Figure UC6.5** – Time distribution of acoustic data throughout 2022, colour-coded by split type (training, validation, and testing).

### Service delivery

The service is operational as an API to do inference on the iMagine marketplace<sup>14</sup>. The user can choose to upload a 10 sec .wav file or embedding to compute the distance to the closest vessel. Currently, the model weights (for both fine-tuning and feature extraction) are available upon request via email. They are not open source at this stage, as the corresponding paper is under review, but will be made publicly available upon its acceptance. The fine-tuning model has a slightly higher performance and is most recommended. The code is available as open source on GitHub under the MIT Licence. It is integrated with the DEEPaaS API. More information about the code can be found on GitHub<sup>15</sup>.

### Findings and Outlook

Working with the AI4OS/iMAGINE platform has been a positive experience overall. The platform proved to be user-friendly and well-suited for scientific development workflows. I used the AI4OS Development Environment to build and train this model. The tools and services provided—particularly the integration with containerized services, dataset handling, and model deployment—matched our expectations and allowed for a smooth development cycle.

I relied mainly on the AI4OS Development Environment for the entire process, from data preparation and model training to testing and packaging. The environment was intuitive, and the platform’s modular structure facilitated reproducibility and scalability. Progress was largely in line with my expectations.

<sup>14</sup> <https://dashboard.cloud.imagine-ai.eu/catalog/modules/audio-vessel-classification>

<sup>15</sup> <https://github.com/ai4os-hub/audio-vessel-classification>

Looking ahead, the current version of the service allows users to submit acoustic recordings and receive predictions on the distance to the nearest vessel, based purely on underwater sound. While this is already a valuable capability, there is still room to enhance the model's performance further, especially under challenging acoustic conditions or with unreported (dark) vessels. At the current stage, the user can also choose between two types of models to execute their predictions, balancing speed/accuracy for their performance. Similar to this API, an automated script can be created for larger batches of data.

Future research will focus on improving both the robustness and interpretability of the model. Furthermore, besides the distance, more information about the vessel can be gathered. This can include the vessel type, speed, engine, etc.

Furthermore, currently only the closest vessel (according to AIS data) is looked at, but it could be that the second closest has a more significant contribution to the observed noise, so perhaps the model can learn from the combination of different vessels close by. The model is currently trained on the Belgian part of the North Sea, which has a maximum depth of around 45m, so perhaps a more generalised model can be created by combining data from different environments.

We plan to engage the user community actively through ongoing scientific publication and dataset publications.

To summarise

- Improving Model Generalizability: by using datasets from different environments
- Integrating Models into Operational Frameworks: coding pipelines to automate the closest vessel predictor in batches
- Extending the model output:
  - Predicting not only on distance, but also other metrics (e.g. vessel type)
  - Looking at the closest two or three vessels.

## UC7 Beach monitoring

### Description

This use case aims to automate beach wrack identification, shoreline extraction, and rip currents detection from beach imaging systems (e.g. video monitoring stations and crowd-sourced smartphone imagery). These are key coastal features with particular importance for coastal research, management, and safety (**Table UC7.1**). Due to their highly dynamic nature, spatially and temporally, field-based monitoring methods are often insufficient. The increasing availability of visual data, however, presents a significant opportunity for enhanced coverage and resolution, but traditional image processing (e.g.

histogram-based or colour space transformation) lacks robustness in the face of complex scenes (e.g. variable lighting and meteoceanic conditions). With deep learning, we can overcome the complexity of detecting these features, enabling the development of consistent and automated systems capable of meeting the spatial and temporal demands required for effective coastal monitoring.

**Table UC7.1** – UC7 sub use cases: coastal features, definition, and target groups.

Coastal feature	Definition and relevance	Target groups
Beach wracks	Beach wracks are onshore accumulations of detached marine vegetation that appear as sparse patches or forming dense wedge structures in the shoreline. They can influence beach dynamics, but often conflict with tourism, leading to frequent removal, and presenting a complex management challenge	Coastal managers, researchers
Rip currents	Rip currents are narrow, fast-moving channels of water that flow seaward from the shore. They are one of the main causes of rescues and drownings all around the world and are considered the principal physical hazard to recreational bathers on surf beaches	Emergency services, coastal managers, researchers
Shoreline	The shoreline position serves as a fundamental proxy of coastal change. It has been monitored using images for over 30 years. However, traditional RGB image analysis methods struggle with complex conditions (e.g. variable lighting, waves, etc.), requiring too much human intervention and limiting the effectiveness of beach imaging systems	Coastal managers, researchers

## Achieved results

For UC7, the initial focus was on the detection of rip currents and the identification of beach seagrass wracks. Over time, the scope expanded to include shoreline extraction, which was successfully incorporated by leveraging existing data and resources available through the iMagine project. This expansion provided additional valuable results to the project. A summary of UC7's activities and achievements are provided in [Table UC7.2](#)

A significant undertaking for UC7 was in data annotation, as the project started without any existing labelled dataset. Various annotation tools, such as CVAT and Label Studio, were tested to identify the optimal platform for UC7's specific requirements. The feedback gathered was then shared with other project partners and the wider community. Then data was carefully selected and, for each coastal feature and imaging system, a

training dataset was created and published in the Zenodo iImagine community open repository ([Table UC7.3](#)).

In terms of model development, a range of approaches was explored. This included testing semantic segmentation, instance segmentation, and detection methods, including various deep learning architectures (e.g. Recurrent Neural Networks (RNNs), convolutional Neural Networks (CNNs)) and models (Bi-LSTM, U-Net, YOLO). Throughout this process, extensive experimentation was carried out with data preprocessing techniques (e.g. class configuration) and training parameters (e.g. number of hidden layers, warmup epochs) to optimize model performance for each specific case.

