

Article

Automated Detection and Tracking of Marine Mammals in the Vicinity of Tidal Turbines Using Multibeam Sonar

Douglas Gillespie *¹, Gordon Hastie, Jessica Montabaranom, Emma Longden¹, Katie Rapson¹, Anhelina Holoborodko¹ and Carol Sparling

Sea Mammal Research Unit, Scottish Oceans Institute, University of St. Andrews, St. Andrews KY16 8LB, UK; gdh10@st-andrews.ac.uk (G.H.); jlm35@st-andrews.ac.uk (J.M.); egl5@st-andrews.ac.uk (E.L.); khr1@st-andrews.ac.uk (K.R.); ah319@st-andrews.ac.uk (A.H.); ces6@st-andrews.ac.uk (C.S.)

* Correspondence: dg50@st-andrews.ac.uk; Tel.: +44-1382540245

Abstract: Understanding how marine animals behave around tidal turbines is essential if we are to quantify how individuals and populations may be affected by the installation of these devices in the coming decades. Our particular interest is in collision risk, and how this may be affected by the fine-scale behaviour of seals and small cetacean species around devices. We report on a study in which multibeam sonar data were collected close to an operational tidal turbine in Scotland continuously over a twelve-month period. The sonars provide high-resolution (a few cm) data over a 120° angle out to a range of 55 m at a rate of 10 frames per second. We describe a system which uses automatic computer algorithms to detect potential targets of interest, verified by human analysts using a sophisticated computer user interface to confirm detections and assign target species. To date, we have identified 359 tracks of marine mammals in the data, as well as several thousand tracks from fish and diving birds. These are currently being parameterised to study how these species react to the moving turbine rotors, and the data are now being used to explore the development of improved automated detection and classification algorithms.

Keywords: marine mammals; multibeam sonar; seals; tidal turbines; collision risk



Citation: Gillespie, D.; Hastie, G.; Montabaranom, J.; Longden, E.; Rapson, K.; Holoborodko, A.; Sparling, C. Automated Detection and Tracking of Marine Mammals in the Vicinity of Tidal Turbines Using Multibeam Sonar. *J. Mar. Sci. Eng.* **2023**, *11*, 2095. <https://doi.org/10.3390/jmse11112095>

Academic Editor: Alberta Mandich

Received: 14 September 2023

Revised: 24 October 2023

Accepted: 29 October 2023

Published: 1 November 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

As increasing numbers of structures associated with marine renewable energy are installed in our oceans, it is important to understand how animals behave in these changed environments. Noise disturbance during the construction phase is clearly a cause for concern [1–3] and is known to cause temporary displacement from construction areas [2,4,5]. Conversely, there is compelling evidence that some structures may be creating artificial reefs, which can make a preferential foraging ground for some marine mammals [6]. Unlike wind turbines, tidal energy generators also have large moving parts underwater, which could injure or kill marine mammals if they collide with them [7]. It is therefore important to understand fine-scale animal movements around structures, particularly for rare and protected species, to understand both the potential risks and benefits of these installations.

Our particular interest is in fine-scale interactions and the potential for collisions between marine mammals and tidal stream turbines. Previous work has successfully tracked harbour porpoises around operational tidal turbines using passive acoustics [8,9]; however, the two seal species present in UK waters, grey seals (*Halichoerus grypus*) and harbour seals (*Phoca vitulina*), rarely vocalise underwater, and although broad-scale surface distributions and diving behaviour have been measured using satellite telemetry [10], no data is yet available on their fine-scale underwater movements in the vicinity of operating turbines. While marine mammals are the species of primary interest to UK regulators, other jurisdictions are also interested in commercially important fish species [11]. Further, relating the presence and movements of top predators to that of potential prey may also

help us to understand the motivation animals may have to interact with these structures, and to predict the potential impact of larger-scale installations, potentially comprising many tens of similar structures [12].

Tracking animals underwater in a highly energetic mobile environment is extremely challenging. Underwater cameras have very limited ranges, don't work during periods of darkness, and can be onerous to process [13,14]. Passive Acoustic Monitoring (PAM) can be used to detect and track soniferous species, such as harbour porpoises [15], but is ineffective for species which either don't vocalise, or only vocalise rarely [16]. Animal-borne motion and sound recording tags, such as the dTag [17], can provide detailed movements of animals underwater; however, there are challenges with accurately geo-referencing underwater position estimates, which must be inferred from surface GPS locations and the integration of accelerometer data. More importantly, however, tagging animals is challenging, invasive, and when the research interest is in a very specific location (such as a turbine foundation), the likelihood of any of a small number of tagged animals visiting the site is generally small. High-frequency multibeam sonar is a relatively new technology, which can be used to detect and track animals underwater irrespective of their behaviour and is unaffected by daylight or meteorological conditions [18–20].

We have recently completed a 12-month data collection period using two high-frequency multibeam sonars (Tritech Gemini 720id: Tritech International Ltd., Westhill, Aberdeenshire, UK) at an operational tidal turbine in Scotland. Each sonar has 512 beams in the horizontal plane and scans a horizontal swath of 120°. The vertical -3 dB beam width is $\pm 20^\circ$ [21]. Range resolution is approximately 3 cm, so the sonars provide extremely high resolution in the horizontal plane, but very limited information in the vertical. The sonars emit an acoustic signal approximately 70 μ s in duration and have a source level of approximately 200 dB re 1 μ Pa (broadband), with a main lobe at 720 kHz [21]. The frame rate across the two sonars is approximately 10Hz. Previous research has shown that these sonars can reliably detect marine mammals up to ranges of approximately 45 m, and the acoustic signals do not appear to elicit overt behavioural responses in marine mammals [19].

