UK Acoustics Network

3rd Underwater Acoustics PhD Symposium Day

13 December 2023, National Oceanography Centre

Abstracts

From 0830			**Arrival, tea and coffee available**
			Session 1 Jonathan Bull
0900	Jonathan Bull	NOC	**Opening**
0905	Sam Dugelay	Thales	**INVITED TALK:** A life in underwater acoustics – in particular a passion for synthetic aperture sonar and all its ramifications
0930	Ellen White	Univ of Southampton	Towards the automation of detecting marine sound sources from passive acoustic monitoring data
0950	Yiyi Whitcelo	Univ of Liverpool	Scattering of acoustic waves from rough seabeds: a comparison of two-and three-dimensional models
1010	Andrea Napoli	Univ of Southampton	Bioacoustic event detection under domain shift
1030			**Break (30 min)**
			Session 2 Yiyi Whitcelo
1100	Gemma Bekki	Univ of Southampton	Towards the localisation of cetaceans using passive acoustic data recorded from gliders
1120	Tom Corner	Newcastle Univ	Sea trial results for an adaptive rake receiver with the Phorcys V0 acoustic communication waveform
1140	Finley Boulton	Univ of Liverpool	Ray and beam models for underwater sound propagation
1200	Gary Benson	Univ of York	Localisation of ad-hoc AUV networks for underwater acoustic ship hull inspection
1220	Ben Ford	NPL	**RESEARCH SOUNDINGS COMPETITION – Introduction**
1235			**Lunch / Networking / Voting (1 h 30 min)**
			Session 3 Ellen White
1405	Ben Ford	NPL	**RESEARCH SOUNDINGS COMPETITION – Winner**
1415	Ben Thomas	Atlas Elektronik UK	**INVITED TALK:** Branching out from underwater acoustics: From synthetic aperture sonar to electromagnetic signature control
1440	Oscar Bryan	Univ of Bath	A diffusion-based super resolution model for enhancing sonar images (lightning talk)
1445	Joshua Wakefield	Univ of Liverpool	Maximising detection performance using high performance processing of multi-sensor data (lightning talk)
1450	Amy Feakes	Univ of Southampton	Estimating the abundance and distribution of marine mammals using passive acoustic monitoring and biotelemetry (lightning talk)
1455	Luke Prentice	Univ of Strathclyde	Bio-inspired ultra-compact acoustic transmitter – Bubbles as a means of weak-source amplification (lightning talk)
1500	Florian Mahieu	Univ of York	Acoustic underwater networking concepts in flooded caves and other constrained underwater environments (lightning talk)
1505	Hanif Sutiyoso	Univ of Southampton	Novel laboratory measurement of acoustic velocity and attenuation at 1-20 kHz frequency for ice bearing sediment (lightning talk)
1510	Ben Dell	Univ of Southampton	Exploring the use of acoustic indices for habitat level monitoring and event detection
1530			**Break (30 min)**
			Session 4 Andrea Napoli
1600	Edward Clark	Univ of Bath	Optimising sensor path planning with reinforcement learning and passive sonar modelling
1620	Abdullah Olcay	Univ of Southampton	Deep learning based underwater sound classification using stacked cepstral features
1640	Francesca Pereira	Dstl / Univ of Southampton	Spice driven sound speed fluctuations in Fram Strait during Summer 2022
1700	Denise Risch	Scottish Association for Marine Science	**INVITED TALK:** Passive acoustics for long-term monitoring of protected species and habitats: current results and future directions
1725	Duncan Williams	Dstl	**Closing**
1730			**Reception**
From 1900			**No host dinner TBC**

UK Acoustics Network
3rd Underwater Acoustics PhD Symposium Day
13 December 2023, National Oceanography Centre

Sam Dugelay

INVITED TALK: A life in underwater acoustics – in particular a passion for synthetic aperture sonar and all its ramifications

This talk will present Dr Dugelay's discovery of underwater acoustics during her PhD in mathematics at the Institute of Research on the Sea (Ifremer) in France and the profound impact on her career. Following on her PhD, synthetic aperture sonar (SAS) has been a constant theme throughout her research which developed from the sensor design to autonomous mine countermeasures. The fundamental concept of SAS will be given but the focus of the talk will be centred around the multiple opportunities that a high resolution imaging sonar can provide.

