

MACHINE LEARNING FOR CLASSIFYING THE UNDERWATER ENVIRONMENT FROM TRANSMISSION LOSS DATA

M Donnelly	Systems Engineering & Assessment Ltd, Beckington, UK
M Tipping	University of Bath, Bath, UK
F Cowie	Airbus Defence and Space, Chippenham, UK
Ph Blondel	University of Bath, Bath, UK

1 BACKGROUND AND MEASURES OF PERFORMANCE

Future sonar operations will be increasingly autonomous, including uncrewed vehicles operating over long periods of time in remote areas with little or no mission planning/replanning input from expert operators. Machine Learning (ML) has the potential to support the autonomous repositioning of vehicles for optimising sonar operations. To investigate this potential, work was conducted to develop and test probabilistic ML algorithms for classifying the underwater environment using only acoustic Transmission Loss (TL) data. Environmental classification would then enable appropriate positioning strategies, for example adjusting spacing between distributed sonar systems in order to maintain detection coverage for environments with differing acoustic propagation characteristics. An incorrectly classified environment could lead to inefficient use of sensors (i.e. too closely spaced when the TL in the environment is actually low but is classified as high) or poor coverage (i.e. the reverse).

Measures of Performance (MoPs) are required in order to assess the ML classification algorithms. The standard metrics of classification accuracy and confusion matrix were used. In addition, a 'detection coverage' MoP was constructed in order to assess the practical impact of misclassification. Given a TL versus range curve, this was defined as the fraction of the ranges for which the TL is less than a fixed Figure of Merit (i.e. the maximum TL which can be tolerated for a given set of sonar, contact and background noise characteristics). It was used to construct a cost matrix which, when combined with the confusion matrix, enabled an overall performance metric to be calculated. The cost matrix was the same size as the confusion matrix, with zeros along the diagonal and the off-diagonal terms calculated from the absolute difference between the detection coverage MoP for the environment classes at the row and column indices (i.e. true versus that selected by the classifier).

2 ENVIRONMENTS AND ACOUSTIC MODELLING

The underwater environments were defined by a temperature and salinity profile, a seabed type (with associated density, compressional wave sound speed and attenuation values), and water depths (both flat and sloping seabed topographies were modelled). The oceanographic profiles and water depths were chosen to be representative of the UK North-West approaches' continental shelf. Temperature and salinity data were obtained from the Copernicus Marine Environmental Monitoring Service (CMEMS). The Atlantic-European North West Shelf Ocean Physics Analysis and Forecast dataset [1] was used, which provided hourly interval forecasts from the Nucleus for European Modelling of the Ocean model, produced by the UK Meteorological Office. Temperature and salinity profiles were downloaded for the location 57.5°N, 8.5°W. Sound speed was calculated from temperature and salinity using the UNESCO standard Chen and Millero equation [2]. Seabed type was taken from the US Naval Oceanographic Office Bottom Sediment Type (BST) version 2 database [3]. The High Frequency Environmental Acoustics dataset was used, which contains 23 descriptive seabed types; the corresponding geo-acoustic parameters were taken from Section IV Table 2 of [4].

Nine classes of underwater environment were defined, corresponding to all permutations of three types of sound speed profile (Upward Refracting, Downward Refracting and Stratified) and three types of seabed (Low Loss, Medium Loss and High Loss). Within each of these classes, multiple

variants were created, in order to generate training and test datasets for the ML algorithms. The temperature and salinity profiles used in calculating the sound speed profiles corresponded to a 48-hour period from midnight on the 14th to midnight on the 16th of January 2018 (Upward Refracting), July 2018 (Downward Refracting) and October 2018 (Stratified) – see Figure 1. The bathymetry profiles consisted of flat and sloping seabed variants with constant gradients for the sloping variants.

For the training data, 1200 variants were created, corresponding to all permutations of:

- 16 sound speed profiles, taken from ocean model forecasts every 3 hours over a 48-hour period starting at 0 hours, i.e. at times 0, 3, 6 ... 45 hours;
- 5 seabed types (similar types within the available BST types);
- 15 bathymetry profiles, both flat and sloping, with water depths between 150 m and 200 m.