Previous work [19] used an automated target detection and tracking module within the SeaTec software supplied by the sonar manufacturer to detect discrete objects in the sonar data, within user-defined size and persistence bounds. This used a flood-fill algorithm approach, e.g., [22], to summarise the shape and intensity patterns exhibited by each target and, if the target was within the user-defined specifications, recorded basic information to text files on timings (hh-mm-ss), locations (X-Y coordinates), ranges from the sonar (m), and kinematic information (speed and trajectory in the X and Y planes) for all mobile targets detected in the data [21]. These outputs were then classified as either seals or non-seals post hoc using Support Vector Machine classification algorithms. More recent work by Williamson et al. [23] and Cotter et al. [24,25] employed similar detection algorithms to detect animal tracks in data collected with an Imagenex 837B Delta T multibeam echosounder and a Teledyne BlueView M900-2250 sonar, respectively. Our intention was to use the SeaTec software, as it is compatible with the Tritech sonars. However, while the SeaTec detection algorithm could run in real-time in a static environment, when deployed close to a turbine the moving turbine rotors created many thousands of false candidate detections per minute, rendering the software chain unstable and unable to collect data. Our data collection therefore reverted to the collection and storage of raw sonar data only, with no online data processing.

Here, we report on methods developed to extract useful target features from the data. First, the data from the sonars were processed fully automatically to detect objects and to link them into tracks. These tracks were then examined manually by a human operator to classify them within biological or non-biological groups. Below, we describe the detection and tracking algorithms, and a sophisticated User Interface (UI) designed to present data succinctly to an analyst to make final decision making as efficient as possible. We also

discuss future developments, which include real-time operation and further automation of the classification process.

2. Materials and Methods

2.1. Equipment and Data Collection

The data were collected using two Tritech Gemini 720is multibeam sonars mounted on a platform positioned approximately 30 m N of an operating tidal turbine in the Pentland Firth, on the north coast of Scotland [26]. The turbine being monitored is a horizontal axis turbine mounted on a 25 × 19 m tripod support structure. The turbine sits at 37.2 m depth, the rotor centre is 14 m above the sea floor, and each of the three blades is 9 m long. The nacelle (generator) rotates horizontally on a yaw mechanism to always face the tidal current. The rotors themselves are 3.7 m in front of the yaw centre, and the ebb and flood currents are offset by 150°, meaning that the position of the rotors changes by several metres four times per day.

The monitoring platform [26] was connected to a subsea structure with power and communication connections for up to four turbines (known as the “subsea hub”) via an umbilical cable, to receive power and send data to shore via optical fibre. These were connected to a desktop computer in the turbine substation, which ran software to control the sonars and save data to disk. The internet connection to the substation was insufficient for data transfer, so data files were automatically backed up to external hard drives and sent by post to our lab for processing.

The sonars operate at a frequency of 720 kHz (2.1 mm wavelength). Each monitors a horizontal swath of 120°, with 512 beams giving an angular resolution of ~0.25° and a range resolution of ~3.3 cm. Individual sonars have a vertical beam half-width of 20°. Note that the angular resolution of these sonars is twice that of the models used in previous work [19], which covered the same 120° angle with only 256 beams. Thirty metres from the turbine, the 120° angle provides monitoring to a nominal 26 m up- and downstream of the turbine centre. The use of two sonars with a 30° vertical offset allows us to view the full vertical height of the turbine. While the individual sonars can provide no vertical information (beyond being detected on the “higher one” or “lower one”), by comparing the relative strength of target reflections on the two sonars, it is possible to gain some vertical information when animals are in the overlap between them [20]. Both sonars were set to a maximum detection range of 55 m.

The data collection used the Tritech Genesis software supplied with the devices (version 1.9.0.73). Raw data from the sonars take the form of “image frames”, which are a rectangular array of unsigned 8-bit (0–255) intensity values, versus angles and ranges, plus associated metadata that includes device gain settings, timestamps, sound speed measurements, and a list of beam angles. Each frame has 512 angles and typically ~1660 ranges, although the latter varied slightly with a measurement of sound speed made by each sonar. To avoid interference between sonars, images were generated alternately at a rate of ~5 per second per sonar. This generates approximately 730 Gbytes of raw data per day. Compression of the data within the files reduces data volumes to around half that. Frames from both sonars are written to a single file, with new files being created when their user-specified size is reached (here ~300 Mbyte), and therefore contain less than two minutes of data. Particularly at small ranges, the resolution is very high, i.e., for an object 1 m from the sonar, the horizontal distance between adjacent beams is only 4 mm; however, at the maximum range of 55 m, this increases to 22.5 cm, though this is still smaller than our primary targets of interests: harbour porpoises with a body size of 90 to 150 cm [27], and harbour and grey seals with typical adult body sizes of 130 to 160 cm and 140 to 190 cm, respectively [28,29].

2.2. Data Processing

All offline processing was carried out using our own software, specifically developed for this purpose. The software was developed as a plugin for the PAMGuard passive acous-

tic software [30] to utilise the existing PAMGuard infrastructure for automatic processing, followed by human operator browsing and annotation of large datasets. Our workflow for processing sonar data is broadly similar to that for passive acoustic data described in [31] for the monitoring of cetaceans, where automatic detectors are initially applied to a large dataset to detect candidate vocalisations, which are then manually checked and further annotated by an operator viewing sophisticated data displays.

For display purposes, the data were transformed from the rectangular image frame to a fan-shaped image, which then relates directly to real-world x and y coordinates. A sample raw image frame and its corresponding fan image are shown in Figure 1.

