Detecting Seals in Active Sonar Images with A.I?

Benjamin Blundell, Doug Gillespie, Gordon Hastie, Emma Longden, Jess Montabaranom & Robert Harris.

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

Dataset

We have collected a large dataset of active sonar images over a period of 12 months. These images are processed by hand with the aid of a simple detection algorithm within the program PAMGuard[1], which draws coloured 'tracks' over areas where the intensity of the received signal is consistently higher than background. The majority of track segments detected are created by the movement of the turbine blades. These occur at an average rate of around one per second at slack water, rising to 10 per second during turbine operation.

Data on the movements of animals around the turbine were collected using a high-frequency multibeam sonar system (Tritech Gemini 720id: Tritech International Ltd, Westhill, Aberdeenshire, UK). This is a forward looking multibeam sonar which provides information on sonar targets in the X-Y plane; it has a temporal resolution of approximately 10 Hz, an angular resolution of 0.5°, and a range resolution of 0.8 cm. The horizontal swath width of the Gemini is 120° and the vertical beam is 20° (-3dB with a 10° downward tilt)[2]

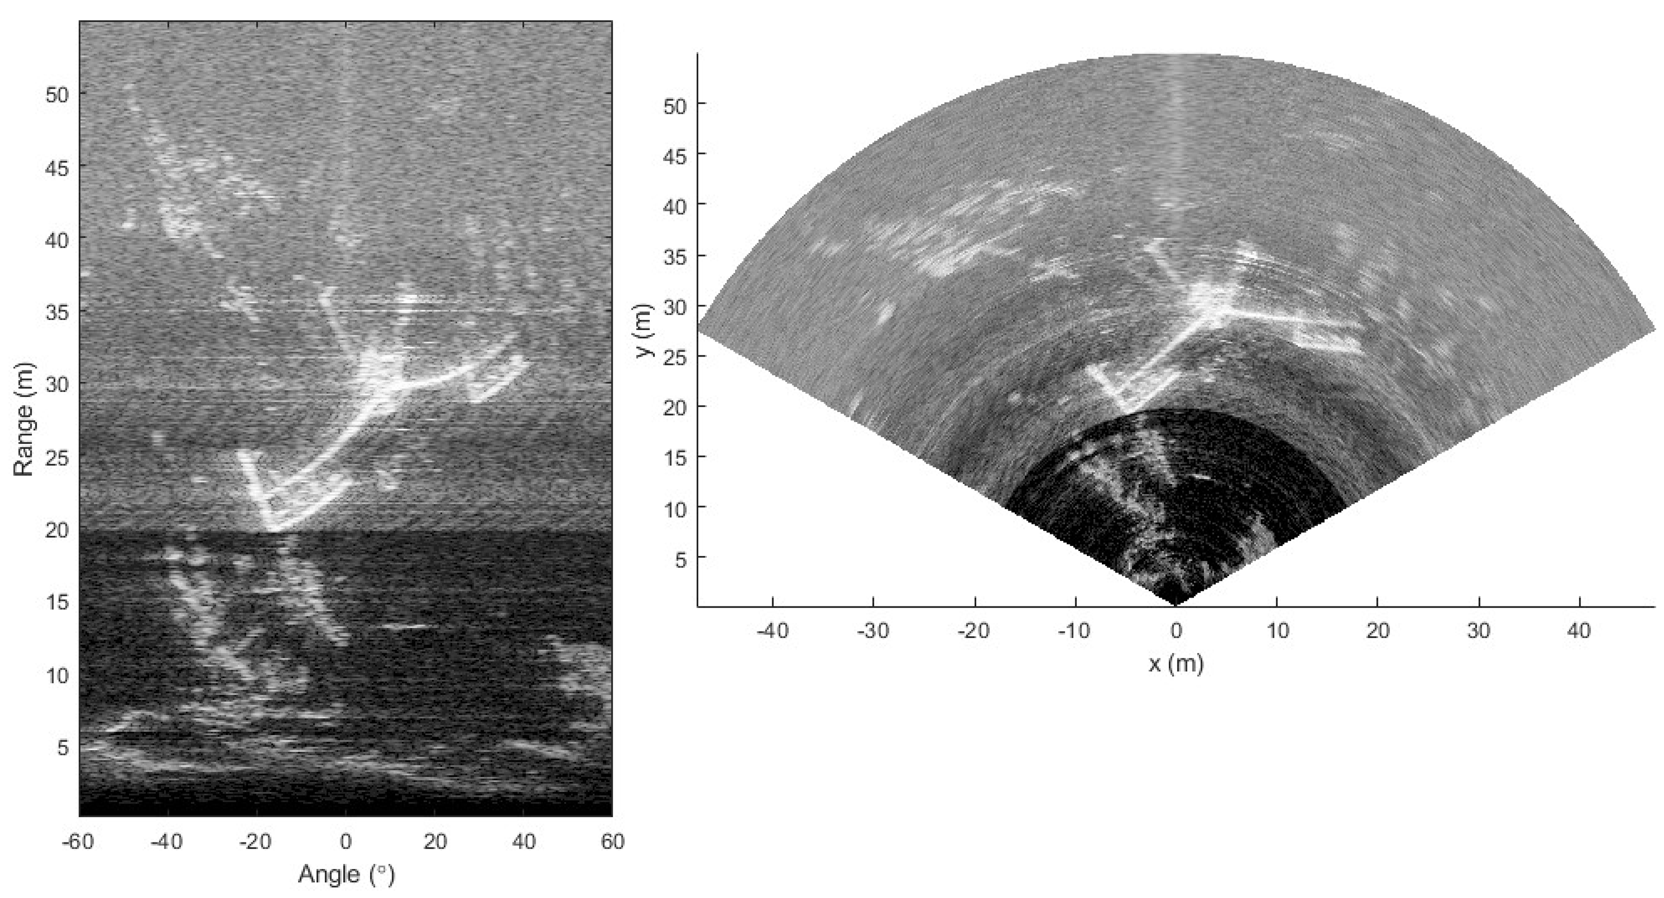

Raw data from a single sonar image, and the same data transformed into a geometrically correct fan image. The triangular tripod base of the turbine is clearly visible as is the central tower, slightly offset to the right of the image centre. This process is computationally expensive and can introduce artifacts if not handled correctly.

The dataset consists of large image files, PAMGuard saved files and SQLite3 databases. Marshalling these data is a considerable task. The following numbers highlight the frequency of certain events – a *'needle in a haystack'*.

Raw datasize:	> 96TB
Dataset Size on Disk:	204GB
Number of Images:	980, 662
Number of Tracks:	229,394
Number of events:	9,377
Mammal Events:	149
Fish Events:	2,737
Unknown Events:	2,411