The work within UC7 also contributed significantly to scientific discourse. The team led one conference paper, contributed to another published journal paper ([Table UC7.2](#)), and is currently developing a journal manuscript for each of the sub-use cases. Dissemination of these findings occurred through several presentations to diverse audiences. These included a presentation to engineering informatics BSc students at the University of the Balearic Islands (UIB) in 2023, aimed at promoting their interest in joining the research. In 2024, a presentation of the rip currents detection prototype was given to end users, including administrative staff, technicians, and lifeguards from the General Directorate of Emergencies and Internal Affairs of the Balearic Islands (DGE) during the 'Official beach opening day'. Additionally, two scientific presentations were delivered: one for young researchers at the Coast2Coast symposium (Polytechnic University of Catalonia) in 2024, and an oral presentation to coastal researchers at the Coastal Dynamics conference in 2025 (Aveiro, Portugal).

Academic development was also a key focus. A collaboration agreement was formalized with the Laboratory for Artificial Intelligence Applications of the UIB (LAIA@UIB). Furthermore, three students were successfully recruited and trained, resulting in the defence of two Bachelor of Science theses ([Table UC7.2](#)) and the ongoing work on two Master of Science projects.

UC7 established external collaborations, forming specific partnerships for each sub-use case. For beach wrack, a partnership with the Bureau de Recherches Géologiques et Minières (BGRM), the French geological survey, facilitated data sharing to extend model capabilities and generalize to other coastal imaging systems, environments, and beach wrack compositions. In rip current detection, an existing collaboration with the DGE of the Balearic Islands was strengthened to meet UC7's requirements. Their interest added value, validated efforts, and provided a two-year database of rip current occurrences on various beaches, aiding in training dataset image selection. For shoreline extraction, UC7 is being tested by the Laboratório Nacional de Engenharia Civil (LNEC), the Portuguese National Civil Engineering Laboratory. LNEC, which works with coastal cameras, is interested in extracting waterlines, aligning with UC7's shoreline extraction approach.

Finally, important synergy was established with the Horizon Europe FOCCUS project (Forecasting and observing the open-to-coastal ocean for Copernicus users), allowing UC7 developments to be integrated into the generation of novel data products, further expanding the impact and application of the research.

**Table UC7.2** – Summary of UC7 activities and achievements. The Presentations field includes the reached audience. Acronyms are defined in the **dedicated section**.

<b>Annotation</b>	Three labelled datasets published ( <b>Table UC7.3</b> )
<b>Model development</b>	One model developed for each sub-use case ( <b>Table UC7.1</b> )
<b>Scientific publications</b>	One conference paper <sup>16</sup> and one journal paper <sup>17</sup>
<b>Presentations</b>	One for academia, one for regional administration, and two for researchers
<b>Academy</b>	One collaboration agreement, three students recruited and trained, two BSc Theses defended <sup>[18,19]</sup> , and two MSc Theses ongoing
<b>External collaborations</b>	One external collaboration for each sub use case (i.e. three)
<b>Related projects</b>	Synergy with FOCCUS <sup>20</sup>

## Data delivery

For each coastal feature and related task, a labelled dataset was created and published in the Zenodo iImagine community open repository (**Table UC7.3**). Each dataset is supplemented with a README file that enhances its comprehension, elaborating on the

<sup>16</sup> Soriano-González, J., et al. (*in press*). Machine learning-driven shoreline extraction and beach seagrass wrack detection from beach imaging systems. In Proceedings of the 10th International Coastal Dynamics Conference (Aveiro, Portugal).

<sup>17</sup> Azmi, E., et al. (2025). Best practices for AI-based image analysis applications in aquatic sciences: The iImagine case study. Ecological Informatics, <https://doi.org/10.1016/j.ecoinf.2025.103306>

<sup>18</sup> Pérez-Cañellas, J.D. (2024). Detección y segmentación de bermas de *Posidonia oceanica* en imágenes [BSc Thesis – Informatics Engineering (UIB)]

<sup>19</sup> Oliver-Sansó, J. (2024) Monitorització costanera mitjançant intel·ligència artificial: detecció de la línia de costa [BSc Thesis – Informatics Engineering (UIB)]

<sup>20</sup> European Commission, CORDIS. (2023). Forecasting and observing the open-to-coastal ocean for Copernicus users (FOCCUS). European Union. <https://foccus-project.eu/>

context and contents, and providing detailed explanations of technical aspects, annotation process, usage recommendations, potential limitations, and related works.

**Table UC7.3** – UC7 open-source labelled datasets and reference (ref).

Coastal feature	Dataset name (number of images)	ref
Beach wracks	BWILD: Beach seagrass Wrack Identification Labelled Dataset (3286 images)	<sup>21</sup>
Rip currents	RipAID: Rip current Annotated Image Dataset (2815 images)	<sup>22</sup>
Shoreline	SCLabels: Labelled rectified RGB images from the Spanish CoastSnap network (1717 images)	<sup>23</sup>

## Model validation

### *Beach wracks identification*

YOLO11m-seg, trained on BWILD v1.1.0 (**Table UC7.3**), was chosen for detecting and segmenting beach wracks with different density degrees (**Figure UC7.1**). Precision exceeded recall in both tasks, yet *Intermediate wracks* ( $mAP^{50} \sim 0.38$ ) were consistently harder to detect/segment than *Dense wracks* ( $mAP^{50} \sim 0.74$ ). Given the subjective border definition of beach wracks, especially *Intermediate wracks* (sparse/sand-mixed), the results are considered an optimum baseline. While similar prior research exists <sup>[24,25]</sup>, limited training datasets (<200 images) hinder model generalisation and fair comparison.