Ellen White

Towards the automation of detecting marine sound sources from passive acoustic monitoring data

Passive Acoustic Monitoring (PAM) data holds the potential to determine spatial and temporal variations in ecosystem health and species presence if automated detection and classification algorithms are available, which are capable of discrimination between marine species and the presence of anthropogenic and environmental noise. This talk will outline my PhD thesis in which we present an open-source state-of-the-art 'small-scale' convolutional neural network (CNN), trained to detect odontocete tonal and broadband call types and vessel noise from broadband acoustic data between 0 – 48 kHz. Performance is evaluated in a diverse array of ambient noise conditions to measure how the network will perform on long-term datasets. The model presents proof-of-concept that a computationally low-cost deep learning model can be used to mine extensive marine acoustic datasets for more than a single signal type, providing information on temporal scales relevant to the management of marine protected areas and the conservation of vulnerable species.

For the CNN to be of use to the wider bioacoustic community, it must be robust within different marine environments where the ambient noise levels vary from the original training set. We explore the adaptability of the CNN to new acoustic fields through 'fine-tuning'. Training sets of differing sizes from each sound source class are curated from data collected in the Gulf of Mexico to determine at what point the network adapts to the ambient environment. The use of fine-tuning improves accuracy by 30% with only 50 spectrograms per class input to the model. When the training set is increased to 500 spectrograms a further 5% improvement in accuracy is achieved. We therefore establish that CNNs can be adapted to new marine soundscapes with only a small-amount of site-specific data, allowing researchers to harness the power of pre-trained bioacoustic CNNs to their own datasets. Before CNNs can become commonplace in the bioacoustic toolbox they must demonstrate performance benefits above that of currently used algorithms.

The final element of this work empirically evaluates the performance of the CNN as a dolphin monitoring tool in comparison to that of a widely popular data logger, the C-POD. A total of 2136 hours of acoustic data were collected at two locations in west Scotland, and analysed by the CNN, the C-POD and also manually to curate a ground-truthed set of detection positive hours (dph). The CNN presents as a useful tool for monitoring dolphin presence in diverse marine soundscapes, achieving an overall accuracy of 0.84 and a very low false-positive rate of 0.05, outperforming the C-POD for all analysis periods which experienced very high rates of missed detections. The CNN accurately describes seasonal and diel patterns in dolphin presence around each of the moorings, even when sensitivity is hindered by adverse weather periods. We present to the research community an open-source CNN for broadband multi-signal detection for monitoring marine species in diverse acoustic environments, which can provide information to management and stakeholders in an appropriate timeframe for conserving vulnerable species and habitats.

UK Acoustics Network
3rd Underwater Acoustics PhD Symposium Day
13 December 2023, National Oceanography Centre

Yiyi Whitcelo

Scattering of acoustic waves from rough seabeds: a comparison of two- and three-dimensional models

The seabed possesses variations in its material properties, layering structure, slope and roughness making prediction and analysis of sound propagation in shallow water highly complex. For certain frequency ranges, the effect of roughness becomes significant. Numerous models have been adopted to account for the scattering of acoustic waves by rough surfaces, including Kirchhoff approximation (for relatively large scales of roughness), small perturbation theory (for relatively small scales of roughness), and the small slope approximation.

In this study, we have used the Kirchhoff approximation to model the scattering of acoustic waves, and the effects of the seabed roughness and slope angle of the seabed on the scattering characteristics are investigated. The results are compared with a numerical graphics processing unit (GPU)-accelerated finite element model incorporating roughness parameters (Root Mean Square height and correlation length) for rough, sloped seabed geometries (with gradients ranging from zero to twenty degrees) in shallow water environments. The effect of large- and small-scale roughness on sound propagation is investigated in two- and three-dimensional space.