For the test data 320 variants were created, corresponding to all permutations of:

- 8 sound speed profiles, taken from ocean model forecasts every 6 hours over a 48-hour period starting at 1 hour, i.e. at times 1, 7, 13 ... 43 hours;
- 4 seabed types, linearly interpolated at the midpoints between the five training variants;
- 10 bathymetry profiles, similar to but not the same as the training profiles.

Hence the test dataset is completely independent of the training dataset.




Figure 1 - Sound speed profiles: upward refracting (left), downward refracting (centre) and stratified (right). The blue curves are training data, the red curves test data. All profiles go to the maximum modelled water depth of 200m. The sound speed varies over depth by approximately 3m/s, 20m/s and 5m/s for the upward refracting, downward refracting and stratified cases respectively.

The Low Loss seabed variants were taken from BST indices 2 to 7 (Rock, Cobble or Gravel or Pebble, Sandy Gravel, Very Coarse Sand, Muddy Sandy Gravel), the medium loss variants from indices 9 to 13 (Medium Sand or Sand, Muddy Gravel, Fine Sand or Silty Sand, Muddy Sand, Very Fine Sand), and the high loss from 18 to 22 (Sandy Mud or Silt, Fine Silt or Clayey Silt, Sandy Clay, Very Fine Silt, Silty Clay). This range of BST types encompasses densities from ~ 1.1 to 2.5 g/cm^3 and sediment-to-water sound speed ratios of ~ 0.98 to 2.5 , i.e. a wide range of seabed reflectivity would be expected. The seabed was modelled as a single isotropic infinite layer.

Acoustic Transmission Loss (TL) was calculated using the PyRAM acoustic model [5], which is a Python adaptation of the Range-dependent Acoustic Model (RAM) [6]. Water volume attenuation and sea surface loss/scattering are not included in PyRAM (nor in the original RAM model). This is not considered a limitation at the low frequencies modelled here, e.g. at 1 kHz (the centre frequency used in the modelling) the water volume attenuation over 50 km equates to around 3 dB (using the Francois-Garrison formula), which is considerably lower than the TL variations between the environmental classes seen below. The acoustic modelling parameters were as follows:

- The source and receiver depth were both 40m, chosen so that the stratified profile has an acoustic effect (the profile changes gradient at 60 m – see the right-hand plot in Figure 1);
- The calculation frequencies were 900 Hz to 1100 Hz every 10 Hz, and TL was averaged over frequency in order to give a smoother, hence more realistic curve;
- TL was calculated every 100 m out to a range of 50 km.

The TL training data is shown in Figure 2. The span of minimum to maximum TL values at each range is shown for each of the 9 classes in separate plots. The FoM values of 75 dB, 80 dB and 90 dB used for the detection coverage MoP (see Section 1) are shown as dashed red lines. Although the High Loss / Downward class is an obvious outlier, there is considerable overlap between many of the other classes. At first glance, with the exception of the High Loss / Downward class, it would not seem a feasible task to classify the environment based upon these TL curves, particularly if only a small segment of the curve was available (which might be the case in practice due to a short vehicle track).