<sup>21</sup> Soriano-González, J., et al. (2025). BWILD. [Dataset] <https://doi.org/10.5281/zenodo.14834611>

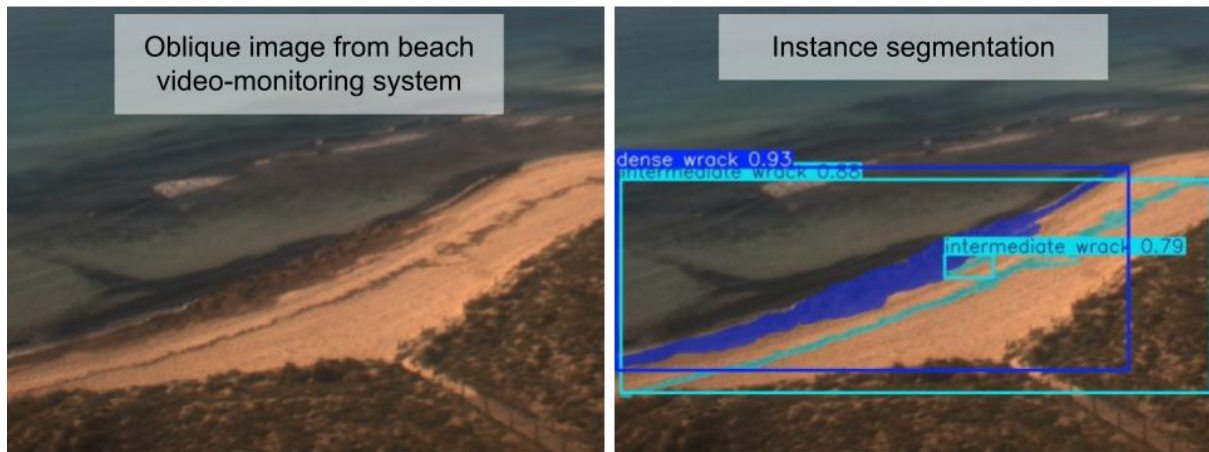
<sup>22</sup> Soriano-González, J., et al. (2025). RipAID. [Dataset] <https://doi.org/10.5281/zenodo.15082427>

<sup>23</sup> Soriano-González, J., et al. (2023). SCLabels. [Dataset] <https://doi.org/10.5281/zenodo.10159978>

<sup>24</sup> Valentini, N., Balouin, Y. (2020). Assessment of a smartphone-based camera system for coastal image segmentation and sargassum monitoring. *J.mar.sci.eng.* 8(1), 23.

<sup>25</sup> Pan et al. (2022). Beach wrack dynamics using a camera trap as the real-time monitoring tool. *Front. Mar. Sci.* 9, 813516.





**Figure UC7.1** – Left: source image. Right: detected beach wracks and masks. 'Dense wrack' indicates high-density, while 'intermediate wrack' signifies low-density or a mix of wrack and sand.

#### Rip currents detection

For rip currents detection, YOLO, trained on a dataset of over 6000 rip current instances (merged RipAID – [Table UC7.3](#) – and reprocessed external datasets), outperformed other architectures (e.g. RTDTR, FCOS), especially with oriented bounding boxes ([Figure UC7.2](#)). The final model is not yet defined, but the strongest candidates with this configuration (YOLOv8n-s, 11n-s, 12s-m) achieve recall and precision  $> 0.8$ , and  $mAP^{50} > 0.9$ . These results are in line with those obtained by YOLO-rip<sup>26</sup> trained on a smaller dataset and indicate that models are confident in detecting most of the rips. YOLO architecture enables high-speed inference, allowing near real-time monitoring for effective complementary beach surveillance systems.



**Figure UC7.2** – Left: source image. Right: detected rip currents.

<sup>26</sup> Zhu, D., et al. (2022). YOLO-Rip: A modified lightweight network for rip currents detection. *Front.Mar.Sci.* 9.

### Shoreline position extraction

The best method for automatic shoreline extraction relied on a 2-step approach: (i) Segmentation of landwards, seawards, and background classes, using a Duck-Net<sup>27</sup> trained on the SCLabels dataset (**Table UC7.3**); (ii) Post-processing to extract the shoreline as the boundary between landwards and seawards classes. The Duck-Net proved to be superior to other architectures (e.g. attention U-Net and Bi-LSTM) for the segmentation task, achieving overall accuracy, precision, recall, and F1-score > 0.98. Since the primary objective of this sub use case was accurate shoreline extraction rather than the segmentation (**Figure UC7.3**), a post-validation exercise was performed. Comparing the extracted shorelines with their labelled counterparts using pixel-distance-based metrics resulted, in most instances, in distances of 2-10 pixels. This establishes a benchmark for the automated extraction of shorelines from complex, smartphone-based rectified imagery, and effectively addresses challenges previously identified in related research<sup>28</sup>.



**Figure UC7.3** – Left: original crowd-sourced submission. Right: rectified image with the extracted shoreline (blue) superimposed.

### Service delivery

The application socib-beach-wracks-identification<sup>29</sup>, which was one of our primary focus, is delivered on the iImagine marketplace and in ai4os-hub GitHub organisation<sup>30</sup>, thus fully available for researchers. The other two services, rip current detection and shoreline extraction, will be published soon following the same service delivery approach.

