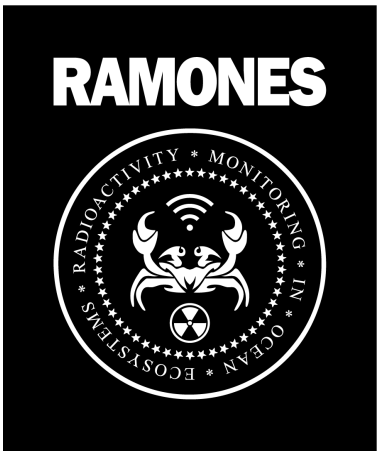

RAMONES
Radioactivity Monitoring in Ocean Ecosystems

Deliverable

Document ID: D4.2 1st release of AI-based solutions for time series and imaging data analysis

05/05/2023

RAMONES funded by European Union under Horizon 2020 FET Proactive programme via grant agreement No.101017808

Document Info

Project Information			
Acronym	RAMONES		
Name	Radioactivity Monitoring in Ocean Ecosystems		
Start Date	1 Jan 2021	End Date	31 Dec 2024
Program	H2020-EU.1.2.2. - FET Proactive		
Call ID	H2020-FETPROACT-2020-2	Topic	FETPROACT-EIC-08-2020 - Environmental Intelligence
Grant No	101017808	Instrument	RIA

Document Information			
Document Id	D4.2		
Document Title	1st release of AI-based solutions for time series and imaging data analysis		
Due Date	28-Feb-2023	Delivery Date	05-May-2023
Lead Beneficiary	NTUA (5)		
Beneficiaries (part.)	NKUA (1), IST-ID (2), PLOA (6)		
Editor(s)	Konstantinos Karantzalos (NTUA)		
Authors (s)	Valsamis Ntouskos, Konstantinos Karantzalos, Melkon Chatsikian (NTUA)		
Contributor (s)			
Reviewer(s)			
Workpackages	WP4 - Multi-modal Data Analytics and Environmental Modelling		
Version	v1.0	Stage	Submission
Distribution	Public		
Keywords	Pathfinder, EIC, FET, Radioactivity, Environmental Intelligence, Underwater Research, Architecture, γ-radiation, Spectrometers	Type	Other

Document Change Record

Version	Date	Change Description	Editor
0.1	1/3/2023	First draft report's structure	NTUA
0.2	20/4/2023	Updates after the Milos Field trip	NTUA
0.3	5/5/2023	Revisions and updates in all subsections	NTUA
1.0	12/7/2023	Document version submitted to EC	

Disclaimer

RAMONES is a European Innovation Council (EIC) FET Proactive project in the Environmental Intelligence Scope B, related to radically novel approaches to resilient, reliable, and environmentally responsible in-situ monitoring, funded by European Union under Horizon 2020 FET proactive programme, via grant agreement No. 101017808.

RAMONES project's main objective is to close the current marine radioactivity under-sampling gap and foster new interdisciplinary research in ocean ecosystems. RAMONES will invest a significant effort to provide tools to enable long-term data acquisition missions, rapid deployments, low cost per information byte, and propose new AI and Robotics-driven and supported methodologies, being ambitious to eventually offer scaled-up solutions to researchers, policy makers and communities. All these may be achieved by combining state-of-the-art (SoA) methodologies and equipment from various disciplines in a well-balanced synergy and designing new and effective methodologies targeting the marine environment, which will provide efficient response to existing natural and man-made hazards, and shape future policies for the global population. RAMONES will additionally contribute to shaping a blueprint on Environmental Intelligence in the EU and worldwide.

This document has been produced through funding from the European Commission. The content of this document is a product of the RAMONES project Consortium, and it does not necessarily reflect the opinion of the European Commission. The editor, author, contributors, and reviewers of this document have taken any available measure in order for its content to be accurate and lawful. However, neither the project consortium as a whole, nor the individual partners that have implicitly or explicitly participated in the creation and publication of this document may be held responsible for any damage, financial or other loss or any other issue that may arise as a result of using the content of this document or any of the project outputs that this document may refer to.

The European Union (EU) was established in accordance with the Treaty on the European Union (Maastricht). There are currently 27 member states of the European Union. It is based on the European Communities and the member states' cooperation in the fields of Common Foreign and Security Policy and Justice and Home Affairs. The five main institutions of the European Union are the European Parliament, the Council of Ministers, the European Commission, the Court of Justice, and the Court of Auditors (http://europa.eu.int/).

List of acronyms

Acronym	Description
AI	Artificial Intelligence
ADC	Analog-to-Digital Converter
API	Application Programming Interface
ASIC	Application-Specific Integrated Circuit
AttMask	Attention-guided Masking
CPS	Counts per second
LLD	Low-Level Discriminator
MIM	Masked Image Modeling
MLP	Multi-Layer Perceptron
NORM	Naturally Occurring Radioactive Material
PSNR	Peak Signal-to-Noise Ratio
(R)MSE	(Root) Mean Squared Error
ROV	Remotely Operated Vehicle
SoA	State of the art
SUGI	Submarine Gamma Imager
TCP/IP	Transmission Control Protocol / Internet Protocol
UDP	User Datagram Protocol
UTC	Coordinated Universal Time
ViT	Vision Transformers

Table of Contents

Abstract

This deliverable presents the AI-based software solutions developed for processing the data produced by the RAMONES instruments. Specifically, software for preprocesing and analysing radioactivity measurements (time-series) collected from the γSniffers have been developed. Additionally, a custom API and dedicated software modules have been developed for collecting and processing individual radioactivity detection events registered by SUGI, towards radioactivity imaging using Compton imaging. The document also presents a high-level simulation environment that has been built to study the ability to detect and map radioactivity in large volumes from the measurements collected by γSniffers mounted on the underwater gliders, under different considerations regarding the trajectories of the gliders, their spacing and the form of the radioactivity distribution.

1. Introduction

1.1. Context

This report is the deliverable D4.2 "*1st release of AI-based solutions for time series and imaging data analysis*" of the H2020 RAMONES project. The report is a product of the work performed in WP4 "*Multi-modal Data Analytics and Environmental Modelling*" in a tight conjunction with the concurrent activities performed also in *WP1 "Design and Develop Novel Underwater Gamma Radiation Instruments*", WP2 "*Autonomous Marine Robotics for Collaborative Radioactivity Mapping*" and WP3 "*Onsite Integration, Testing and Pilot Demonstrators*".

1.2. Contents and Rationale

The main purpose of the present document is to present the software solutions that have been developed for analysing the data collected from the instruments developed so far in the project. The developed solutions consider the characteristics of each instrument and the corresponding data modalities. Specifically, γSniffers produce time-series reporting number (counts) of detected gamma ray events in a specific time interval, potentially with associated photon energies. SUGI, besides number of events, provides imaging information, thanks to its pixelated detector modules.