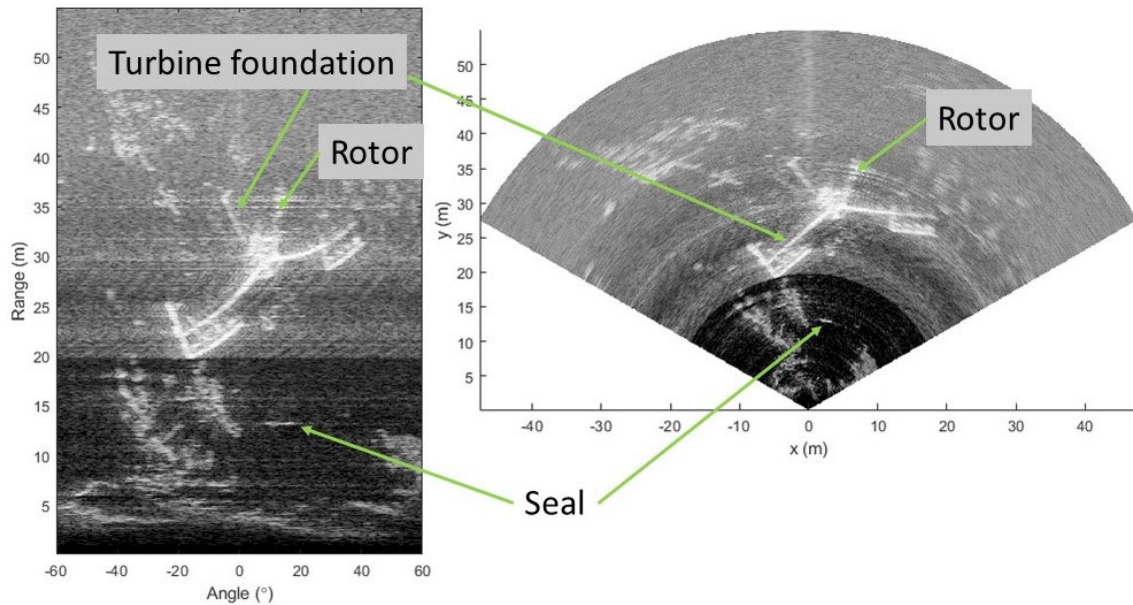


Figure 1. Raw data from a single sonar image and the same data transformed into a geometrically correct “fan” image. The triangular tripod base of the turbine is clearly visible, as is the central tower, slightly offset to the right of the image centre. Also visible is a seal at a position of 11°, 13.4 m in the untransformed data or $xy = (2.4, 13.1)$ m in the fan image. Other high-intensity targets in the image are believed to be rocks on the seafloor.

2.3. Automatic Processing

Detection of targets in the sonar data is a two-stage process. The first stage (“region detection”) measures background (constant) levels in the data and picks out regions of each image frame rising above some threshold. The second stage (“track detection”) searches for track segments, i.e., nearby regions in space and time which appear to be moving in a consistent direction. Region detection takes place on the rectangular image frames, though outputs are transformed to real-world $(x,y$ or $\theta,r)$ coordinates for track detection.

2.3.1. Region Detection

The background, or constant, level in the data is measured independently for each sonar on a pixel-by-pixel basis using a single-order low-pass filter, e.g., if $S_{\theta,r,t}$ is the raw data for beam angle θ , range r and time t , then the background, $B_{\theta,r,t}$, is described by Equation (1).

$$B_{\theta,r,t} = (1 - \alpha)B_{\theta,r,t-1} + \alpha S_{\theta,r,t} \tag{1}$$

where α is a small number, typically 0.05, meaning that the background updates over a time of around 4 s.

For region detection, a double-threshold method is used, whereby at least one pixel in each region must be above the higher threshold, but the outer limits of the region are determined using a slightly lower threshold. The thresholds also consist of two components,

a scale term, c , which is a multiplier of the background noise, and a constant term, Th , i.e., for a point to be above the threshold, it is required that

$$S - B \times c > Th \quad (2)$$

where c is set to 1.2, and Th to 20.

Points above the threshold are assigned to the same region if they are adjacent (i.e., one bin different in r or θ) or if they touch on the diagonals (i.e., one bin different in r and θ). The practical implementation of the software also has simple limits on region size, with any region (following transformation to real-world x and y coordinates) outside those defined size limits being discarded (0.5 to 4.0 m for this study). For each region, the details of which pixels were within the region are currently discarded, with the software retaining the minimum and maximum angles and ranges for the region, the angle and range to the point of highest intensity in the region, the mean intensity, and the occupancy, i.e., the fraction of pixels within the stored limits which were above threshold.

2.3.2. Track Detection

We use the term “track segment” to refer to an automatic track detection from the algorithm and the word “track” to mean a group of track segments marked by the human reviewer and considered to be the trajectory of a single animal or a nearby group of animals (e.g., within a fish school).

For automatic track-segment detection, data from both sonars are combined. For our current installation, where there is a vertical, but no horizontal, offset, an object appearing in both sonar beams should have the same x and y coordinates, with only the intensity (or indeed, presence) of the signal varying in each sonar. Detected regions are joined to a track segment if the detected points are within 0.5 s of one another, the distance between them represents a speed of <5 m/s, and the size of the detections varies by less than a factor of 3. Many short tracks containing only a small number of points are initially generated but are only retained if they have 10 or more connected points.

In the track detection stage, it is common for an animal track to be broken into several parts, each logged separately by the detector. There are several reasons for this: firstly, an animal may move in and out of the sonar beam, making it temporarily unavailable for detection; additionally, small (and therefore weaker) targets (single fish, or bubbles/turbulence in the water) may be detected, then not detected, for multiple short periods, and targets such as fish schools, by their nature, form many short track segments of individual animals.