<sup>27</sup> Dumitru, RG., Peteleaza, D. & Craciun, C. Using DUCK-Net for polyp image segmentation. Sci Rep 13, 9803 (2023). <https://doi.org/10.1038/s41598-023-36940-5> ; <https://github.com/RazvanDu/DUCK-Net>

<sup>28</sup> Soriano-González, J., et al. (2024): From a citizen science programme to a coastline monitoring system: Achievements and lessons learnt from the Spanish CoastSnap network. *Ocean Coast. Manag.* 256, 107280 (2024).

<sup>29</sup> <https://dashboard.cloud.imagine-ai.eu/catalog/modules/socib-beach-wracks-identification>

<sup>30</sup> <https://github.com/ai4os-hub/socib-beach-wracks-identification>



## Findings and Outlook

The iMagine AI platform was used for all the development cycle, from data preparation to model training, tracking, and packaging. A CVAT instance<sup>31</sup>, deployed as a tool in the iMagine platform, was used for collaborative data annotation. Model development experiments were conducted by leveraging virtual machines with specific hardware resources, such as GPUs, deployed directly from the iMagine platform. For experiment management and tracking, we used the MLFlow cloud-instance dedicated to the iMagine project<sup>32</sup>. The platform has been a positive and user-friendly experience for scientific development workflows. Resources were granted to both researchers and students involved in UC7, offering a framework that was otherwise unavailable at our institute or the university. By centralizing maintenance and democratizing access to hardware and services, the iMagine platform streamlines the research process and eliminates unnecessary hardware/expertise over-redundancies, facilitating knowledge generation.

The progress made in UC7 exceeds our initial goals. We achieved the expected results for the prototype, but also developed and matured several models, making them almost ready for integration into an operational and application-oriented frameworks. Our next steps involve enhancing these models, converting data into actionable information, and increasing user engagement. Specific plans include:

- Improving Model Generalizability: Re-training models with expanded datasets (see external collaborations in **Table UC7.2**).
- Generating Ready-to-Use Data Products: e.g. Beach wracks time series<sup>33</sup> (see also 'Related projects' in **Table UC7.2**).
- Implementing (Near)Real-Time Inference: Extending the applicability of modules such as the rip currents detection one.
- Integrating Models into Operational Frameworks: An example is the automatic extraction of the shoreline to enhance current manual procedures.

## UC8 Freshwater diatoms identification

### Description

Diatoms are unicellular microalgae which are recognised as bioindicators of freshwater ecological health within the EU Water Framework Directive (WFD)<sup>34</sup>. They are traditionally identified through microscopic examinations, a time-consuming process requiring a high

---

<sup>31</sup> <https://dashboard.cloud.imagine-ai.eu/catalog/tools/ai4os-cvat>

<sup>32</sup> <https://mlflow.cloud.imagine-ai.eu>

<sup>33</sup> Soriano-González, J., et al. (2025). Beach-imaging derived beach wracks [Data set]. <https://doi.org/10.5281/zenodo.14808058>

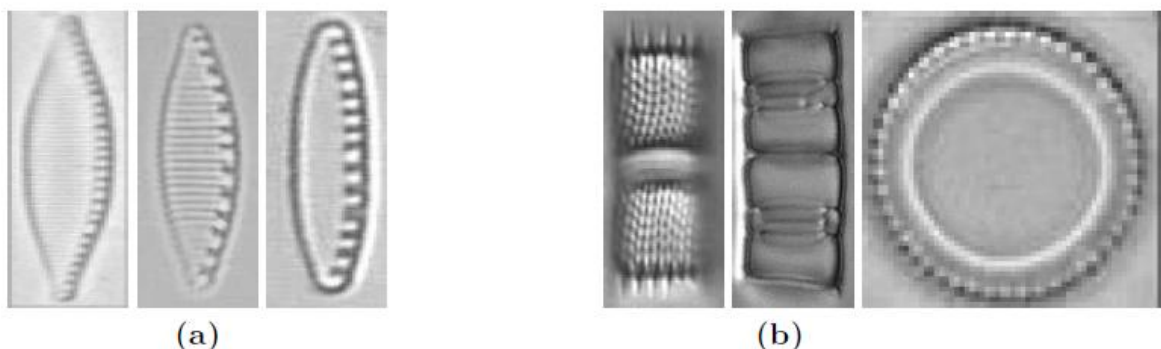
<sup>34</sup> [https://environment.ec.europa.eu/topics/water/water-framework-directive\\_en](https://environment.ec.europa.eu/topics/water/water-framework-directive_en)

expertise in taxonomy. In this context, the use case aimed to develop a prototype using automatic pattern recognition algorithms for microscope images from freshwater benthic diatom samples.

From a computer vision perspective, the main challenge for developing a pipeline for diatom identification lies in the fact that our target organisms (diatoms) include a considerable number of often very subtly different categories, in comparison with more standard problems (see **Figure UC8.1**). Another challenge was the development of sufficiently large and diversified datasets (in terms of numbers of images and categories, respectively) enabling the upscaling of the already existing deep learning models to real-world situations. Finally, beyond the identification of diatom species, the aim of our UC was to explore models for defining new metrics based on morphological criteria (size, deformations) that could be used in bioindication in the future. To achieve this, we employed Convolutional Neural Networks (CNNs). In this context, our objectives were the following:

- Building a prototype end-to-end detection, classification, and trait quantification pipeline, including performance metrics meaningful for diatom experts
- Assembling an extensive quality-controlled dataset for tuning the CNNs
- Deploying the service on the iMagine AI platform

In comparison to existing alternatives, the aim was thus to provide tools that are fast and robust while being reliable with a high level of repeatability, but also with training sets allowing transferability to the greatest number of end-users. The expected impact is the streamlining of diatom identification, addressing the labour-intensive nature of the process required by the WFD, and facilitating wider use by stakeholders and educators.