A novel high-level simulation environment is described first, which allows to simulate the data collected by the γSniffers mounted on the underwater gliders, considering basic models for radioactivity interactions with matter under reference distributions, for typical forms of glider trajectories used to scan the volume of interest. Subsequently, AI-based methods developed for time-series data analysis are first described, followed by methods for processing the data collected by the imaging sensors. The document also presents the application of the developed software solutions for the analyses of data collected from γSniffer and SUGI during the field deployment at Milos Island, Greece.

1.3. Structure of the document

Chapter 2 presents the high-level simulation developed for studying detection and mapping of radioactivity using underwater gliders. Chapter 3 describes the AI-based solutions developed for the processing and analyses of time-dependent radioactivity measurements, while Chapter 4 discusses the algorithms developed to support radioactivity imaging. Finally, Chapter 5 provides concluding remarks and discusses the steps towards the final release of the AI-based algorithms.

2. High-level Simulation for Multi-modal Data Analytics and Environmental Modelling

In this chapter, a high-level simulation environment is presented that has been developed for providing simulated radioactivity measurements for different types of radioactive sources considering typical trajectories followed by underwater gliders, based on their kinematic model. In particular, simplified models are considered both for the radioactivity measurement and the movement of the gliders (e.g., absence of localization uncertainty). The goal is to provide an environment where different radioactivity detection and mapping scenarios can be tried out, allowing to choose suitable scanning patterns and better understanding limitations due to challenges from radioactivity shielding of the water and the significantly small coverage of the area of interest. The simulator has been developed in Python, both for the graphical representations, and for the numerical simulations of the measured radioactivity values.

2.1. Environment and Glider Trajectories

A reference volume of interest has been defined in the form of a three-dimensional cuboid with a size of 3000x3000x300 meters (x, y, z axes, respectively), as shown in Figure 1. The simulation is performed in a 1:1 scale.

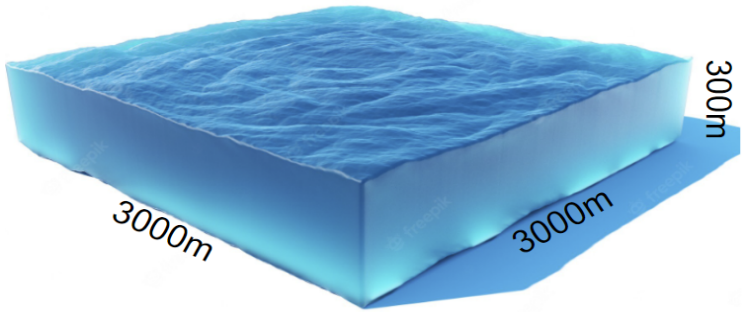

Figure 1: Volume of interest

In this reference volume, two types of underwater glider trajectories are considered. The first corresponds to diagonal (zigzag) vertical motion and the other to spiral movements, as shown in Figure 2, Figure 3 and Figure 4. The simulation considers different spacing between the trajectories. For the diagonal movement, the horizontal distance between two trajectories is

defined as trajectory spacing. For the spiral movement, the original volume is divided into smaller sub-volumes, based on a horizontal spacing. The radius of the spiral movement is defined as ¼ of the size of this horizontal spacing. The vertical movement of the underwater gliders is limited to the range [25, 285], considering suitable margins for safe operation.