Due to the high levels of detected regions that were not assigned to a track segment, only those forming a track segment were retained and written to PAMGuard data storage files for further processing by a human operator. A new output file was automatically started by the software for each hour of data.

2.4. User Interface and Manual Processing

With 24/7 data collection over an extended period, it is essential that a human operator can view and confirm detections post hoc in the data many times faster than in real time. To do this, the operator must be able to view both the detection data as well as the underlying sonar images.

A typical day's data consists of several hundred raw data files and 24 PAMGuard output files of detected track segments. The key motivations in the design of the software user interface were to provide rapid access to any point within the dataset and to minimise the number of keyboard or mouse actions required to annotate a track, so that the operator could direct all their attention to viewing both the sonar images, the detected track-segment data, and the “playbacks” of sequences of images. PAMGuard contains a two-tier scroll system to efficiently navigate large datasets: first, a map of available data is created, which can be held within memory, thus allowing the software to rapidly identify which files contain data for a specified time period. The user can then page through that map of data, typically loading 15 min of data into memory for viewing. While the detection data for a

15-min period requires very little memory, even for a short 15-min period, loading all the sonar raw data (approximately 9000 frames) would require 10's gigabytes of memory and would be slow. Therefore, for the image data, only a secondary fine-scale map of the times, file names, and positions within the files of every data frame within that 15 min is loaded. Individual frames are then rapidly loaded as required by the display.

A PAMGuard sonar display (Figure 2) shows the transformed sonar images for both sonars, with a scroll bar at the bottom of the display allowing movement through the loaded 15 min of data. Detection data can be shown over this, either for the full 15 min, the time of the single displayed frames, or for a short time period (typically a few seconds) immediately before the time of the displayed frame. More than one display can be created within the software, which allows the user to view one with and one without the detection overlays, generally spread across multiple monitors. Other features designed to make viewing data as easy as possible are:

1. The operator can zoom into any part of the image by moving the scroll wheel of the mouse.
2. If the mouse is hovered over a detection, a small pop-up image of the underlying sonar data at the time of that detection is shown automatically. This allows the user to sweep the mouse over a track and get a quick view of all the underlying images without having to scroll the images themselves.
3. Clicking on a detection allows the user to position the scroll bar at the exact time of that detection so that the full underlying image can be viewed.
4. The "play" button will automatically advance the scroll bar, showing a video-like display of the sonar image data. This can be at normal speed or 0.5, 2, 5, or 10 times normal speed.
5. Additional scroll arrows can advance or wind back the display a single frame at a time, which can prove easier than trying to move a scroll bar by one or two pixels.
6. The image data can be shown with the background (Equation (1)) subtracted prior to fan transformation, which can make tracks more visible. The displays also offer a variety of colour schemes both for the sonar images and for the detection overlays, including colour-blind palettes, which are a standard PAMGuard feature.

Once 15 min of data are loaded into memory, a quick look at the detection data is generally enough to determine if any tracks of interest are present. If there are no tracks or only very short track segments in regions where false detections frequently occur, then the operator can move immediately to the next 15 min of data. If tracks are present, then the operator uses the above software features to examine the underlying image data to make a manual classification. They can then mark the track with the mouse and enter their classification, which, along with the details of which individual detected the track segments they include in the track, is written to a relational database.

A typical workflow is therefore as follows:

- Load the catalogues for a day's data.
- Step through the data in 15-min sections, viewing all the detection data for that period.
- If no tracks are visible on the displays, step to the next 15 min.
- If tracks are visible, view the underlying sonar images by hovering the mouse over detections, scrolling through each track, and playing the tracks as videos. Zoom in as required to get the best view of the underlying data, etc., to make a decision.
- Mark the track, or group of tracks, with the mouse and fill in a form with species information.
- Repeat steps 4 and 5 as required before moving to the next 15 min.




Figure 2. PAMGuard sonar display. The left figures show transformed images from the two sonars, and on the right are the same images overlaid with 15 min of detection data (during most analyses, we would spread these images across two monitors to increase their resolution). An important feature of the display is the scroll system, which allows easy movement through the 15 min of loaded data, irrespective of the fact that the raw data will be in many different raw data files. The controls in the top right of each plot allow fine-scale scrolling of 1 or 10 frames at a time. In the bottom right of the display, the sideways triangular button will cause the data to “play” either in real time, or faster or slower. The blue double arrows allow easy movement to the next (or preceding) 15 min of data. Several features have been incorporated to make it as simple and fast as possible for the operator to come to a decision regarding the obvious track across the bottom of the display. Hovering the mouse over any detection in the track will show a small pop-up window with a $3\text{ m} \times 3\text{ m}$ clip of that part of the raw sonar image corresponding to that detection. Right-clicking on the detections can move the scroll bar of the displays to exactly the start of that track so that it can be viewed in the left display without the detection overlay. The display can be zoomed by using the mouse’s scroll wheel to examine the images at higher resolution. Once the track is confirmed, it is marked with the mouse and appropriate identifying metadata, and the operator’s classification is added to an underlying database.

3. Results

Annotation of 24 h of dual sonar data typically takes around 2–3 h for an experienced operator. To date, we have completed manual auditing of 266 days of data.