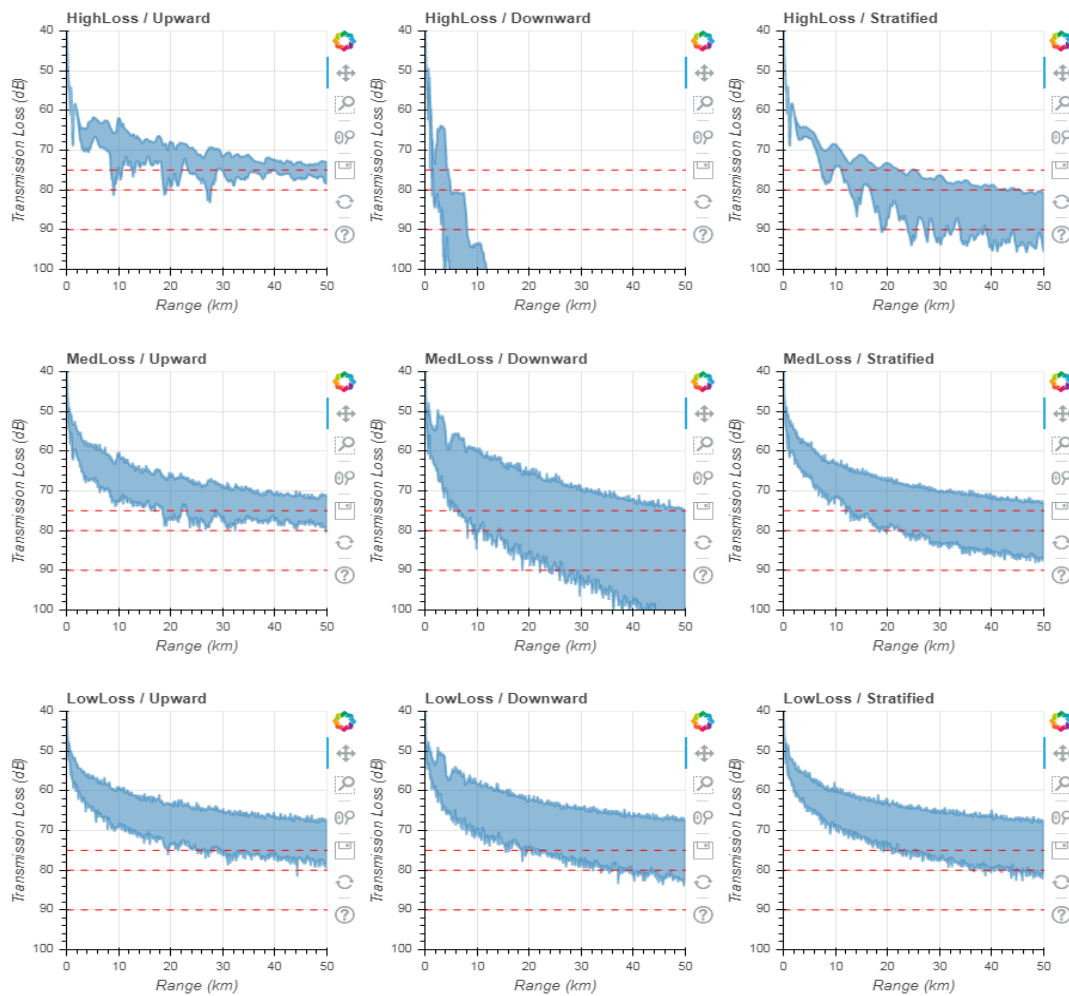


Figure 2 - TL training data for all nine classes, with FoM values of 75dB, 80dB and 90dB used for the detection coverage MoP shown as dashed red lines.




Figure 3 - Cost Matrix with FoM values of 90 dB (left), 80 dB (centre) and 75 dB (right).

Figure 3 shows the cost matrix referred to in Section 1 for each of the three different FoM values, produced from the TL data by calculating the detection coverage MoP for each of the 1200 training data variants and then averaging over the variants. The indices 1 to 9 follow the nine classes in row order as per Figure 2, i.e. index 1 is HighLoss-Upward, 2 is HighLoss-Downward, etc. The cost matrix clearly depends strongly on the selected FoM value (to be expected, as the FoM value determines the relevant portion of the TL curve). For the FoM=90 dB case, the only significant cost is misclassifying the HighLoss-Downward class as any other class. This is because only the HighLoss-Downward class has the majority of TL values greater than 90 dB (beyond about 10 km all TL values exceed 90 dB). As the FoM reduces, more structure becomes visible in the Cost Matrix. At 80 dB the highest costs are for misclassifying HighLoss-Downward as MedLoss-Upward or as any LowLoss classes, with a similar trend but more structure at 75 dB.

3 MACHINE LEARNING ALGORITHMS

The problem was cast as one of multi-dimensional classification with the number of dimensions given by the number of ranges at which TL data was available (500). Each class in the training data set comprises 1200 samples, and each sample is therefore a point in 500-dimensional space. Beyond direct examination of the TL curves (Figure 2), to visualise how the data is distributed both within and across classes, a number of methods were applied for projecting the data down to two dimensions where relevant structure might be visually inferred (although information is inevitably lost during the dimension-reduction process). The examples shown in Figure 4 are principal component projections - the orthogonal linear projection (down to two dimensions in this case) that retains maximum variance. As this is a linear projection, groups of points (e.g. representing classes) that are separable in the projection must likewise be separable in the data space. However, separable classes in high-dimensional space may overlap in the projection.