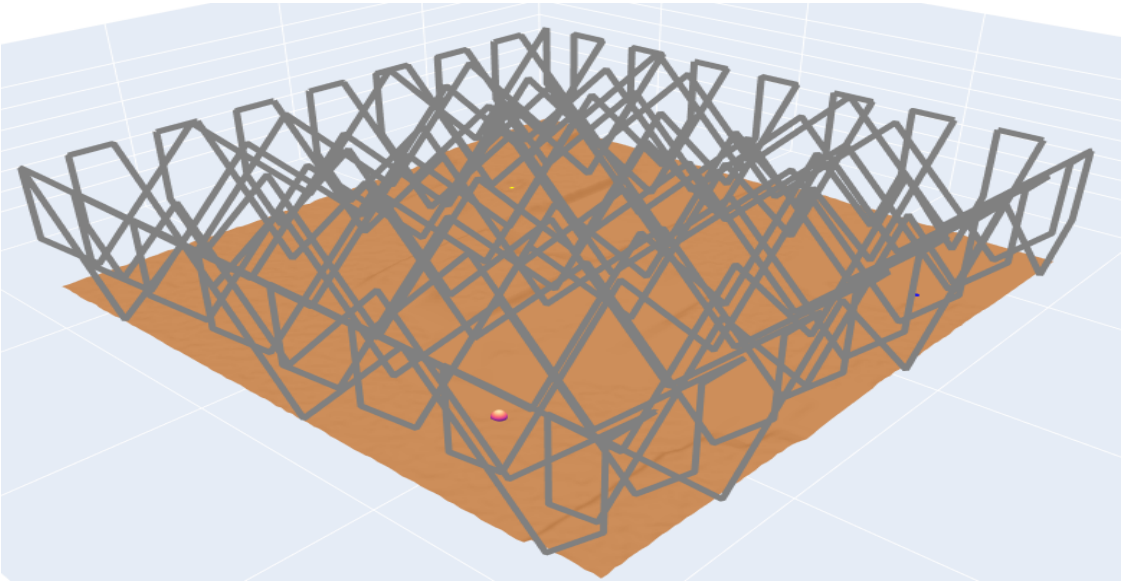

Figure 2: Presentation of diagonal trajectories in the simulation environment

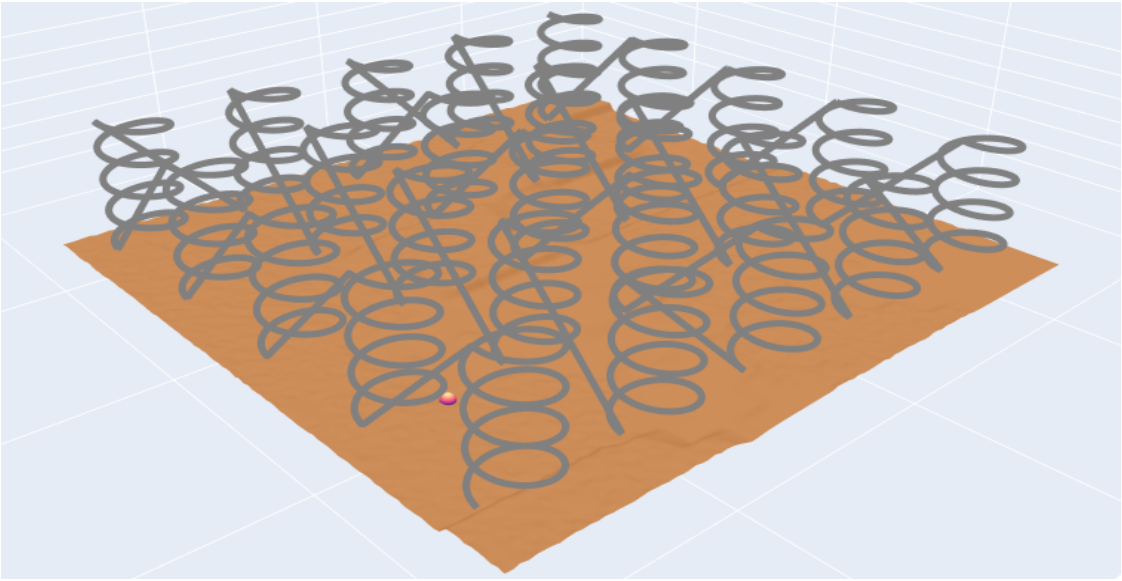

Figure 3: Presentation of spiral trajectories in the simulation environment

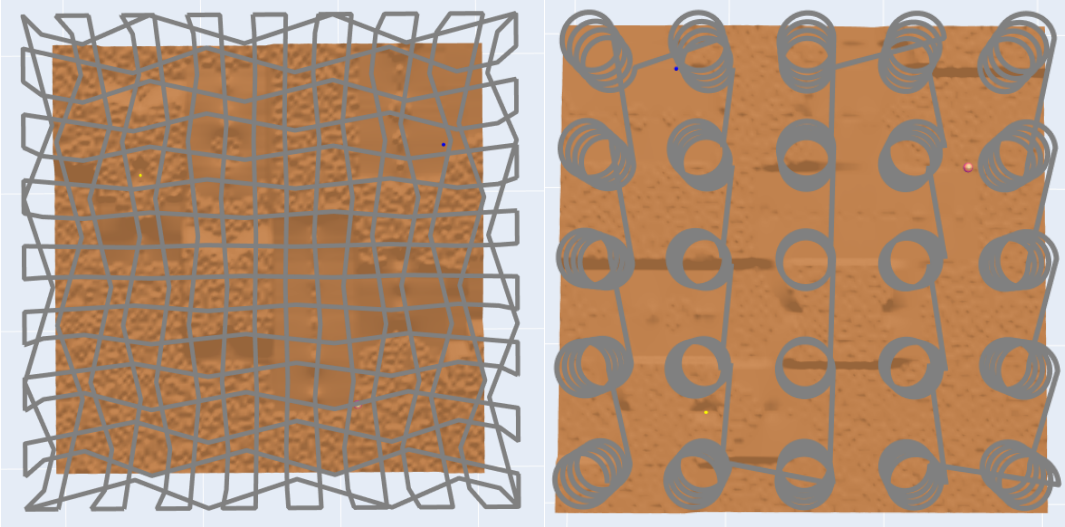

Figure 4: Top-view of diagonal and spiral trajectories in the simulation environment

The type of glider movement has an important impact on the time required to execute the trajectory and to the percentage of coverage of the volume of interest. The simulation environment allows to estimate the time needed to scan the entire volume. Given a specific type of movement, providing the value of the trajectory spacing, the volume is filled with the planned trajectories, allowing to compute total length of the resulting scan-path. Dividing the scan-path length by the speed of underwater gliders (e.g., v=0.5 m/s) the duration of the scan can be estimated. Additionally, considering a maximum detection range for the radioactivity detectors mounted on the gliders (e.g., r=1m based on preliminary results for point sources), the volume of the space measured can be reported in relation to the total scan volume. Table 1 and **Error! Reference source not found.** report indicative values for v=0.5 m/s and r=1m.

Table 1: Characteristics of diagonal trajectories

Trajectory Spacing (m)	Total Distance (km)	Estimated Time	Coverage Rate (%)
100 x 100	221.2	5d & 2:53h	0.025
150 x 150	151.8	3d & 12:19h	0.017
200 x 200	117.1	2d & 17:02h	0.013
250 x 250	96.3	2d & 5:28h	0.011

Table 2: Characteristics of spiral trajectories

Voxel Spacing (m)	Total Distance (km)	Estimated Time	Coverage Rate (%)
200 x 200	362.2	8d & 9:13h	0.042
300 x 300	229.6	5d & 7:32h	0.026
400 x 400	146.7	3d & 9:30h	0.017
500 x 500	133.2	3d & 2:00h	0.015

2.2. Types of Radioactivity sources and Measurement Models

Different types of radioactivity sources are considered, ranging from point sources to radioactivity distributions covering partly or entirely the volume of interest.

Point sources

This simplified model can be used as a good approximation when the size of the source is much smaller compared to the distance at which measurements are recorded. Due to the reduction of the solid angle in which the detector observes the source, the intensity of radioactivity emitted by a point source and reaching the detector is subject to the inverse square law. Therefore, as we move away from the point source, the intensity of the radioactivity decreases according to the square of the distance. Another factor that significantly affects radioactivity detection is shielding from the material between the source and the detector. Due to shielding, the intensity of radioactivity decreases exponentially according to an attenuation factor that depends on the density and the physical properties of the material in which it transports. There are other factors that affect radioactivity detection, such as build-up factor, which were not considered as they contribute less to the simulated measurements.

Based on these factors, the intensity of the measured radioactivity I_R at a certain point at a distance R (for R a few times larger than the detector size) from a point source of intensity I_0 can be approximated by the following formula:

$$I_R = \frac{I_0}{R^2} \cdot e^{-\mu \cdot R} \quad (1)$$

Table 3: Comparison of radioactivity attenuation in air and in water

Distance [m]	Attenuation factor	
	Air (ρ=1.225 kg/m^3) (μ/ρ=0.0636cm^2/g, $\mu \simeq$0.0078m^{-1})	Water (ρ=1.02 g/cm^3) (μ/ρ=0.0707cm^2/g, $\mu \simeq$7.2m^{-1})
0.25	0.998	0.165
0.5	0.996	0.027
1	0.992	7.4e-4
2	0.985	5.4e-7
5	0.962	2.2e-16
10	0.925	4.8e-32

Table 3 provides a comparison of the attenuation factor in the air and inside the water, considering shielding alone. One can see that detection of point sources in the water is extremely challenging due to the very effective shielding of radioactivity from water.

Distribution Sources

Another scenario of interest is the mapping of radioactivity in the water column in large areas. To examine this possibility, we consider different distributions of radioactivity in the water, and use the simulator to examine how accurately these distributions can be reconstructed from radioactivity measurements acquired from an underwater glider traversing through them, considering different scan strategies.

The following reference distributions were examined, which offer different characteristics in terms of geometry and spatial frequencies. For simplicity, the detector is assumed to directly measure the quantity of interest, without any loss or noise affecting the measurement. In case the distribution does not cover the entire space, a static background 10 times smaller than the signal is assumed.

Analytical distributions

Analytical distributions have been considered, as they are defined by mathematical formulas, hence their properties can be accurately recovered.