UK Acoustics Network
3rd Underwater Acoustics PhD Symposium Day
13 December 2023, National Oceanography Centre

Andrea Napoli

Bioacoustic event detection under domain shift

Machine learning methods, and deep networks in particular, often underperform on data which lies outside the training distribution. Changes to the input data distributions (known as domain shift) are particularly prevalent in bioacoustics, where many external factors can vary between datasets, although the effects of this are often not properly considered. We construct a benchmark for out of distribution (OOD) performance based on the detection of humpback whales in underwater acoustic data. Several humpback whale detectors from the literature are implemented as baselines, along with our own detector based on a convolutional neural network (CNN). Then, a set of domain generalisation (DG) and unsupervised domain adaptation (UDA) algorithms are compared.

Results show that UDA can significantly improve OOD performance when few distinct sources of training data are available. However, this is not a substitute for better data, as negative transfer (where the adapted models actually perform worse) is commonly observed. On the other hand, we find that training on a variety of distinct sources of data (at least 6) is sufficient to allow models to generalise OOD, without the need for advanced UDA or DG algorithms. This allows our model to outperform all the baseline detectors we test, despite having 10,000 times fewer parameters and 100,000 times less training data than the next-best model.

Gemma Bekki

Towards the localisation of cetaceans using passive acoustic data recorded from gliders

Sperm whales (*Physeter macrocephalus*) can be found in every ocean, mainly along continental shelf edges and canyon systems. Visual surveys have provided vital information about this species' geographical distribution for many years, while individual whales have been physically tagged to better understand their movement within the water column. Another method to track their activity is using the acoustic signals they produce. So far, studies have used acoustic data recorded on static or towed hydrophone arrays.

This project aims to locate Sperm whales using data from underwater gliders which offer the potential for greater understanding of their geographical and depth ranges. Four SeaExplorer gliders, developed by Alseamar, were deployed in the PELAGOS protected sanctuary in the Mediterranean. Information about a whale's position can be determined by combining the bearing given by several techniques such as beamforming, Time Difference Of Arrival (TDOA) and the sea surface reflection of the clicks. By collating this information from each glider, a better estimate of the cetacean's location can be calculated.

Tom Corner

Sea trial results for an adaptive rake receiver with Phorcys V0 acoustic communication waveform

Over the last few decades signal processing methods traditionally applied to Radio Frequency (RF) communication have been attempted and validated in the acoustic channel. Two such techniques are adaptive equalisation and rake receiving. Adaptive equalisation combats Intersymbol Interference (ISI) and short arrival multipath by dynamically updating a filter and applying this to the received signal. In channels with multiple dispersed path arrivals it is possible to use raking to exploit the temporal profile of the channel, the receiver can combine additional rake ”fingers” with the direct path to increase Signal to Interference Noise Ratio (SINR) at the demodulator’s input.

Here these techniques are combined into an Adaptive Rake Receiver (ARR) and this structure is capable of removing ISI, combining multipath arrivals and dynamically modifying its filter taps throughout a communications packet. The structure also utilises closed loop Doppler tracking to track the relative velocity between communicating parties. This structure has been tested on data collected from the North Sea in January 2021 over ranges from 160 m to 20 km. In this trial the Phorcys V0 waveform was used as the acoustic communication scheme. This waveform is a Direct Sequence Spread Spectrum (DSSS), M-ary orthogonally modulated scheme with multiple packet lengths (1-byte to 256-bytes), spreading ratios (16 to 128) and frequency bands (8 – 10 kHz and 20 – 28 kHz) defined for versatile military and civilian user operation.

The results show that in severely time spread and highly dispersive channels the combination of the rake receiver with the adaptive filtering provides up to 25% increase in packet delivery rate over a structure with a single adaptive filter on the direct path. The most improvement occurs at longer ranges where the Signal to Noise Ratio (SNR) is lower. The raking is performed on a channel estimate provided by the normalised cross-correlation (NXC) already performed during the detection of the packet and no additional adaptive filter taps are used over the single path variant, making the combination of the two methods reasonably efficient.

Supported by the NATO Science for Peace and Security programme under the project Cybersecurity for Safe Underwater Acoustic Communications and Sonardyne International Ltd.