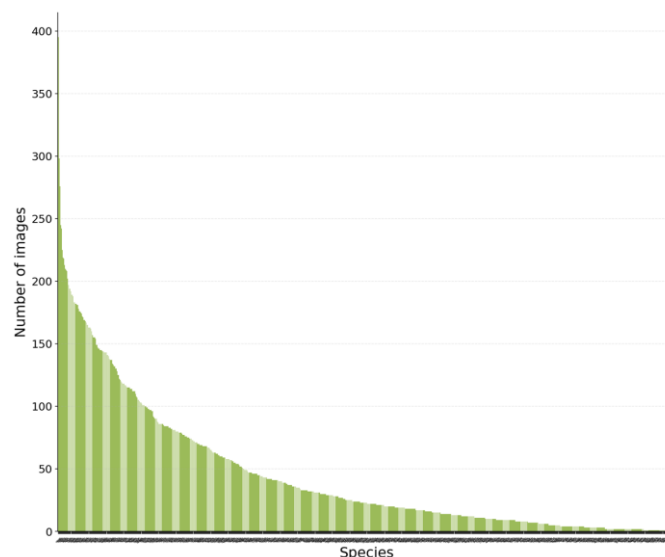
**Figure UC8.1** – Examples of diatom images usually misclassified due to high inter-class similarity (a) or intra-class variance (b). (a) Three different diatom species of the genera *Nitzschia*: *N. soratensis*, *N. subacicularis* and *N. costei*, from left to right. (b) A single diatom species *Aulacoseira pusilla* seen from the side or from the top, from left to right. From Venkataramanan et al. 2021 [https://doi.org/10.1007/978-3-030-87156-7\\_8](https://doi.org/10.1007/978-3-030-87156-7_8)

## Achieved results

### Data delivery

Several datasets were generated: as training sets were limited at the beginning of the project, a first dataset of individual diatom images (“individual dataset”) were gathered to generate virtual raw microscope images (synthetic dataset) simulating real images (“real dataset”).

*Individual dataset* – Individual images of diatom (thumbnails) were gathered from taxonomic atlases which were available online as open-source PDF (e.g. atlas Rhône-Alpes<sup>35</sup>). In total, the dataset represents a collection of 38,719 images (ca. 30kb/image) of various sizes and various scales, representative of 119 genus and 799 species, with 1 to 395 images per category (highly unbalanced dataset: **Figure UC8.2**). This dataset was used to generate a synthetic dataset to pre-train detection and classification algorithms in the initial phase of our project (see below). It can also be used directly to train a classifier. A first subset (9230 images, 166 diatom species) was released on our institutional repository (DOREL)<sup>36</sup>. The full version of the dataset was released on Zenodo<sup>37</sup> and will be updated beyond the lifetime of the project as soon as new atlases are available.



**Figure UC8.2** – Image distribution across the 799 diatom species of our “Individual” dataset.

*Synthetic Dataset* – This dataset was generated to pre-train the CNNs as an alternative to the lack of annotated real images. It was created by:

<sup>35</sup> [http://webissimo.developpement-durable.gouv.fr/IMG/pdf/Tome1\\_2013\\_print\\_cle1aalad.pdf](http://webissimo.developpement-durable.gouv.fr/IMG/pdf/Tome1_2013_print_cle1aalad.pdf)

<sup>36</sup> <https://doi.org/10.12763/UADENQ>

<sup>37</sup> <https://doi.org/10.5281/zenodo.16260887>

1. collecting the individual images of diatoms (individual dataset) but also debris from real images,
2. applying data augmentation by varying the size and orientation of the thumbnails, and
3. generating as many virtual microscope images (ca. 250 kb/synthetic image) as we need using seamless techniques to paste the thumbnails onto a gray background, thereby mimicking a realistic set of microscope images containing various diatoms and debris<sup>38</sup>.

*Real dataset* – This dataset consists of real raw microscope images acquired from field samples. It is currently made of 2,867 images (ca. 1.1 MB/image) of size 2080 x 1042. We used BIIGLE to annotate ca. 10k oriented bounding boxes at genus and species level. This dataset will be released on Zenodo. This training set was used to train the first detection and classification algorithms. The same images are currently annotated with instance segmentation masks using LabelBox and, more recently, SAM2. However, this dataset has not yet been publicly released.

#### Model validation

Several models were developed in the framework of Aishwarya Venkataramanan PhD thesis<sup>39</sup> using on-premise infrastructure. UC8 has mainly focused on the development of an end-to-end pipeline for diatom detection and taxonomic classification using a probabilistic approach, which was validated internally. YOLO-based detection and classification models were further deployed on the iMagine platform and were validated by diatom experts, including a few external users. Regarding the approach for morphological analysis, a first pipeline based on instance segmentation has been developed, thereby enabling the quantification of handcrafted size-related descriptors (major axis, minor axis). This approach will be compared with unsupervised methods (e.g. GANs) capable of exploring the morphological variability of diatoms in an unsupervised manner.