One reference distribution examined is the following:

$$I(x,y,z) = \sin\sqrt{\left(\frac{x-1500}{150}\right)^2 + \left(\frac{y-1500}{150}\right)^2 + \left(\frac{z-150}{15}\right)^2} \quad (3)$$

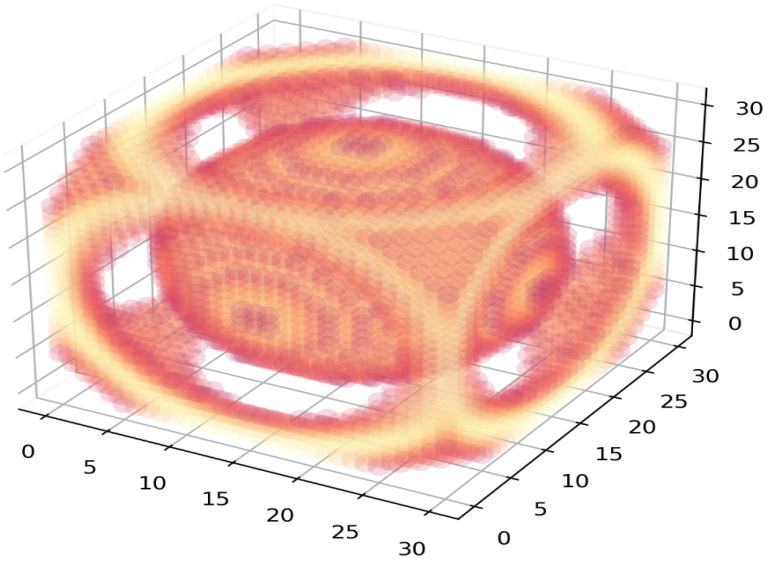

Figure 5: Graphical representation of relation (3)

In Figure 5, the graphical representation of the reference distribution of relation (3) is presented. We notice that the specific function completely covers the area of interest, having several changes in terms of the returned values depending on the position.

One more reference distribution examined is the one presented in Figure 6 and relation (4). This function also covers the entire volume of interest. It follows a pattern as the values increase as the z-axis value increases.

$$I(x,y,z) = \sqrt{\left(\frac{x}{3000}\right)^2 + \left(\frac{y}{3000}\right)^2 + \left(\frac{z}{300}\right)^2} \quad (4)$$

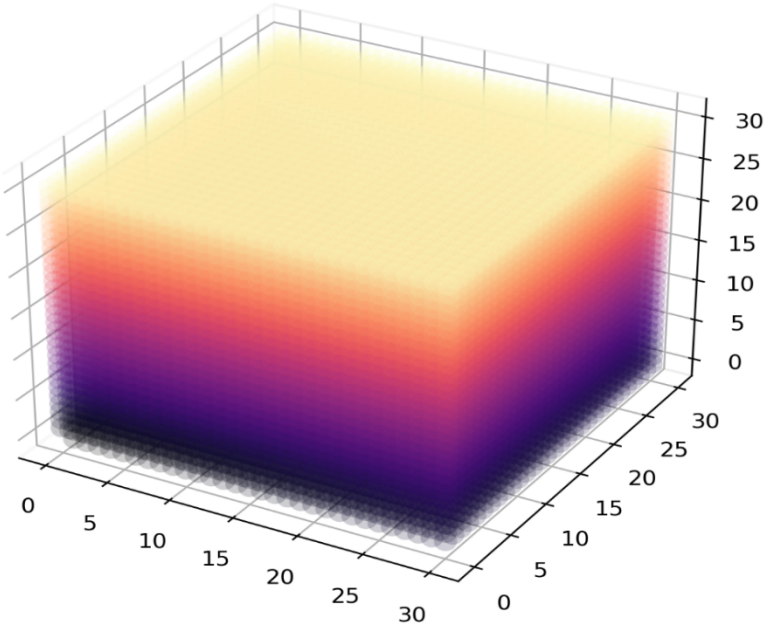

Figure 6: Graphical representation of relation (4)

Diffusion-based sources

Another important scenario considers the diffusion of radioactive material around a source, as for example in the case of a hydrothermal plume. Fick's law describes the flow of particles or some substance in a propagation medium. According to Fick's law, the rate of diffusion depends entirely on the degree of concentration of the substance. Letting J be the diffusion rate, D the diffusion coefficient, C the concentration of the substance and x the distance along the diffusion axis, Fick's law is expressed mathematically by the following expression:

$$J = -D \cdot \frac{dC}{dx} \tag{5}$$

In real world, the concentration of a substance is not constant, since it changes with respect to time and space, due to the continuous production or consumption of the substance. To express this phenomenon, Fick's law is modified as follows:

$$\frac{dC}{dt} = D \cdot \frac{d^2C}{dx^2} \tag{6}$$

Relation (6) constitutes Fick's second law. The analytical form of Fick's second law, presented in relation (7), is considered to simulate diffusion of radioactivity in the water. $C(x, t)$ is the concentration of the substance at point x and time t, C_0 is the initial concentration of the substance, C_S is the concentration of the substance at point x, while $\mathrm{erf}(x)$ is the error function.

$$C(x,t) = C_s - (C_s - C_0) \cdot \mathrm{erf}\left(\frac{x}{2\sqrt{Dt}}\right) \quad (7),$$

where $$\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x \exp(-u^2)\, du \quad (8)$$

In the simulations, the radioactivity diffusion that corresponds to 4000 timesteps around a cylinder with constant concentration equal to 100 is considered. These results to a radioactivity distribution with values in the range [0, 100], as can be seen in Figure 7.

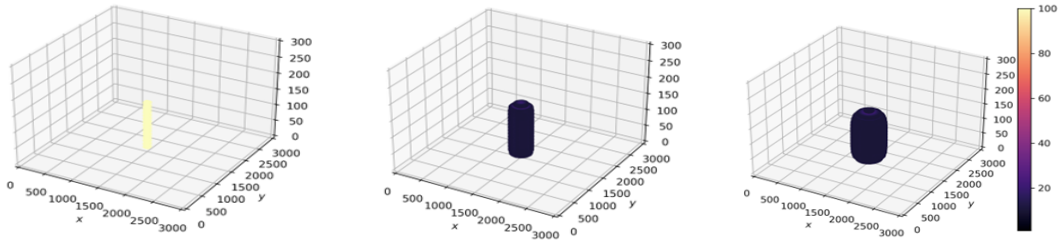

Initial state (t=0) t=4000 t=15000

Figure 7: Diffusion effect

Plume

A geometric representation of a hydrothermal plume is also considered. This is modelled as a truncated conical shape. A three-dimensional inverted cone is considered placed at the center of the bottom of the simulation environment (x = 1500, y = 1500). Inside the cone, the returned values are equal to 10, while outside it is equal to 1. Finally, the radius of the cone at its base is Rb = 80, while the radius at the highest point is Rt = 300 (Figure 8).