The majority of detected track segments are created by the movement of the turbine blades. These occur at an average rate of around one per second in slack water, rising to 10 per second during turbine operation. False detections are also caused by reef areas on the seabed, and sediment and/or seaweed moving in the flow. These occur near-constantly in the same location, resulting in the same track pattern, and are easy to distinguish as false positives and ignore. At high flows and rotation speeds, turbulence from the blade tips can cause large numbers of detections running downstream from the rotors (Supplementary Materials Figures). At high flows, we believe that our platform vibrates slightly in the current, so the apparent movement in the images can cause further false detections. Finally,

during bad weather and certain tidal states, large “clouds” of high intensity pass through the images, clearly adding noise to the data and reducing our ability to detect. These are most likely caused by air being sucked into the water column, though may also be in part caused by sediment movement.

Tracks of interest within the current dataset include marine mammals, elasmobranchs, individual fish, fish schools, and diving birds. An example seal track is shown in Figure 2. Examples of other track types, including videos of these, are available in Supplementary Materials. The example images shown in Figures 1 and 2 are from a period of slack water, when the turbine was “yawing” from its flood to ebb tide position. In addition to the obvious track at a 13–20 m range, there are a large number of false detections close to the turbine structure.

From 266 days of analysed data, a total of over 28 million candidate track segments were detected. Of these, a total of 14,324 tracks were labelled by a reviewer, comprising 156,884 detected track segments, which is just 0.5% of all detections (Table 1). The mean number of track segments grouped into a track is 4.8 for marine mammal tracks and 29 for fish schools; this reflects the single vs multiple nature of marine mammal and fish school targets. As with fish schools, non-biological tracks (turbulence and bubbles) consist of a large number of track segments, due to the broken nature of the targets.

Table 1. Numbers of tracks by type and numbers of segments making up those tracks.

Track Type	Human-Annotated Tracks	Track Segments	Segments per Track
Marine mammals ¹	359	1720	4.8
Fish schools	4582	132,572	29
Individual fish	5554	8014	1.4
Birds	215	5898	27
Unidentified	3614	8680	2.4
Total	14,324	156,884	11

¹ includes seals and harbour porpoises (although the majority are seals—only 4 HP in 131 days).

“Unidentified” tracks are generally small and are therefore unlikely to be marine mammals. Most are probably small single fish, but they may also be other particulate matter in the water column.

4. Discussion

Our platform successfully collected data 24 h a day for a one-year period. Further analysis of this dataset will provide insights into the fine-scale movement of marine mammals in the vicinity of an operational tidal turbine. Due to movement restrictions during the COVID-19 pandemic, it was not possible to test the platform in a high current area prior to the year-long deployment, leading to some difficulties with the data. First, the acquisition software was found to be unstable when confronted with the moving parts of the turbine, and, secondly, the platform appeared to vibrate slightly, again creating false detections. Future data collection will likely use PAMGuard for real-time operation, which has an inbuilt reset function which will help recover from any instabilities in data collection. Future platforms will also be made more rigid in an attempt to reduce vibration.

One of the main purposes of our work was to efficiently process a very large dataset. A manual review of this volume of data, by replaying it as a “video” and having a human operator review it either in real-time or perhaps at 2× speed, would have been impossible without a large team of operators, and such a system would still need a method to extract track details when an object of interest was spotted. Our system reduced this workload to around 2 working days per week of data, or approximately 11× real time, which is still onerous but achievable. Further, all the output was recorded in a database, making the extraction of parameters such as swim speed and changes in direction straightforward. The

technology and software used in this application demonstrated good capability of detecting a wide range of animal species, including marine mammals, elasmobranchs, birds, single fish, and fish schools.

Our current processing algorithm is intentionally simple, using no information on animal shape or movement behaviour, and the processing chain is highly dependent on the work of human operators to select and label what they consider to be genuine animal tracks, as opposed to false detections from the moving turbine. However, future projects may require fully automated real-time data analyses to trigger mitigation actions should an animal approach an operating turbine. Further, even in applications where real-time triggers are not required, real-time data processing, or more rapid offline processing, during long-term monitoring programmes has a number of benefits, including data quality control, reduced data storage costs (e.g., by only storing raw data when a detection occurs), and reduced data turnaround times.

As we build up a database of detected animals, there is scope for developing improved algorithms, both for the detection and classification of tracks. Modern machine learning methods, such as convolutional neural networks, would appear well suited to this task and have recently achieved much improved results compared to earlier algorithms for processing passive acoustic data [32,33]. Such algorithms learn (or train) directly from data, so it would not have been possible to employ such algorithms for an initial pass through this dataset. However, now that we are building a large set of human-annotated detections, we are preparing to test machine learning methods on our data in the future. We do, however, approach this task with caution, and while we believe that accurate real-time detection of marine mammals is possible, we also assume that, for the foreseeable future at least, there will continue to be a role for human analysts checking at least a subset of detections. First and most importantly, for this application, we are looking for rare events and unexpected behaviours; since machine learning algorithms learn from example data, they must not be over-trained to the point where they reject detections because they are too dissimilar to the data used for training. Secondly, our data are all derived from a single turbine site; other turbine designs may create quite different false targets in the sonar images, so it is likely that any AI algorithms would need at least some retraining when applied to a different turbine design.

The semi-automated workflow presented here is not dissimilar to that of similar image processing tasks such as cancer screening, where a qualified radiologist remains the ultimate decision maker, but algorithms are used to reduce their workload [34–36]. The purpose of the algorithm is often not so much to find the object of interest, whether it be a tumor or marine mammal, but to reject images that are clearly of no interest to allow the expert to concentrate on images where there may be something of importance. In their review paper, “What the radiologist should know about artificial intelligence”, Visser et al. [35] ask the question “Will radiologists be replaced by AI?” and give the very simple answer “No”; we believe that the same is true for the processing of sonar and other similar (e.g., passive acoustic) data. However, we do believe that AI will make an important contribution, reducing the amount of data requiring expert review, and our aim is to reduce the time taken to manually verify a week’s data to considerably less than one day’s work.