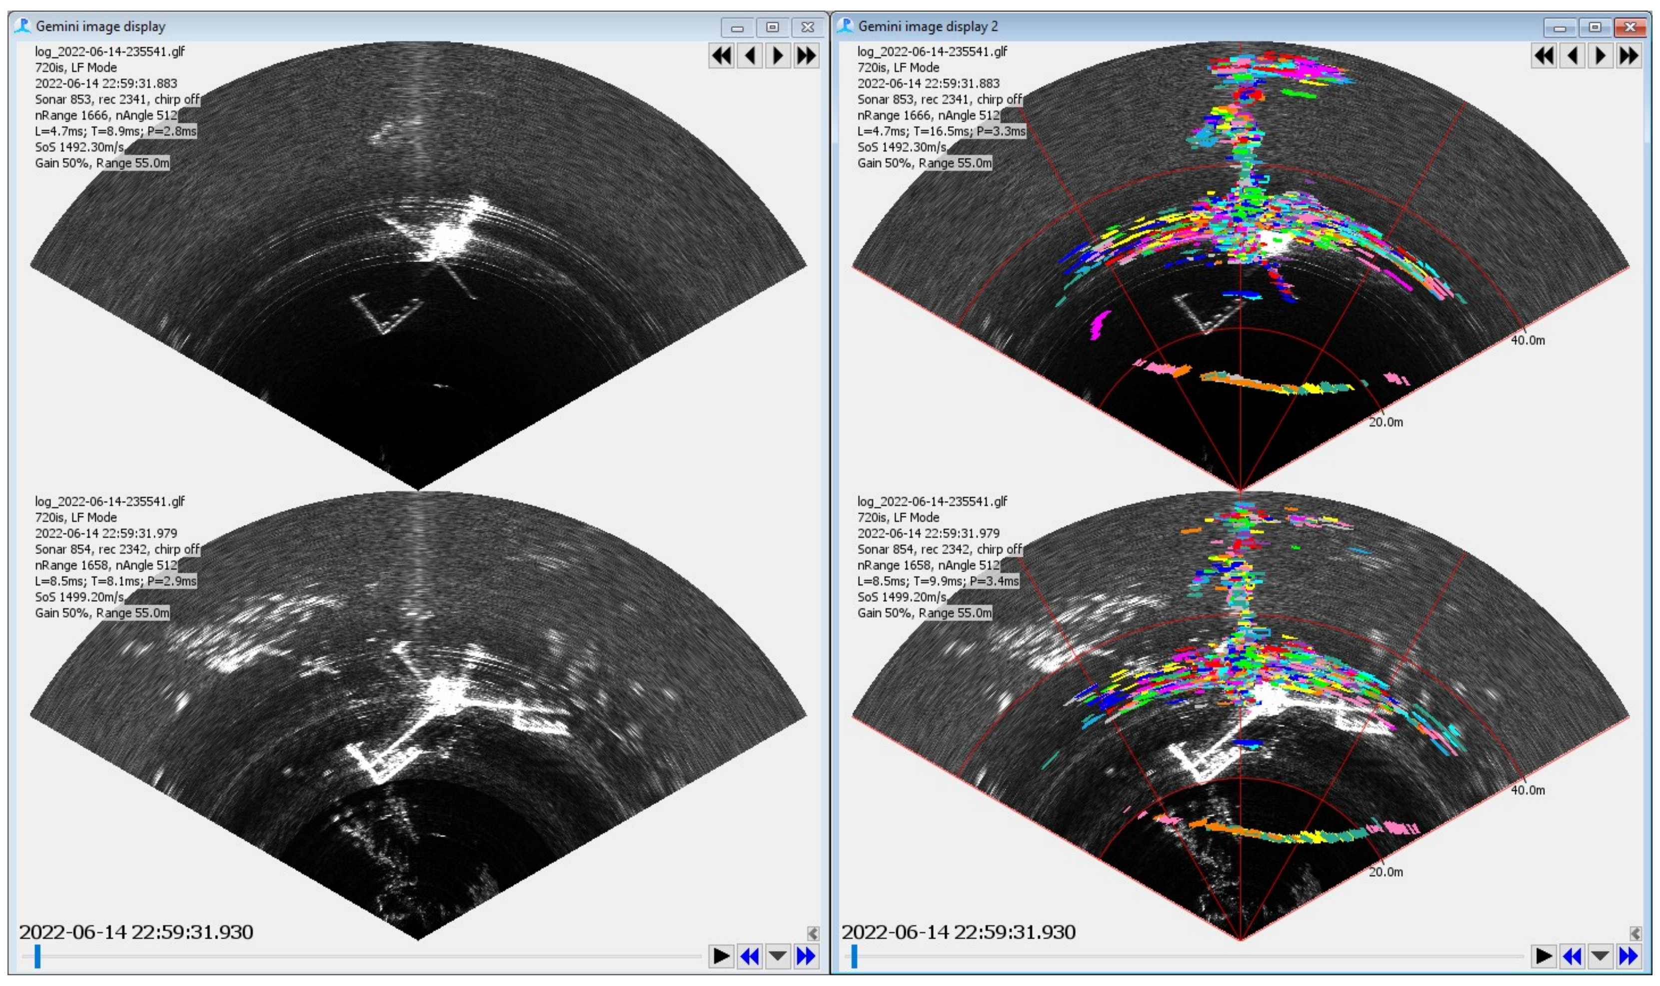

PAMGuard sonar display. The left figures show transformed images from the two sonars, and on the right are the same images overlayed with 15 minutes of detection data. Hovering the mouse over any detection in the track will show a small popup windows with a 3mx3m clip of that part of the raw sonar image corresponding to that detection, right clicking on the detections can move the scroll bar of the displays to exactly the start of the track, so that it can be viewed in the left display without the detection overlay. The display can be zoomed 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 meta data and the operators classification added to an underlying database.

Potential approaches

YOLO (You Only Look Once)[3] is one potential approach, though this algorithm only considers single frames; we believe that considering a time period will yield the best results.

Convolutional neural networks (CNNs) can be trained with an extra dimension, operating over volumes instead of 2D images.

Transformers have recently been applied to time-series data other than text[4].

Several issues need to be overcome when automating an identification process. Some of these are:

1) Clutter in the image from various sources, along with noise and effects that are not of interest such as the movement of debris on the sea bed.