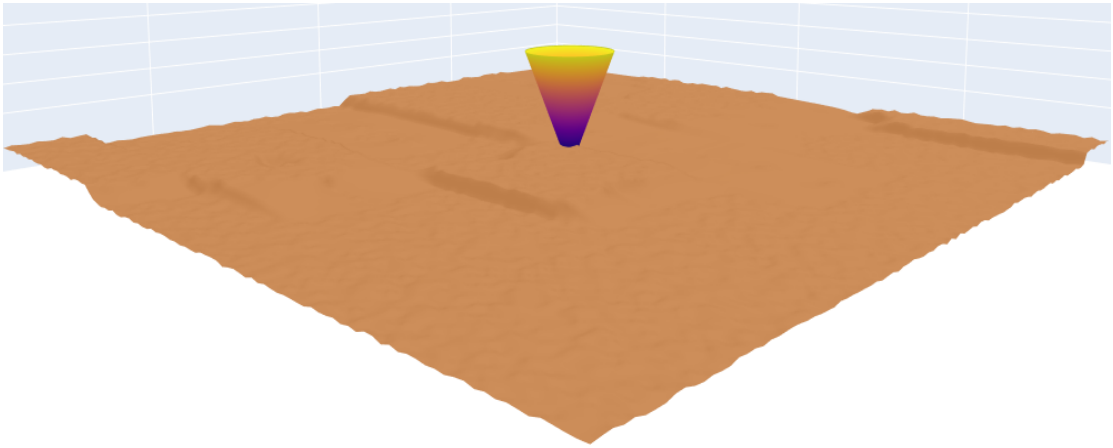

Figure 8: Graphical representation of plume in the environment of simulation

3. AI-based solutions for time series analysis

We report on the developed software for the analysis of time series data acquired mainly from the γSniffers instruments. The considered datasets are derived from the developed simulated environment as well as the real data that were acquired during the Milos field campaigns.

3.1. Processing of measurements

Measurements from γSniffers are stored in .spe files in regular intervals, containing details about the energy of the events recorded in the corresponding temporal interval, as well as meta-data regarding the operating parameters and conditions of the detector. The detected events are correlated with the position of the glider at the same time interval together with the corresponding uncertainties, based on the localization algorithms developed in the project. Radioactivity mapping considers these combined time-series (detected events and localization) in order to produce a map of the radioactivity distribution in the underwater volume being scanned. Even without considering uncertainties in the estimated location of the gliders, the problem is notoriously challenging, as measurements cover an infinitesimal portion of the scanned volume, as shown in Section 3.1. This is inherent to the problem having co-dimension 2, as measurements are virtually 1-dimesional, due also to the locality of measurements caused by water shielding, and the scanned volume being 3-dimensional. The most obvious solution to propagate measurements to the volume of interest in order to provide a radioactivity map is through interpolation of the recorded values. The simulation environment described in Section 3, was used to assess the accuracy of the maps reconstructed through interpolation, as well as their efficiency. Although the interpolation results were satisfactory, their computation is very demanding due to the very high number of measurements involved and the large extent of the volume needed to be filled through interpolation. To address these challenges, a novel MLP-based sparse measurement interpolation method was developed, described in the following section.

3.2. MLP-based sparse measurements interpolation

Model

The idea of using an MLP for interpolating values from a highly complex function builds on the work of NeRF [1], where a method for synthesizing complex scenes was developed using a continuous function defined based on a sparse set of input values. In fact, a function of interest can be approximated in arbitrary precision by an MLP, according to the universal approximation theorem. In this context one can optimize the weights of an MLP by providing

the function argument as input and using the measured (or computed) value of the function as supervision. In practice, it has been found that MLPs are not capable of learning high degree functions from low-dimensional inputs. This difficulty can be overcome by applying positional encoding to the input, which artificially increases the dimensionality of the input through the use of Fourier features [2]. In the context of radioactivity mapping, using an MLP as a function approximator and, subsequently as an interpolator through querying values at specific locations, leads to an improvement in the accuracy as well as a significant reduction of the required execution time.

As a reference architecture, an MLP with four hidden layers was considered, of which the first three consist of 256 neurons, while the third one consists of 128 neurons. The ReLU activation function was considered for the hidden layers. The MLP is optimized considering the mean-squared error with respect to the measured values in the corresponding spatial locations. A small portion of the available data is considered as validation data and an early stopping strategy is implemented, with patience equal to 7 seasons. The early stopping helps to prevent the model from overfitting, thus allowing better approximation of values far from the measured locations. Based on the discussion, the three-dimensional input coordinates are transformed using Positional encoding based on Fourier Features.

Results

To assess radioactivity mapping algorithms, the simulation environment presented in Section 3 was used, considering different types of underwater glider trajectories and different radioactivity distributions. A grid of query coordinates was considered consisting of 31x31x8 = 7688 regularly spaced points. The points are equally spaced in 31 intervals along the x and y axes in the range [0, 3000] and in 8 equally spaced intervals along the z axis in the range [0, 300]. The grid size was chosen with the rationale of having enough points to produce a reconstruction accurate and detailed enough, but also with a number of points as small as possible, as typical implementations (e.g., function ‘griddata’ of SciPy in Python) are very demanding computationally and an excessive number of query points is prohibitive due to space and time constraints. Table 4 and Table 5 present the root-mean square error (RMSE) and peak signal to noise ratio (PSNR) metrics, typical metrics used for assessing signal reconstruction quality, for the estimated radioactivity maps computed using linear interpolation for diagonal (zigzag) and spiral glider trajectories, respectively. The name of the experimental setup reports the corresponding trajectory spacing and the type of radioactivity distribution considered.

Table 4: Radioactivity mapping using diagonal (zigzag) trajectories

Setup	RMSE	MSE	PSNR	Time (sec)
100x100_cone	0.435	0.19	26.309	4517.82
150x150_cone	0.448	0.201	26.058	3685.19
200x200_cone	0.455	0.207	25.922	6208.8
250x250_cone	0.485	0.235	25.373	2474.45
300x300_cone	0.465	0.217	25.729	2489.56
100x100_diffused	1.254	1.573	37.946	4714.56
150x150_diffused	1.282	1.643	37.757	3189.47
200x200_diffused	1.995	3.979	33.915	6514.14
250x250_diffused	1.467	2.153	36.582	2365.43
300x300_diffused	1.756	3.085	35.02	2126.22
100x100_distrib	0.392	0.153	14.164	4781.69
150x150_distrib	0.455	0.207	12.867	3081.06
200x200_distrib	0.526	0.276	11.609	6601.45
250x250_distrib	0.526	0.276	11.607	2692.13
300x300_distrib	0.562	0.315	11.031	2432.46

Table 5: Radioactivity mapping using spiral trajectories

Setup	RMSE	MSE	PSNR	Time (sec)
200x200_cone	0.364	0.132	27.863	8903.57
300x300_cone	0.451	0.203	26.003	2063.71
400x400_cone	0.556	0.309	24.191	1566.98
500x500_cone	0.584	0.341	23.757	901.47
600x600_cone	0.617	0.381	23.275	1350.45
200x200_diffused	0.966	0.933	40.213	9747.65
300x300_diffused	1.928	3.718	34.209	1775.19
400x400_diffused	1.83	3.351	34.661	1691.73

500x500_diffused	2.9	8.413	30.663	863.36
600x600_diffused	2.732	7.465	31.182	1166.25
200x200_distrib	0.2	0.04	19.996	9425.7
300x300_distrib	0.273	0.074	17.31	2059.14
400x400_distrib	0.353	0.124	15.073	1696.59
500x500_distrib	0.366	0.134	14.746	893.72
600x600_distrib	0.392	0.153	14.162	1308.94

We note that the radioactivity distribution based on diffusion is better approximated. This can be attributed to the fact that changes in the radioactivity values are quite small. Plume (conical) distribution gives the second-best results, while the analytical distribution of relation (4) is the most challenging, most probably due to the large extent of spatial frequencies that the signal contains. In terms of the glider trajectories used for scanning the volume, the best results are obtained with the spiral motion and voxel spacing equal to 200 x 200, at the cost of excessively long time to complete the mapping. Diagonal movement, on the other hand, provides slightly worse results but with the important advantage of reduced scan time.