Finley Boulton

Ray and beam models for underwater sound propagation

Propagation of acoustic waves in the ocean is complex due to environmental factors causing spatial and temporal sound speed variations. Understanding how sound propagates in the ocean is critical for effective communication, source localisation, and navigation. To this end, we develop ray and beam methods to understand how acoustic waves propagate in the underwater environment. We use finite difference methods to solve the ray equations with spatially varying sound speed in three dimensions. Further, we solve the dynamic ray equations and construct a beam around each ray to obtain transmission loss within an oceanic waveguide.

Gary Benson

Localisation of ad-hoc AUV networks for underwater acoustic ship hull inspection

Ship hull inspections are currently conducted using divers, dry docks or remotely operated vehicles. My PhD seeks to develop techniques to allow a network of AUV nodes to conduct underwater acoustic ship hull inspection quickly. My study focuses on developing and managing AUV organisation and localisation techniques where the ship hull is used as a position reference. The study focuses on using new techniques such as underwater acoustic one way travel time. The benefits of my techniques will be speed of operation and simplicity of equipment required plus removal of skilled operators or specialist facilities.

Ben Thomas

INVITED TALK: Branching out from underwater acoustics: From synthetic aperture sonar to electromagnetic signature control

In this talk, Ben will cover his early career experience, moving from his PhD and postdoctoral research in synthetic aperture sonar (SAS) micronavigation, to his recent transition to industry research at Atlas Elektronik UK. During his PhD and postdoctoral studies, Ben developed an interest in simulation, algorithm development, optimisation and unbiased estimation. Although SAS micronavigation is a seemingly niche topic, Ben will demonstrate how the skills learnt in his academic career have been applicable to his industry research, including estimation and control of the electromagnetic signature of naval platforms.

Oscar Bryan

A diffusion-based super resolution model for enhancing sonar images (lightning talk)

Improved hardware and processing techniques such as synthetic aperture sonar has led to imaging sonar with centimetre resolution. However, practical limitations and old systems limit the resolution in new and legacy datasets. We propose using single image super-resolution based on a conditioned diffusion model to map between images at different resolutions. We demonstrate improved performance compared to interpolation, convolutional neural networks, and generative adversarial networks. We also propose two new sonar specific evaluation metrics based on acoustic physics and utility to automatic target recognition.

Joshua Wakefield

Maximising detection performance using high performance processing of multi-sensor data (lightning talk)

This work utilises ray tracing techniques to simulate sound propagation in complex underwater environments, enabling the accurate prediction of acoustic signal paths. Conventional array beamforming techniques are then employed to process the received signals and extract directional information for source localisation. A particle filter tracking algorithm is implemented to estimate the trajectory of underwater objects based on bearing-only measurements.

Ray-based methods solve the equations governing propagation of sound waves by tracing the paths of individual sound rays, considering interactions with the seafloor and sea-surface. The speed and direction of the rays are determined based on the local properties of the medium, such as temperature, salinity, and density. These interactions are modelled using principles from geometrical and physical optics, incorporating mathematical expressions for the laws of reflection, refraction, and diffraction. This is achieved using 3D vectorised Snell's Law where the speed of sound at either side of the boundary is calculated by an appropriate equation for deep or shallow water, e.g., the Munk or Mackenzie equations.

Conventional beamforming is a signal processing technique employed to enhance desired acoustic signals from a specific direction while minimising noise and interference originating from other directions. By employing time delays and weightings of signals received by an array of hydrophones, conventional beamforming creates a spatial filter that focuses on a particular direction, effectively improving the signal-to-noise ratio. The power spectral density (PSD) for a specific steering direction can be derived by computing the squared magnitude of the complex signal's Fourier transform. The steering direction at which the PSD is maximum serves as the estimated direction of arrival and the measurement to be used in the update procedure of the particle filter.

A particle filter is a sequential Monte Carlo method used for state estimation in dynamic systems. By employing a set of weighted particles to represent possible source positions, the filter can effectively track the source's movement over time. At each timestep the particles are propagated forward in time using the system's dynamic model to predict their positions. When a measurement becomes available, the weights of the particles are updated based on the likelihood of the measurements given the predicted positions. To prevent the degeneracy of the particle set, resampling is often performed.