*Diatom detection and classification* – The proof of concept for the prototype was initially established by pre-training the models using a synthetic dataset and then fine-tuning them with a more limited real dataset. Using the synthetic dataset, the performance of the detection network (YOLOv5) improved by up to 25% in precision and 23% in recall at an IoU threshold of 0.5 ([Venkataramanan et al. 2023](#)).

For taxonomic classification, a first version of the individual dataset (ca. 20,000 individual diatom images representing 197 diatom species) was used to train a deep learning classifier with EfficientNet as the backbone. The classifier was evaluated based on the

---

<sup>38</sup> [Venkataramanan et al. 2023](#)

<sup>39</sup> [Venkataramanan 2023](#)

accuracy score, which was 94%. Of the approximately 200 diatom species included in the dataset, 113 were classified with 100% accuracy. Additionally, other classification methods were explored to better address the high inter-class similarity and intra-class variance among the different diatom species. These factors can make it challenging for traditional classification methods to accurately distinguish between species. To tackle this issue, a method was proposed for learning feature representations that group visually similar images of each class together while ensuring that inter-class features are widely separated<sup>40</sup>. Furthermore, a method for estimating uncertainty in classification performance was introduced. This method uses the proximity of a data point to different class features to estimate uncertainty in the network's prediction. It also demonstrates how this approach can provide a reliable estimate of prediction confidence and detect out-of-distribution samples<sup>41</sup>. This method has also been tested on another dataset of diatom images recently published by external collaborators<sup>42</sup>.

The code base for diatom detection and classification was then improved for the development of a first GUI prototype based on YOLOv5 with oriented bounding boxes and a web frontend<sup>43</sup>. The project was finally ported to the iMagine platform with the ultralytics<sup>44</sup> implementation of YOLOv8. For each training, the dataset was split into 80% / 20% for training/validation. Validation of the models was performed from the validation metrics (20% of the data). The quantitative metrics reported below are always on the validation folds. For the detection task, the real dataset was used, with the labels being binarized (single object class). Several models were trained (YOLOv8m, YOLOv8x, YOLOv11n, YOLOv11x) and training took less than 3 hours on a RTX 3090 reaching a mAP50-95 from 0.8 to 0.84. For classification at either the genus or species level, the individual dataset was used after excluding classes with fewer than 10 images, which resulted in 108 classes (ranging from 11 to 4390 images/class) and 547 classes (ranging from 10 to 395 images/class), respectively. As these images were taken at different resolutions, a preprocessing step was necessary. All the images have been resized to 512 x 512 with a resolution of 0.1  $\mu\text{m}$ /pixel. Each training performed on an RTX 3090 took between 1 and 2 hours. At the genus level, a YOLOv8l-cls reached a TOP5 accuracy of 99.8% and a TOP1 accuracy of 97.6%. At the species level, a YOLOv8l-cls reaches a TOP5 accuracy of 99.4% and TOP1 accuracy of 94.2%.

*Diatom segmentation* – The instance segmentation dataset is currently under development. It will be employed to test a pipeline based on instance segmentation using YOLO and morphological parameter extraction from the obtained segmentation masks.

---

<sup>40</sup> [Venkataramanan et al 2021](#)

<sup>41</sup> [Venkataramanan et al 2023](#)

<sup>42</sup> [Venkataramanan et al 2024](#)

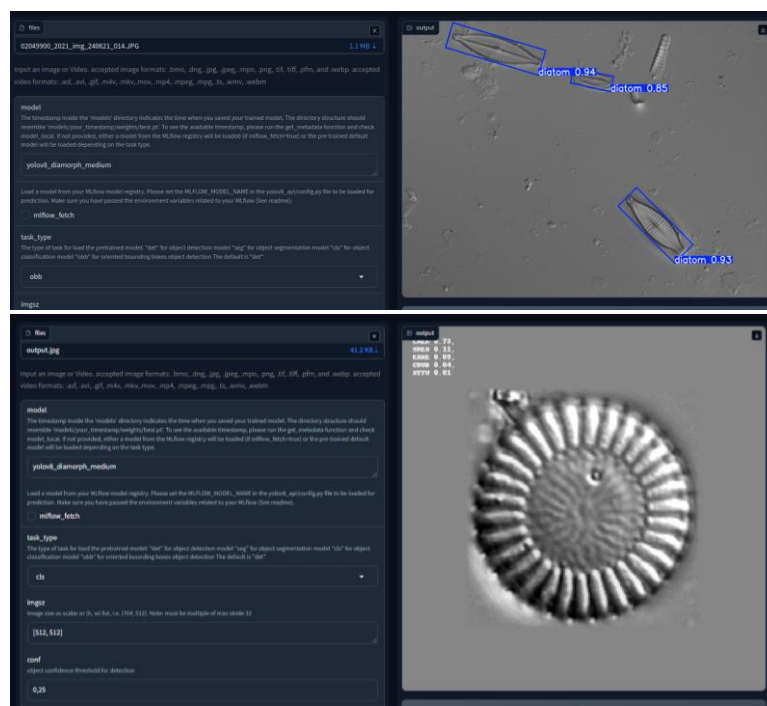
<sup>43</sup> <https://www.ins2i.cnrs.fr/fr/diamorph-identification-automatique-des-diatomees-par-apprentissage-profond-partir-dimages>

<sup>44</sup> <https://docs.ultralytics.com/>

As an example, a first dataset focusing on a single diatom species (*Nitzschia palea*)<sup>45</sup> has been used to illustrate how this approach will further support biologists in their morphometric analyses (Chéron et al. 2025<sup>46</sup>). In parallel to this pipeline, a first approach to quantify diatom deformations has been explored. It is based on Generative Adversarial Networks (GAN) models (StyleGAN3+encoder4editing<sup>47 48</sup>) for unsupervised estimation of the variation of morphological features.