The last column of Table 4 and Table 5 reports the time required to execute the linear interpolation on the 31x31x8 grid on a typical workstation. This time ranges from 15 minutes to approximately 3 hours. The time required to complete the interpolation task is directly related to the trajectory and the number of query locations. These extended computation times are also one of the reasons we have developed the MLP-based interpolation method.

Figure 9 and Figure 10, show the performance comparison of the MLP-based interpolation with respect to the linear interpolation baseline using the PSNR metric, for the diagonal and spiral movements, respectively. One can see that MLP gives comparable performance with the linear based interpolation in the case of diagonal trajectories, while they lead to significantly better performance when spiral trajectories are used. As expected, the smallest voxel spacing, i.e., 200m x 200m, leads to the best results, from the cases considered.

Importantly, the execution time of MLP-based interpolation is sped-up multiple times with respect to linear-based interpolation. In particular, the time required for MLP-based interpolation for the scenarios we considered range from 18 to 369 seconds with an average of 101 seconds. Compared to the corresponding linear interpolation times, this corresponds to a speedup in the worst case of 8 and in the best case of 26 times.

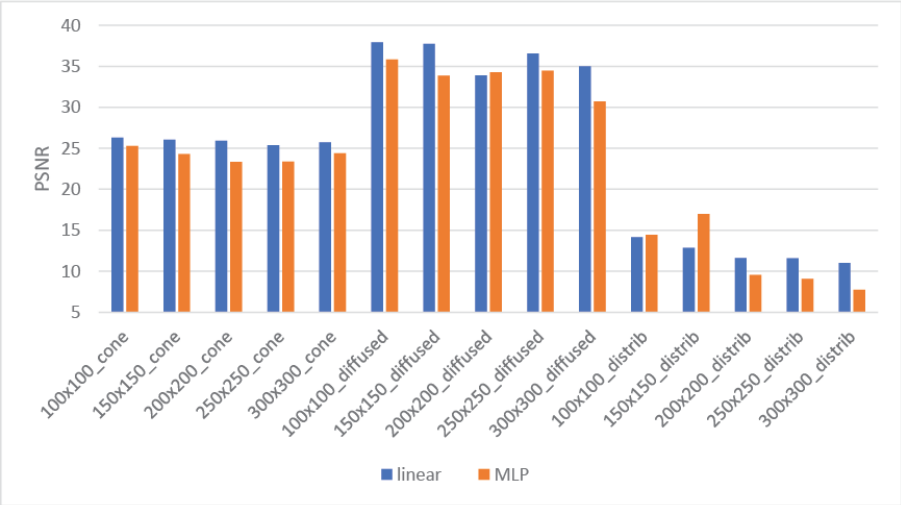

Figure 9: Comparison of linear and MLP-based interpolation results for diagonal movement

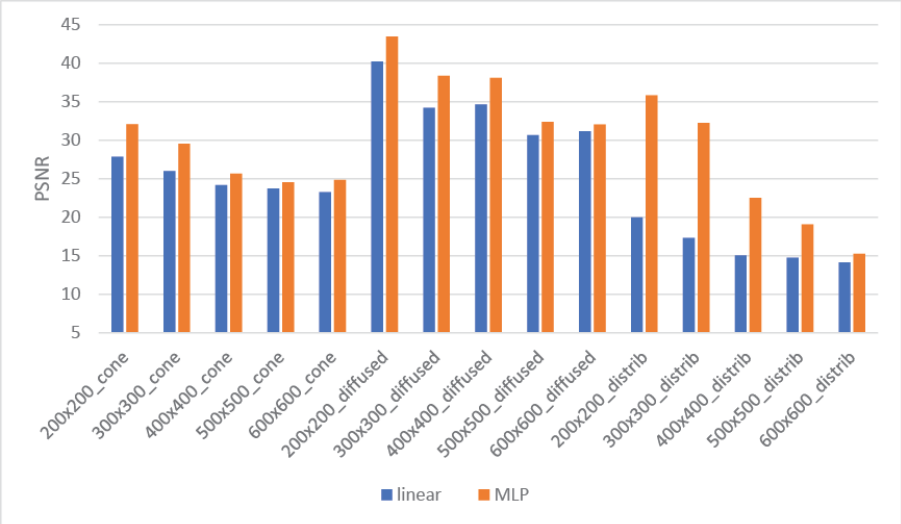

Figure 10: Comparison of linear and MLP-based interpolation results for spiral movement

3.3. Field deployment at Milos Island

From the 25th to the 27th of March 2023 field experiments have been performed at the island of Milos, Greece. Radioactivity measurements with the γSniffer and SUGI have been collected

from hydrothermal fields at the Palaiochori and Alykes locations (Figure 11). The two instruments were mounted on an ROV. During the experiments the ROV has visited multiple gas emitting sources (vents), collecting radioactivity measurements at different depths and distances from the vents, in order to estimate radioactivity levels in the water column and on the sea bottom. Multiple radioactivity hot spots have been identified around the vents, as can be seen by the bubbles and the seabed discolorations in Figure 12.

Figure 11: Map of Milos Island (Greece) where the field deployment sites are indicated

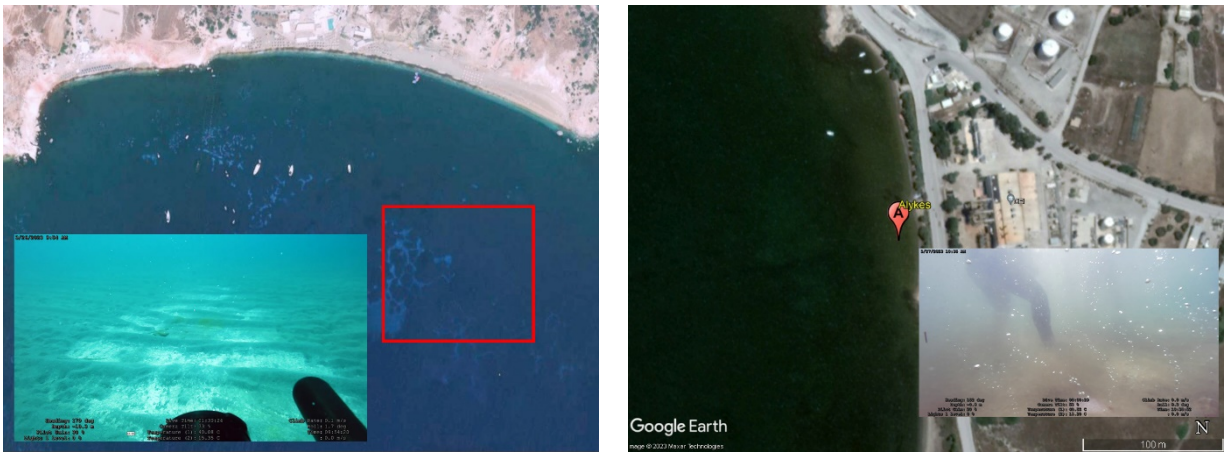

Figure 12: Locations where field deployments took place at Palaiochori (left) and Alykes (right) beaches. Images from the ROV are shown in the inserts.