2) The turbine itself causing many false detections.

3) Training a detector that can be applied to multiple sites easily.

4) Detection versus classification. Are there enough data to train a classifier?

5) Sparsity – individual detections are rare and take up a small proportion of the image.

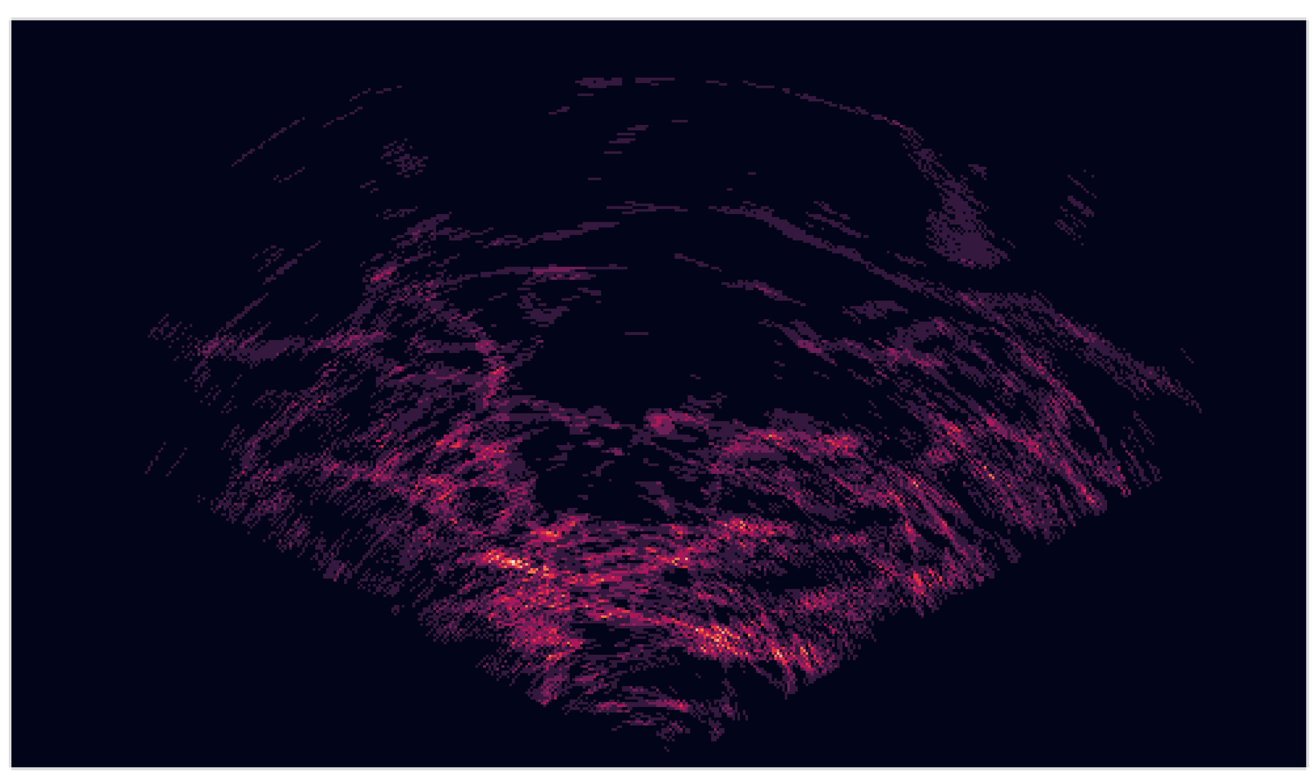

A visualisation of the Seal tracks collected to date. Brighter areas are frequented more often than darker areas.

Scottish Oceans Institute
Institiud Chuantan na h-Alba
University of St Andrews
Gatty Marine Laboratory
East Sands St Andrews
KY16 8LB
Scotland
UK

[1] https://www.pamguard.org/
[2] Parsons, M.J.G., Fenny, E., Lucke, K., Osterrieder, S., Jenkins, G., Saunders, B.J., Jepp, P. & Parnum, I.M. (2017) Imaging Marine Fauna with a Tritech Gemini 720i Sonar. *Acoustics Australia,* **45,** 41-49.
[3] https://pjreddie.com/publications/
[4] https://arxiv.org/abs/2202.07125