Amy Feakes

Estimating the abundance and distribution of marine mammals using passive acoustic monitoring and biotelemetry (lightning talk)

Commercial whaling prior to the IWC moratorium reduced global whale populations, with fin whales among the most heavily exploited species. In recent years, the recovery of southern hemisphere fin whales has been observed with the return to ancestral feeding grounds - Elephant Island; a krill-rich area, allowing efficient feeding.

There is still limited understanding of this population. We aim to carry out research to gain a further understanding of this population's abundance and distribution, specifically at the feeding grounds through integrating data types. We will use both passive acoustic monitoring (PAM) and biologging, common methods of monitoring marine mammals. These two approaches differ in the insights that they provide on marine mammals, Eulerian versus Lagrangian, respectively.

The experimental design involves deploying passive acoustic recorders on the shelf around Elephant Island at 150 m depth. Then, applying machine learning techniques to discriminate the vocalisations of fin whales. Concurrently, fin whales will be tagged to gain an understanding of their movements during the feeding season. This tracking data will be used to develop density surface models. Then, investigation will begin to develop methods of fusing this data to aim to provide improved abundance and distribution estimates of the Southern Hemisphere fin whale population.

Luke Prentice

Bio-inspired ultra-compact acoustic transmitter – Bubbles as a means of weak-source amplification (lightning talk)

An investigation into the relevant literature in the field of bubble acoustics has been carried out in service of better understanding the mechanisms behind the incredibly loud mating call of *Micronecta scholtzi*. It is found that while the field of bubble acoustics is deep and wide there is a gap in the literature when it comes to the understanding of larger-scale bubbles as harmonic oscillators, the shell effects in cm scale bubbles and weak-source bubble coupling. Due to these gaps many possible experimental procedures have been developed to define elastically enclosed bubble behaviour, bubble pressure output when stimulated within a superhydrophobic cage, frequency and amplitude output of a volumetrically dynamic bubble and source/bubble coupling in different orientations, with the end goal of creating a marine transducer device of a much smaller size and mass of any existing low-frequency transmitter.

There is a growing need for underwater acoustic devices that are small in size compared to the wavelength that they produce. Current methods for high-power transmissions at lower frequencies rely on the use of large displacements or a large radiating area which is not feasible for a truly compact device. Work has already been done at the Strathclyde Department of Electronic and Electrical Engineering in the Centre for Ultrasonic Engineering studying the lesser water-boatman male insect (*Micronecta scholtzi*) which communicates through efficient generation of underwater sound despite being much smaller than the wavelength of its output. This is achieved by the animal using its air-supply as a secondary amplifier. As with other Corixinae the initial sound is performed by stridulation although in this case the striated organ and plectrum are located on the right paramere of the genital appendage, sending vibrations throughout the body of the insect. These vibrations in the body of the insect create a reactive acoustic near field in the surrounding water which, it is thought due to the tiny size of this insect compared to other of its species, the stridulitrum is nearer to the carried bubbles which allows for greater pressure transfer into the bubble, stimulates the air supply to oscillate at its Minnaert frequency. This frequency then represents the de facto "tone" of the insect's mating call at a much larger call volume than the insect would be able to produce through its own physiology.

As shown in "Extreme call amplitude from near-field acoustic wave coupling in the stridulating water insect *Micronecta scholtzi* (Micronectinae)," the paper that set the groundwork for this research, insects in nature have managed to use the phenomenon of Minnaert frequency to greatly amplify their call amplitude. However this phenomenon has not been observed elsewhere which raises multiple questions which this research seeks to answer – can this phenomenon be scaled up both in terms of amplitude and geometry? Practical applications of this mechanism would include sea-floor scanning sonar which requires much lower frequency output and therefore a much larger bubble. Another practical consideration is to explore the possibility of a multi-spectral output rather than the monochromatic output observed in *Micronecta Scholtz*. Another area of interest that has been identified currently is the comparison of 'natural' direct air-water bubble output versus bubbles with a manufactured 'skin', for which losses will be incurred through the elastic shell of any bubble housings that may be considered and these shells will affect source-bubble coupling and bubble-bubble coupling. Lastly, source positioning is a key aspect of this research and will be explored in depth in tandem with the above problems.