Data from the γSniffers were collected with an acquisition time of 10 seconds (live time), resulting in a SPE file containing the spectrum corresponding to the detection events registered in this interval, together with meta-data regarding the setup of the detector and all relevant parameters. Preprocessing algorithms have developed in order to extract the data regarding detection events from the SPE files in relation to the UTC time of the acquisition and correlate them with measurement logs filled during the deployments as well as with the depth and other data regarding the status of the ROV. Figure 13 presents processed radioactivity measurements in terms of counts-per-second (CPS) vs UTC time considering different low-level discriminator (LLD) values for one of the deployments that took place on the second day of the field tests (26th of March), where four different vents were visited as can be seen by the elevated CPS values around 8:50, 9:25, 10:00 and 10:10. The source corresponding to 10:00 is shown in the insert of Figure 12 (left).

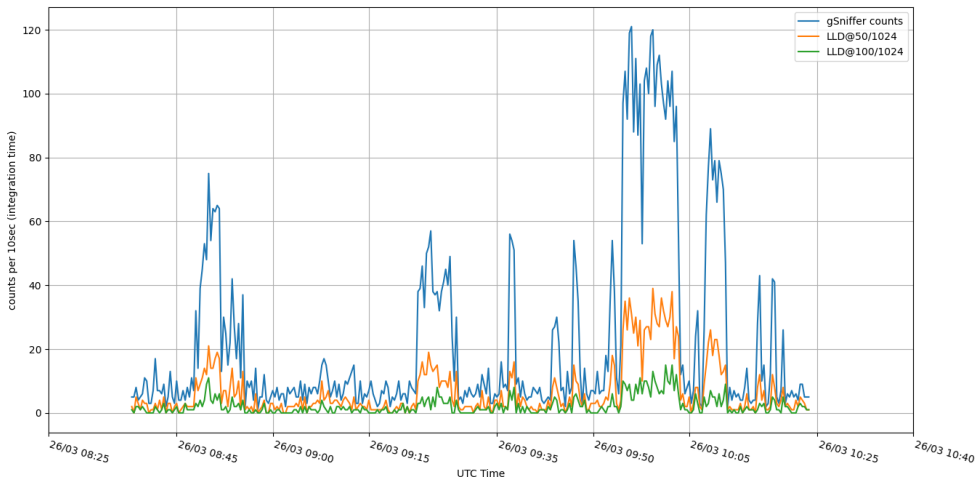

Figure 13: γSniffer measurements from a deployment at Palaiochori beach for different LLD values

Similarly, Figure 14 shows the processed γSniffer data collected during the field deployment at Alykes beach on the 27th of March. In this graph annotations from the field log are showing, indicating different placements of the sensor around the vent being monitored.

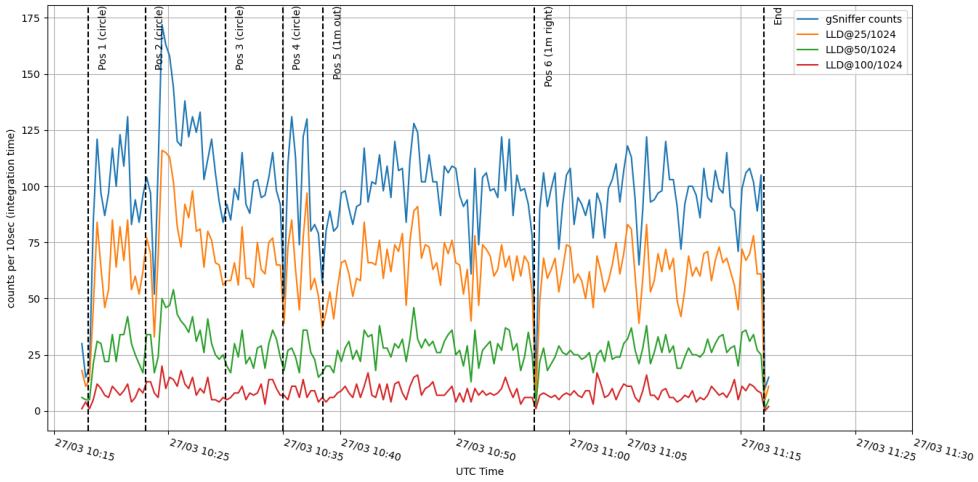

Figure 14: Annotated ySniffer measurements from the Alykes beach for different LLD values

4. AI-based solutions for imaging

We report on the developed software for the analysis of imaging data acquired mainly from the SUGI instrument. The considered datasets are derived from testing exercises in the laboratory as well as during the Milos field campaign. Moreover, state of the art AI methods have been designed and developed able to exploit mixup and powerful attention mechanisms in imaging data with deep vision transformers.

4.1. Metric learning and cutting-edge attention mechanisms

Initially, based on the estimation that the number/size of RAMONES imaging datasets will not be huge in size and adequate to train from scratch deep neural network architectures with numerous parameters, we considered and studied state of the art data augmentation techniques, In particular, based on the argument that metric learning is binary classification of pairs of examples into “positive” and “negative”, we have introduced a direct extension of mixup from classification to metric learning. The developed methodology [3] is generic, applying to a large class of loss functions that separate positives from negatives per anchor and involve component functions that are additive over examples. Those are exactly loss functions that require less mining. We have introduced a novel way of interpolating labels, such that the interpolation factor affects the relative weighting of positives and negatives. Other than that, our approach is completely agnostic with respect to the mixup method, opening the way to using more advanced mixup methods for metric learning. Based on an extensive qualitative and quantitative evaluation we have demonstrated that we consistently outperform baselines using a number of loss functions on a number of benchmarks and we improve the state of the art using a single loss function on all benchmarks, while previous state of the art was not consistent in this respect. Moreover, due to the fact that metric learning is about generalizing to unseen classes and distributions, our work may have applications to other such problems, including transfer learning, few-shot learning and continual learning.