In our data collected at a turbine site, the majority of false detections were caused by moving turbine parts. The location of the turbine rotors is well known, so it would be straightforward to blank out (spatially veto) the regions of the image where the rotors are for a particular tidal flow. Other false detections were created at high flows by turbulence shedding from the tips of the blades. These were less clearly defined spatially than the rotor positions, so could not easily be removed with a spatial veto around the blades. However, how data are processed depends somewhat on the application and the questions being asked. For instance, a key interest for regulators is the possibility of animals approaching and entering the turbine from upstream and being injured by collision with the turbine blades. Similarly, there is interest in triggering mitigation actions, such as the use of an acoustic deterrent, if animals are approaching an operating turbine. If the sole purpose

was therefore to detect animals upstream (i.e., where the flow is carrying them towards the blades), then it would be reasonable to extend a spatial veto over the entire blade and downstream region. This may miss animals downstream but may be a realistic paradigm for automated mitigation actions and to further reduce human workload for processing future datasets. Again, as with the application of machine learning, our current dataset will be used to assess such ideas.

Software Implementation Availability and Performance

Running on a modern PC (12th Gen Intel i7-12700, 2.1 GHz, 32 Gbyte RAM, Windows 11), the automatic algorithm can process archived data read from an external USB drive at around $28\times$ real time. As stated above, offline review of data using experienced reviewers proceeds at around $11\times$ real time.

For efficient data loading, the first time each raw data file is accessed, a catalogue of the file contents is generated and written to the same file directory for future use. When processing offline, this allows access to any given record without having to read the preceding records in the file. Without this step, finding records near the end of a file could take >1 s. When processing data offline, 15-min sets of detection data typically load in a few seconds (depending on the number of detections). Individual image frames can be accessed at random from a USB drive with typical read times of 10 ms and image transform times of 12 ms. Load times reduce to around 5 ms when accessed from an internal solid-state hard drive. With a total of 22 ms to load and transform a sonar image, this is fast enough to provide smooth scrolling/video-type playback, since the interval between frames is 100 ms.

Overall, the speed of the algorithm and offline data loading means that the system can be used in real-time on a modest PC or laptop, and that when browsing data offline, smooth scrolling and data viewing are easily achieved. The system was not deployed for real-time operation at the turbine site, since our priority was stable data collection. The system has, however, been used with a single sonar for real-time detection of seals in Scottish rivers and has also been deployed on a tidal turbine under test at the EMEC facility in the Orkney Islands. In all instances, the software is stable, archiving raw data and running the automatic detectors in real time.

The implementation within PAMGuard comprises three separate Java software packages. One package contains a Java reader for Trittech data files and methods for rapidly cataloguing multiple files. This package has been kept free of PAMGuard-related classes and can be built for Java 8, meaning that its methods can be used to import sonar data into Matlab[®] if desired. The second package is a Java wrapper around the Trittech Software Development Kit (SDK) version 2.0.37.9, making the C and C++ library functions in the SDK available within Java, which is required for real-time operation within PAMGuard. The third package is the PAMGuard plugin, which combines the functionality of the other packages with that of PAMGuard to generate displays, control detector function (e.g., through parameter control dialogues), provide data storage, data annotation, real-time acquisition, etc. PAMGuard is inherently a modular software package, and, although we have not investigated processing data from other sonar manufacturers, if the sonar comes with suitable libraries for data acquisition, then it would be relatively straightforward to replace the Trittech acquisition module with a module to acquire from a different sonar device, and then feed those data into the same processing modules, just as the passive acoustic detection modules in PAMGuard can work with data from a wide variety of acoustic acquisition devices.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/jmse11112095/s1>.

Author Contributions: G.H., D.G. and C.S. developed the overall monitoring program for the turbine site in Scotland. D.G. developed the PAMGuard software and the additional software modules for processing the sonar data. J.M., E.L., A.H., K.R. and G.H. worked on data annotation. All authors have read and agreed to the published version of the manuscript.

Funding: The monitoring platform was developed with funds from the Natural Environment Research Council (Grant Nos. NE/R015007/1 and NE/R014639/1). Software development and data analysis were funded by the Scottish Government as part of the Marine Mammal Scientific Support Program (Grant No. MMSS/002/15). Umbilical cables to the turbine infrastructure were funded and developed by SIMEC Atlantis.

Institutional Review Board Statement: The animal study was reviewed and approved by the University of St Andrews Animal Welfare and Ethics Review Body (Ref number: SEC18017).

Informed Consent Statement: Not Applicable.

Data Availability Statement: All the software is available under the GPL3 opensource software licence at <https://github.com/douggillespie/TritechFiles>, <https://github.com/douggillespie/GenesIsJavaJNAInterface>, <https://github.com/douggillespie/TritechAcquisitionPlugin>, respectively (URL accessed on 20 October 2023). Sample data are available on request to the authors.