If a bubble transducer can be designed this will be a breakthrough for underwater acoustics – presently when a low frequency sonar pulse is required these devices are almost prohibitively thick, heavy and expensive to manufacture as their centre frequency is determined by the speed of sound within the material from which they are composed and the thickness of the resonator. A bubble device would be comparatively miniscule and light in mass allowing for the attachment of this device to miniaturised unmanned submersibles.

Florian Mahieu

Acoustic underwater networking concepts in flooded caves and other constrained underwater environments (lightning talk)

As a yet poorly known and barely explored space, underwater cave systems, and other confined underwater environments pose a steep challenge for review through automated means. But networking can help alleviate these issues by letting devices share in real-time information about their state and their immediate surroundings. However, underwater acoustic networking in these conditions brings its own set of complications. This talk presents these different challenges as well as the research direction taken to overcome them.

Hanif Sutiyoso

Novel laboratory measurement of acoustic velocity and attenuation at 1-20 kHz frequency for ice breaking sediment (lightning talk)

Natural gas hydrate reservoirs have emerged as a potent energy source or carbon-dioxide storage alternative. In addition, due to global warming, the permafrost environment is at risk. Therefore, understanding gas hydrate and permafrost is urgent. We present ongoing experimental measurements to study the acoustic properties of ice-bearing sediment as an alternative for ice-bearing and gas hydrate-bearing sediment measurement. We used ice as an analogue for measurement because ice and gas hydrate are geomechanically similar, resulting in similar acoustic wave speeds. Using a custom-built acoustic pulse tube, we made a novel measurement of acoustic velocity and attenuation in the frequency range of 1 – 20 kHz. This pulse tube is made of stainless steel and has a length of 4.5 m and inside and outside diameters of 0.07 m and 0.252 m, respectively.

The samples are frozen water-saturated sand pack and synthetic sandstone in a cylindrical PVC jacket. We measured the sample’s acoustic properties, P-wave complex velocity and attenuation 1/QP, from frozen to fully melted at a controlled temperature in the water-filled pulse tube. We also monitored the ice distribution inside the sample using a remote microwave method that determines the dielectric constant of the sample. Our initial results showed that P-wave velocity increased with ice saturation. The frozen sand and sandstone velocities were nearly 100% and 40-60% higher than the fully melted ones. Attenuation decreased with ice saturation in the frozen sand sample to about 50% difference, while it was quite stable in the frozen sandstone sample. These initial results provide insights into how changes in velocity and attenuation can be used to remotely monitor ice and gas hydrates distribution in reservoirs.

Ben Dell

Exploring the use of acoustic indices for habitat level monitoring and event detection

Effectively monitoring the marine environment is an urgent challenge. Climate change, growing exploitation of marine resources and increasing interest in habitat protection and restoration create a rapidly changing and diverse environment with varied stakeholder interests. Acoustics can provide a cost-effective means of monitoring biophonic, geophonic and anthrophonic activity over long time periods. Efficiently processing this data into intuitive and generalisable metrics which can be communicated to non-expert stakeholders remains a key challenge.

Deep learning methods are hindered by the limited number of available labelled data sets and are yet to allow transparent and generalisable habitat-level assessments. Acoustic Indices are statistical measures which capture key features of the soundscape and have been used extensively in terrestrial habitats to detect species in acoustic data, measure habitat gradients and estimate ecological parameters such as species diversity. In the marine environment, habitats and restoration gradients have been successfully classified using simple machine learning with acoustic indices as input features.

Our project applies acoustic indices to labelled data from the COMPASS array in Scotland, building on previous work done around the key acoustic drivers in this region. We will attempt to demonstrate the utility of acoustic indices to distinguish between labelled classes to demonstrate their use in targeting labelling efforts. The project will also use Acoustic Indices to distinguish key soundscape drivers, using simple machine learning methods with Acoustic Indices as features to distinguish classes such as location, season, and diel variations. This work aims to provide accessible methodologies for researchers who are not acousticians, potentially allowing acoustics to complement established methods in marine habitat assessment such as eDNA analysis and visual surveys.