Furthermore, we have studied state of the art deep learning frameworks based on vision transformers (ViT). Considering that transformers are data-hungry, many studies advocate pre-training them on unsupervised pretext tasks, determined only by raw data. A prominent paradigm is to mask a portion of the input tokens—words in text or patches in images—and train the transformer to predict these missing tokens. Towards this direction masked image modeling (MIM) based on self-supervised methods have shown impressive results, however, an important aspect that has not been well explored so far is how to choose which image tokens to mask. Thus, we have designed and developed a method [4] that is able to exploit the

intrinsic properties of ViT and in particular, its self-attention mechanism. Given an input sequence of image patches, we forward it through the transformer encoder, thereby obtaining an attention map in its output. We then mask the most attended tokens. This strategy, which we call attention-guided masking (AttMask), is an excellent fit to popular distillation-based self-supervised objectives, because it is the teacher encoder that sees the entire image and extracts the attention map, and the student encoder that sees the masked image and solves the reconstruction task. AttMask thus incurs zero additional cost. In a nutshell: we have introduced a novel masking strategy for self-supervised learning (AttMask), that exploits the intrinsic properties of ViT by leveraging its self-attention maps to guide token masking. We have demonstrated how to efficiently incorporate this above masking strategy into teacher-student frameworks that use a MIM reconstruction objective and demonstrate significant performance improvements over random masking. Moreover, we performed extensive experimental evaluation, and validated that AttMask offers several benefits: it accelerates the learning process; it improves performance on a data-limited regime and on a variety of downstream tasks; it increases the robustness against background changes, thus revealing that it reduces background dependency.

4.2. Preliminary results on anode activations of ASIC imaging modules

Significant efforts have been dedicated in building a custom API in Python for communicating and collecting data from the GDS-100, which is the detector module used for SUGI. Data logging, processing, and communication with the GDS-100 is achieved using the TCP/IP protocol in a custom format developed by the detector manufacturer (Doppio format) with TCP packets being used for system configuration and UDP packets for data readout. The developed API allows the real-time configuration of the detector, monitoring of the system health (e.g., key component voltage, temperatures, etc.), as well as real-time readout of data corresponding to individual detection events. Preprocessing software modules have been developed to achieve basic tasks as reporting of total event counts and counts per second, the assignment of recorded events to ADC channels and the construction of the corresponding histogram, as well as the specific anodes (i.e., pixels) where the event(s) were registered. Figure 15 shows the raw data registered for a single detection event, while Figure 16 shows cumulative pixel activations due to background radioactivity registered in the lab using all four ASIC modules of SUGI corresponding to 1000 registered events. Based on these data, the generation of spherical heatmaps based on Compton imaging will be pursued next.

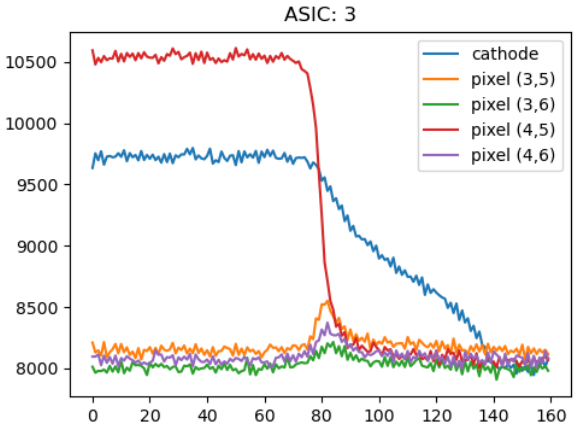

Figure 15: Data from an event registered by GDS-100

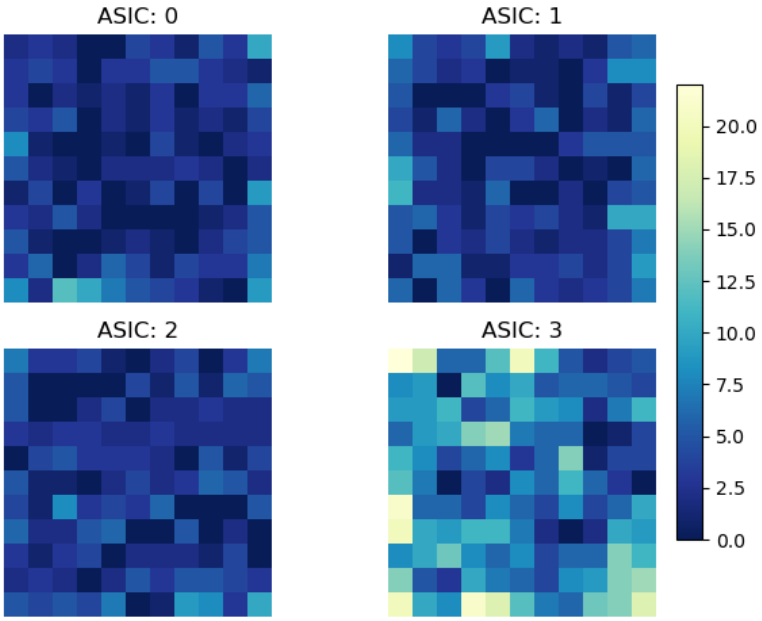

Figure 16: Anode activations of the four ASIC modules of GDS-100

5. Closing Remarks and Next Steps

This deliverable provides an overview of the 1st release of AI-based solutions for processing and analysing data collected from the RAMONES instruments developed up to this point. To anticipate challenges of detector deployment using underwater gliders in large scan volumes a high-level simulation environment has been developed, to study the characteristics of the radioactivity measurements (time-series) resulting from different combinations of glider trajectories and radioactivity sources. The simulation results indicate that detection of point sources lying on the sea bottom is virtually impossible in normal operating conditions of underwater gliders, due to radioactivity shielding from the water. On the other hand, the results are encouraging for the effective mapping of large areas/volumes. The simulations show that mapping accuracy depends on the particular radioactivity distribution, the localization accuracy, the type of scan trajectories and their spacing, as well as the method used to propagate the measurements to the entire scan volume. Regarding the latter, a novel MLP-based interpolation method has been developed that achieves higher accuracy in most cases with notable speedup with respect to typical linear interpolation implementations. An important outcome of the simulations is that they allow to better understand the trade-off between mapping accuracy and time required to execute the selected scan paths.

Regarding processing of time-series, besides the MLP-based interpolation method for highly sparse measurements, algorithms have been developed for preprocessing and analysing the data collected by the γSniffers and the SUGI instruments. Data from the field deployments at Milos Island (Greece) were analyzed using these tools. In the next version, attention will be given on the aspect of localization uncertainty and how this can be accommodated in the developed mapping and time-series processing algorithms.

Finally, cutting-edge deep learning frameworks have been developed along with suitable tools were developed to properly process the imaging data collected by SUGI, to enable the application of Compton Imaging methods at the next stage, in order to compute spherical heatmaps indicating directionality of incoming gamma rays.

References

[1] Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., & Ng, R. (2021). Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, *65*(1), 99-106.

[2] Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J. & Ng, R. (2020). Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, *33*, 7537-7547.

[3] Venkataramanan S., Psomas B., Avrithis Y., Kijak E., Amsaleg L., Karantzalos K., 2022. It Takes Two to Tango: Mixup for Deep Metric Learning, International Conference on Learning Representations (ICLR).

[4] Kakogeorgiou I., Gidaris S., Psomas B., Avrithis Y., Bursuc A., Karantzalos K., Komodakis N.2022. What to Hide from Your Students: Attention-Guided Masked Image Modeling, European Conference on Computer Vision (ECCV).