Acknowledgments: The deployment of the platform and twelve months of data collection would not have been possible without the hard work and enthusiasm of the SIMEC engineering team. We are also grateful to the editor and two anonymous reviewers for their useful comments on earlier drafts of this paper.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Benhemma-Le Gall, A.; Graham, I.M.; Merchant, N.D.; Thompson, P.M. Broad-Scale Responses of Harbor Porpoises to Pile-Driving and Vessel Activities during Offshore Windfarm Construction. *Front. Mar. Sci.* **2021**, *8*, 664724. [[CrossRef](#)]
2. Graham, I.M.; Merchant, N.D.; Farcas, A.; Barton, T.R.; Cheney, B.; Bono, S.; Thompson, P.M. Harbour Porpoise Responses to Pile-Driving Diminish over Time. *R. Soc. Open Sci.* **2019**, *6*, 190335. [[CrossRef](#)] [[PubMed](#)]
3. Graham, I.; Gillespie, D.; Gkikopoulou, K.; Hastie, G.; Thompson, P. Directional Hydrophone Clusters Reveal Evasive Responses of Small Cetaceans to Disturbance during Construction at Offshore Windfarms. *Biol. Lett.* **2023**, *19*, 20220101. [[CrossRef](#)] [[PubMed](#)]
4. Russell, D.J.; Hastie, G.D.; Thompson, D.; Janik, V.M.; Hammond, P.S.; Scott-Hayward, L.A.; Matthiopoulos, J.; Jones, E.L.; McConnell, B.J. Avoidance of Wind Farms by Harbour Seals Is Limited to Pile Driving Activities. *J. Appl. Ecol.* **2016**, *53*, 1642–1652. [[CrossRef](#)] [[PubMed](#)]
5. Whyte, K.F.; Russell, D.J.; Sparling, C.E.; Binnerts, B.; Hastie, G.D. Estimating the Effects of Pile Driving Sounds on Seals: Pitfalls and Possibilities. *J. Acoust. Soc. Am.* **2020**, *147*, 3948–3958. [[CrossRef](#)] [[PubMed](#)]
6. Russell, D.J.; Brasseur, S.M.; Thompson, D.; Hastie, G.D.; Janik, V.M.; Aarts, G.; McClintock, B.T.; Matthiopoulos, J.; Moss, S.E.; McConnell, B. Marine Mammals Trace Anthropogenic Structures at Sea. *Curr. Biol.* **2014**, *24*, R638–R639. [[CrossRef](#)] [[PubMed](#)]
7. Onoufriou, J.; Brownlow, A.; Moss, S.; Hastie, G.; Thompson, D. Empirical Determination of Severe Trauma in Seals from Collisions with Tidal Turbine Blades. *J. Appl. Ecol.* **2019**, *56*, 1712–1724. [[CrossRef](#)]
8. Gillespie, D.; Palmer, L.; Macaulay, J.; Sparling, C.; Hastie, G. Harbour Porpoises Exhibit Localized Evasion of a Tidal Turbine. *Aquat. Conserv. Mar. Freshw. Ecosyst.* **2021**, *31*, 2459–2468. [[CrossRef](#)]
9. Palmer, L.; Gillespie, D.; MacAulay, J.D.; Sparling, C.E.; Russell, D.J.; Hastie, G.D. Harbour Porpoise (*Phocoena phocoena*) Presence Is Reduced during Tidal Turbine Operation. *Aquat. Conserv. Mar. Freshw. Ecosyst.* **2021**, *31*, 3543–3553. [[CrossRef](#)]
10. Onoufriou, J.; Russell, D.J.; Thompson, D.; Moss, S.E.; Hastie, G.D. Quantifying the Effects of Tidal Turbine Array Operations on the Distribution of Marine Mammals: Implications for Collision Risk. *Renew. Energy* **2021**, *180*, 157–165. [[CrossRef](#)]
11. Viehman, H.A.; Zydlewski, G.B. Fish Interactions with a Commercial-Scale Tidal Energy Device in the Natural Environment. *Estuaries Coasts* **2015**, *38*, 241–252. [[CrossRef](#)]
12. Williamson, B.; Fraser, S.; Williamson, L.; Nikora, V.; Scott, B. Predictable Changes in Fish School Characteristics Due to a Tidal Turbine Support Structure. *Renew. Energy* **2019**, *141*, 1092–1102. [[CrossRef](#)]
13. Verfuss, U.K.; Gillespie, D.; Gordon, J.; Marques, T.A.; Miller, B.; Plunkett, R.; Theriault, J.A.; Tollit, D.J.; Zitterbart, D.P.; Hubert, P. Comparing Methods Suitable for Monitoring Marine Mammals in Low Visibility Conditions during Seismic Surveys. *Mar. Pollut. Bull.* **2018**, *126*, 1–18. [[CrossRef](#)] [[PubMed](#)]
14. Hutchison, I.; Tait, C.; Sheehy, J.; Morgan, P. *Review of Underwater Video Data Collected around Operating Tidal Stream Turbines 1225*; NatureScot: Perth, UK, 2020; Volume 33.
15. Gillespie, D.; Palmer, L.; Macaulay, J.; Sparling, C.; Hastie, G. Passive Acoustic Methods for Tracking the 3D Movements of Small Cetaceans around Marine Structures. *PLoS ONE* **2020**, *15*, e0229058. [[CrossRef](#)] [[PubMed](#)]