## Service delivery

The detection (oriented bounding boxes) inference pipeline is released under the AGPL Licence in the ai4os-hub GitHub organisation under AGPL-3.0 Licence under the diatom detection application<sup>49</sup>. It is also available on the iImagine marketplace<sup>50</sup>. The application offers several endpoints outputting the detected objects as a JSON file or even directly drawing the detected objects on the image (Figure UC8.3). The classification inference pipeline at the species level is released in the AI4OS catalogue under the diatom classification application<sup>51</sup> (see Figure UC8.3 of the corresponding “Try-me” deployment).



**Figure UC8.3** – Screenshots of the diatom detection and classification applications as available online

<sup>45</sup> <https://doi.org/10.5281/zenodo.15866234>

<sup>46</sup> <https://hal.univ-lorraine.fr/hal-05168372>

<sup>47</sup> <https://doi.org/10.48550/arXiv.2106.12423>

<sup>48</sup> <https://doi.org/10.48550/arXiv.2102.02766>

<sup>49</sup> <https://github.com/ai4os-hub/diamorph-detection>

<sup>50</sup> <https://dashboard.cloud.imagine-ai.eu/catalog/modules/diamorph-detection>

<sup>51</sup> <https://github.com/ai4os-hub/diamorph-classification>

## Findings and Outlook

Feedback on the experience with the platform by end users depends on their level of IT expertise. Several limitations have been identified, which can be overcome in the near future.

Regarding deployment, the AI4OS platform was used only for deploying the inference endpoints, as we already had the resources to train our models. The deployment of the two applications (bounding box detection and classification) was straightforward using `ai4os-yolov8-torch`<sup>52</sup> child modules. The two applications are currently disconnected, as they consider different types of images (full microscope images for detection, thumbnails for classification). Our future plans include re-training models with expanded real datasets, thereby enabling the pipelining of the two applications. Future plans also include the implementation of the approaches developed within the framework of A. Venkataramanan's PhD thesis (hierarchical classification, probabilistic approach).

For domain scientists with some degree of fluency in IT, the experience is overall positive, but several limitations have been identified. The default values of the parameters on Gradio<sup>53</sup> are not the right ones, and some parameters are listed on Gradio but not necessary for classification. Nevertheless, it is possible to create a custom Gradio interface and configure it to run on port 80 on the dashboard, allowing for better control over the user experience and parameter selection. In that sense, the ergonomics of the classification marketplace application could be improved. Also, a required preprocessing step must be performed by the user before submitting the image to be classified. These limitations could be addressed by implementing a lower-level DEEPaaS application rather than a `ai4os-yolov8-torch` child module.

For our other targeted end users, i.e. diatom experts with minimal IT experience, the platform has demonstrated its usefulness in showcasing our project (i.e. during workshops or conferences) but remains too complex to engage more actively with this community of users. However, even if the platform is not being used directly at the moment, it could be envisaged to connect it to applications with a more user-friendly frontend, such as the ones that have been developed in other parallel but related projects. For example, our prototype GUI<sup>54</sup> could benefit from the platform's computing environment if connected via an API that would trigger the deployment and execution of the application. Another example is the use of our datasets and the resulting trained models, which have enabled the development of a web platform aiming at collecting pictures of deformed diatoms<sup>55</sup>.

---

<sup>52</sup> <https://github.com/ai4os-hub/ai4os-yolov8-torch>

<sup>53</sup> <https://www.gradio.app/docs>

<sup>54</sup> <https://www.ins2i.cnrs.fr/fr/diamorph-identification-automatique-des-diatomees-par-apprentissage-profond-partir-dimages>

<sup>55</sup> Vallanzasca et al. 2023 <https://doi.org/10.1080/23818107.2023.2183898>



Ultimately, the new images collected on this online platform will contribute to the continuous improvement of the algorithms, particularly for deformation analysis.

Overall, the use of the platform has led to significant advances in our use case. These advances have certainly contributed to the leverage effect that led to the funding of a follow-up project by the French National Research Agency<sup>56</sup>. This funding allows our use case to take a significant step forward in terms of maturation, using the iMagine platform.

## Additional use cases

In addition to the project use cases of different maturity, three open calls were organised in the course of the project to engage external users and to support a broader aquatic community. Here we summarise the main findings of these additional use cases.

### EyeOnWater

The use case was brought by one of the iMagine partners, MARIS, and enhanced the EyeOnWater citizen science platform by using AI to automatically validate smartphone images of water bodies. To have an extra helping hand in determining whether an image meets the criteria for inclusion in the app, improve the Citizen Science workflow and enhance the credibility of the data, EyeOnWater has incorporated AI image recognition technology into its operational infrastructure. The YOLOv8 model was chosen, and with the help of this model, all uploaded pictures are assessed. To train this model, a training dataset containing a large pool of different images was created. The trained model is made available on the iMagine marketplace<sup>57</sup> and shared on GitHub, making it possible for any user to validate their water surface images on the fly.