UK Acoustics Network
3rd Underwater Acoustics PhD Symposium Day
13 December 2023, National Oceanography Centre

Edward Clark

Optimising sensor path planning with reinforcement learning and passive sonar modelling

This talk presents a solution for optimising sensor path planning in marine sensor management using reinforcement learning (RL). RL is a type of machine learning where an intelligent system, known as an agent, learns how to make effective decisions by interacting with its environment. Acoustic propagation modelling is integrated into the RL framework using a standard python RL library gym and using algorithms from the stablebaselines3 library. The observation space encompasses signal-to-noise (SNR) information, platform position, bathymetry, and sound speed data. The action space is discretised into 16 horizontal directions and 3 vertical levels, resulting in a 49-dimensional action space. The reward function combines penalisation for movement and rewards for navigating to high SNR regions.

SNR is calculated using PyRAM, a Python implementation of the RAM parabolic equation solution to the Helmholtz Equation. The RL agent uses proximal policy optimisation to learn the management policy. The learnt policy is compared against a gradient ascent policy and an 'oracle' policy which can use perfect knowledge of the source location for direct navigation. The learning process converges to a stable median between 2.5 and 3 million learning steps. The results demonstrate that the learnt policy closely matches the 'oracle' policy in both reward distribution and behaviour. It also outperforms the gradient ascent policy in a realistic environment.

Abdullah Olcay

Deep learning based underwater sound classification using stacked cepstral features

This study evaluates the impact of Mel-frequency cepstral coefficients (MFCCs) stacked with their dynamic features, in contrast to the use of linear frequency spectrograms, as input to deep learning models. While this approach has seen success in music genre classification, speech recognition, and vessel noise tasks, its potential in classifying diverse underwater sound sources remains underexplored. Using data from the COMPASS project's PAM moorings in Western Scotland, this study scrutinizes recordings from three different locations—Tolsta, Hyskier, and StoerHead—across various seasons, capturing diverse underwater sounds reflective of regional dynamics. Employing some state-of-the-art Convolutional Neural Network (CNN) models such as ResNet50, MobileNetV1, and EfficientNetB0 with stacked MFCCs, preliminary results show that acoustic data representation in the network front-end has the potential to mitigate the effect of differences in back-end models. Early findings demonstrate that the stacked MFCCs approach provides a more generalisable classification across acoustic sites and relative insensitivity to DL model variations.

Francesca Pereira

Spice driven sound speed fluctuations in Fram Strait during Summer 2022

The thermohaline structure of the Arctic Ocean is predicted to change drastically over the 21st century, with potential consequences for underwater sound speed and acoustic propagation. Density-compensated fluctuations in temperature and salinity, known as spice, can manifest in the form of fine-structure thermohaline intrusions and spice fronts. This study seeks to quantify the impact of spice on underwater sound speed and acoustic transmission loss by forcing an acoustic ray trace model with in-situ observations of temperature and salinity, collected in the Fram Strait during summer 2022. The effects of internal waves and spice on underwater sound speed are separated through analysis of isopycnal depth displacement. Spice is found to be the largest source of variability in sound speed in the upper 100 m, and can cause sound speed anomalies of up to 15 m/s. These results suggest that predicted future changes to Arctic thermohaline structure will have a profound influence on underwater acoustics.

UK Acoustics Network
3rd Underwater Acoustics PhD Symposium Day
13 December 2023, National Oceanography Centre

Denise Risch

INVITED TALK: Passive acoustics for long-term monitoring of protected species and habitats: current results and future directions

The INTERREG VA COMPASS and MarPAMM projects have established a comprehensive passive acoustic monitoring array with the aim to better understand underwater sound and marine mammal distribution patterns in the border region of Scotland and Northern Ireland. In addition, in 2021 the SAMOSAS project monitored the waters to the west of the Outer Hebrides, while the SEAMONITOR project acoustically monitored the Islay Front region. Together, over a period of five years (2017-2022) the projects generated an impressive data legacy. The data from the COMPASS project alone comprised of more than 100 GB of click detector data and more than 120 TB of acoustic broadband data across its ten monitoring locations.

The output of these projects to date have provided new information on harbour porpoise, baleen whale (i.e., minke and humpback whales) and delphinid species distribution patterns as well as average underwater sound levels in the region. This talk will give an overview of current results and highlight what we have planned for the future.