16. Sparling, C.; Gillespie, D.; Hastie, G.; Gordon, J.; Macaulay, J.; Malinka, C.; Wu, M.; McConnell, B. *Scottish Government Demonstration Strategy: Trialling Methods for Tracking the Fine Scale Underwater Movements of Marine Mammals in Areas of Marine Renewable Energy Development*; Scottish Marine and Freshwater Science, The Scottish Government: Edinburgh, UK, 2016.
17. Johnson, M.P.; Tyack, P.L. A Digital Acoustic Recording Tag for Measuring the Response of Wild Marine Mammals to Sound. *IEEE J. Ocean. Eng.* **2003**, *28*, 3–12. [[CrossRef](#)]
18. Francisco, F.; Bender, A.; Sundberg, J. Use of Multibeam Imaging Sonar for Observation of Marine Mammals and Fish on a Marine Renewable Energy Site. *PLoS ONE* **2022**, *17*, e0275978. [[CrossRef](#)] [[PubMed](#)]
19. Hastie, G.D.; Wu, G.-M.; Moss, S.; Jepp, P.; MacAulay, J.; Lee, A.; Sparling, C.E.; Evers, C.; Gillespie, D. Automated Detection and Tracking of Marine Mammals: A Novel Sonar Tool for Monitoring Effects of Marine Industry. *Aquat. Conserv. Mar. Freshw. Ecosyst.* **2019**, *29*, 119–130. [[CrossRef](#)]
20. Hastie, G.D.; Bivins, M.; Coram, A.; Gordon, J.; Jepp, P.; MacAulay, J.; Sparling, C.; Gillespie, D. Three-Dimensional Movements of Harbour Seals in a Tidally Energetic Channel: Application of a Novel Sonar Tracking System. *Aquat. Conserv. Mar. Freshw. Ecosyst.* **2019**, *29*, 564–575. [[CrossRef](#)]
21. Parsons, M.J.; Fenny, E.; Lucke, K.; Osterrieder, S.; Jenkins, G.; Saunders, B.J.; Jepp, P.; Parnum, I.M. Imaging Marine Fauna with a Tritech Gemini 720i Sonar. *Acoust. Aust.* **2017**, *45*, 41–49. [[CrossRef](#)]
22. Law, G. Quantitative Comparison of Flood Fill and Modified Flood Fill Algorithms. *Int. J. Comput. Theory Eng.* **2013**, *5*, 503–508. [[CrossRef](#)]
23. Williamson, B.J.; Fraser, S.; Blondel, P.; Bell, P.S.; Waggitt, J.J.; Scott, B.E. Multisensor Acoustic Tracking of Fish and Seabird Behavior around Tidal Turbine Structures in Scotland. *IEEE J. Ocean. Eng.* **2017**, *42*, 948–965. [[CrossRef](#)]
24. Cotter, E.; Polagye, B. Automatic Classification of Biological Targets in a Tidal Channel Using a Multibeam Sonar. *J. Atmos. Ocean. Technol.* **2020**, *37*, 1437–1455. [[CrossRef](#)]
25. Cotter, E.; Polagye, B. Detection and Classification Capabilities of Two Multibeam Sonars. *Limnol. Ocean. Methods* **2020**, *18*, 673–680. [[CrossRef](#)]
26. Gillespie, D.; Oswald, M.; Hastie, G.; Sparling, C. Marine Mammal HiCUP: A High Current Underwater Platform for the Long-Term Monitoring of Fine-Scale Marine Mammal Behavior around Tidal Turbines. *Front. Mar. Sci.* **2022**, *9*, 283. [[CrossRef](#)]
27. Kastelein, R.A.; Hardeman, J.; Boer, H. Food Consumption and Body Weight of Harbour Porpoises (*Phocoena phocoena*). In *The Biology of the Harbour Porpoise*; De Spil Publishers: Woerden, The Netherlands, 1997; pp. 217–233.
28. Härkönen, T.; Heide-Jørgensen, M.-P. Comparative Life Histories of East Atlantic and Other Harbour Seal Populations. *Ophelia* **1990**, *32*, 211–235. [[CrossRef](#)]
29. Twiss, S.D.; Pomeroy, P.P.; Thomas, C.J.; Mills, J.P. Remote Estimation of Grey Seal Length, Width and Body Mass from Aerial Photography. *Photogramm. Eng. Remote Sens.* **2000**, *66*, 859–866.
30. Gillespie, D.; Gordon, J.; Mchugh, R.; McLaren, D.; Mellinger, D.; Redmond, P.; Thode, A.; Trinder, P.; Deng, X.Y. PAMGUARD: Semiautomated, Open Source Software for Real-Time Acoustic Detection and Localisation of Cetaceans. *J. Acoust. Soc. Am.* **2008**, *125*, 2547. [[CrossRef](#)]
31. Webber, T.; Gillespie, D.; Lewis, T.; Gordon, J.; Ruchirabha, T.; Thompson, K.F. Streamlining Analysis Methods for Large Acoustic Surveys Using Automatic Detectors with Operator Validation. *Methods Ecol. Evol.* **2022**, *13*, 1765–1777. [[CrossRef](#)]
32. Madhusudhana, S.; Shiu, Y.; Klinck, H.; Fleishman, E.; Liu, X.; Nosal, E.-M.; Helble, T.; Cholewiak, D.; Gillespie, D.; Širović, A. Improve Automatic Detection of Animal Call Sequences with Temporal Context. *J. R. Soc. Interface* **2021**, *18*, 20210297. [[CrossRef](#)]
33. Shiu, Y.; Palmer, K.J.; Roch, M.A.; Fleishman, E.; Liu, X.; Nosal, E.-M.; Helble, T.; Cholewiak, D.; Gillespie, D.; Klinck, H. Deep Neural Networks for Automated Detection of Marine Mammal Species. *Sci. Rep.* **2020**, *10*, 607. [[CrossRef](#)]
34. Fry, H. *Hello World: How to Be Human in the Age of the Machine*; Random House: New York, NY, USA, 2018.
35. Visser, J.; European Society of Radiology. What the Radiologist Should Know about Artificial Intelligence—An ESR White Paper. *Insights Into Imaging* **2019**, *10*, 44.
36. Katzen, J.; Dodelzon, K. A Review of Computer Aided Detection in Mammography. *Clin. Imaging* **2018**, *52*, 305–309. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.