### Cold water coral reefs (call 1)

The use case aimed to develop a non-destructive method for evaluating the coverage of cold-water coral reef species *Desmophyllum pertusum* and *Madrepora oculata* in the North Atlantic, specifically within Natura 2000 network areas of the Bay of Biscay. During the platform usage period, significant progress was made in developing a robust and efficient pipeline for annotating and training deep learning models for cold-water coral reef analysis, using a YOLOv8 model. The results have been compiled into a paper titled 'Deep Learning Based Characterization of Cold-Water Coral Habitat at Central Cantabrian Natura 2000 Sites Using YOLOv8'<sup>58</sup> and the trained model is made available on the iMagine marketplace<sup>59</sup>.

---

<sup>56</sup> <https://anr.fr/Projet-ANR-24-CE04-0345>

<sup>57</sup> <https://dashboard.cloud.imagine-ai.eu/catalog/modules/eyeonwater-yolo>

<sup>58</sup> <https://www.mdpi.com/2077-1312/12/9/1617>

<sup>59</sup> <https://dashboard.cloud.imagine-ai.eu/catalog/modules/cold-coral-segmentation>



## Satellite-derived bathymetry (call 1)

Nearshore bathymetry is one of the most important parameters for coastal studies, whose application extends to different areas. Satellite-derived bathymetry offers an attractive alternative to conventional collection methods for acquiring freely available, reliable, and frequent high-resolution data from space satellite missions such as Sentinel-2A/B. The objective was to fully automate the process from the download of satellite images to the final production of bathymetric maps, using the latest machine learning (ML) and deep learning (DL) techniques. The automation of the bathymetric mapping estimation process was successfully established, and the results were published in the paper entitled 'Blending physical and artificial intelligence models to improve satellite-derived bathymetry mapping'<sup>60</sup>.

## Age reading from fish otoliths (call 2)

Age determination of fish is fundamental to the evaluation of fish stock status as age data underpins the growth and mortality rates, as well as maturity patterns. However, age determination is a highly subjective process. Otolith images have been increasingly used for scientific studies. Having a reliable AI algorithm for age determination means that the subjectivity of the process will be removed. The main objectives of this case study are to develop and test such AI algorithms in an image-based age reading calibration event with expert age readers and carry out a statistical analysis showing levels of precision and bias. Two balanced training datasets were prepared for improved model training. YOLOv8 was trained on these two datasets for detecting and classifying otolith ages in images. The newly trained model performed better than the original approach based on either the Sequential Model, Inception, or VGG16 neural network.

## DEcentrAlised Learning (DEAL) (call 3)

The DEAL (DEcentrAlised Learning for automated image analysis and biodiversity monitoring) project<sup>61</sup> will create an application that allows owners of biological image data to participate in decentralised, collaborative networks, where they can leverage the data and expertise of all participants to obtain better, higher efficacy classification results for their data. DEAL partnered with iImagine to in early 2025 through the creation of a FlowCam Working Group with the following objectives:

- Incorporate LifeWatch observatory data in DEAL
- Incorporate the iImagine best practice framework into the DEAL architecture
- Run a prototype / MVP<sup>62</sup> or DEAL node to classify FlowCam data on the iImagine Platform

---

<sup>60</sup> <https://doi.org/10.1016/j.jecoinf.2025.103328>

<sup>61</sup> <https://pml.ac.uk/projects/deal-decentralised-learning-for-automated-image/>

<sup>62</sup> Minimum Viable Product

A major outcome of the collaboration was the discovery and use of the VLIZ LifeWatch FlowCam dataset. Introduced via the iImagine network, this dataset was combined with PML's own FlowCam data, creating a unified and technically compatible resource. The iImagine platform itself also served as a key enabler. With its integrated use of NextCloud for cloud-based data storage and Zenodo for dataset archiving, the platform offered a collaborative and accessible environment for managing data and coordinating work across DEAL's distributed team. Looking forward, the DEAL team plans to continue leveraging the iImagine platform to support development and collaboration.

## Sea Wave and coastal inundation detection Methodology (SWiM) (call 3)

SWiM aims at developing a methodology to extract information from images taken by cameras installed overlooking a beach. Specifically, SWiM will determine the waterline, the wave runup, and the wave breaking line. The new procedure for extracting this information from images is being now tested on the iImagine Cloud Infrastructure, and using a batch of high-resolution, high-frequency, sequential images obtained from an existing camera at São Pedro de Moel in Portugal, from a 6-month period (January–June 2022). The model is being developed in close collaboration with coastal engineers, who assess the results. In addition, several meetings and discussions were held with a Spanish research group (SOCIB, UC7) working along similar research lines.

## Conclusion

The deliverable summarizes the work and validation of the three prototype use cases of iImagine: UC6 (Underwater Noise Identification), UC7 (Beach Monitoring), and UC8 (Freshwater Diatoms Identification), which started as low maturity prototypes, and successfully developed to AI-powered applications delivered to aquatic community and put it into action as a part of the iImagine project. In addition, the project onboarded and helped further six external use cases in the domain.

Based on the use cases reports, UC6 made full use of the iImagine platform from development to delivery and reported a positive experience. The dataset is open-access on the Zenodo iImagine community page, and the model is also published on GitHub (ai4os-hub) and it is available for inference on the iImagine dashboard.

UC7 used a wide range of iImagine tools across the entire pipeline from data annotation to training, experiment tracking, and service deployment. UC8 used the platform for the deployment of models trained on their own resources. Their experience helped highlight areas for UI and usability improvements. Overall, the use of the platform has led to significant advances in UC8 and helped to receive further funding in France.

The iImagine platform supports each use case by providing services and tools for every stage of the AI model development lifecycle and complements them with best practices and activities of the Competence Centre. The feedback and knowledge exchange from each use case helped all use cases to advance in their development and the platform developers to improve their support and refine the tools and services further.