Figure 4 - PCA projection of all classes (left) and of the upward-refracting profile classes (right).

The linear projections in Figure 4 show separation of the data points both across and within classes (plots were also produced for the downward-refracting and stratified profile classes which also showed separability). Due to the implied separability in the data space, this suggests at least the feasibility of classification.

The key question for a ML approach to classifying the TL curves was in respect of which direction the data was modelled. There are two options: as a mapping from curves to class label (discriminative modelling) or from class label to curves (generative modelling). Generative models seek to approximate the “causal” (“forward”) direction of the observed data, while classification is effectively an “inverse” modelling process. The underlying motivation for generative modelling is that the forward mapping should be smoother than in the inverse direction, and so should be easier to capture in any model (and so, potentially requiring less data). A class-specific generative model would map from label to observed data; classification is then done indirectly, by measuring and comparing the probability those various generative models assign to new data.

A generative modelling approach was chosen for three reasons:

- Given the complexity and dimensionality of the problem (1200 training points vs 500 dimensions), we did not have sufficient data (by at least one order of magnitude) to be able to apply many of the more effective discriminative techniques, especially for the ubiquitously popular deep neural network class of models.
- A key feature of the problem is that, in a practical scenario, we will only have access to a limited number of points on the TL curve, and our models need to be able to make predictions in that scenario. For discriminative models, this is a serious challenge. Conversely, given an appropriate choice of generative model, it is in principle straightforward to compute the probability over a subset of points (a statistical calculation known as “marginalisation”).
- It is possible to detect novel or unusual data that has not been seen before (the classifier system would have the capacity to say “don’t know”).

To perform classification in a generative modelling context, we first develop, based on the training set, a class-conditional probabilistic model of each of the nine classes, $P_k(\mathbf{t}|k)$, with $k = 1 \dots 9$. Then, given new data (a TL curve, or a subset of it), we may compute the predictive probability $P_k(\mathbf{t}|k)$ given by each of the nine class-conditional models. This may be trivially converted (using Bayes’ rule) into a *a posteriori* probability of belonging to each class, $P(k|\mathbf{t})$, and this can be very readily used for classification (i.e. to assign the curve to the class giving the highest posterior probability).

The question is then what form the generative models $P_k(\mathbf{t}|k)$ should take. There are a wide range of options, of which three were identified as follows (in increasing order of complexity):

1. “Probabilistic PCA” (PPCA) [7]: This is a multivariate Gaussian model, applied directly in the 500-dimensional signal space, using a principal component parameterisation to constrain complexity. A persuasive feature of this “Gaussian process” class of approach is that it is straightforward to compute the marginalisation required to make predictions for limited numbers of observation points on the curve.
2. “Latent Variable Gaussian Process” (LVGP). Developed specifically for the study, this is an alternative form of Gaussian process that is defined implicitly in terms of a weighted expansion of fixed basis functions. It has similar expressive power to the PPCA approach, and marginalisation is again straightforward, but this model has the additional benefit that it can be evaluated at arbitrary ranges and resolutions.
3. “Generative Neural Network” (GNN). This may be considered a non-linear extension of the above LVGP approach, based on adaptive basis functions. The non-linearity of this approach in principle offers greater flexibility in signal representation, despite the considerable computational cost.

Regardless of model selected, the high dimensionality of the data (and the comparatively few training examples available for each class) leads to a danger of over-fitting during model development. It was therefore necessary to carefully tune model complexity based on the training data, and for this purpose a standard cross-validation procedure was adopted.

The primary focus of the work was on implementing and testing the PPCA model. The LVGP model was also implemented, with promising preliminary results. The GNN model proved computationally problematic so was not pursued (in part due to the success of the other algorithms – see Section 4). The starting point for implementing the class-conditional models $P_k(\mathbf{t}|k)$ was the multivariate Gaussian model given by Equation (1). The multivariate Gaussian model is parameterised by its mean vector \mathbf{m}_k and covariance matrix \mathbf{C}_k . Note that these parameters are specific to each of the classes: it will be necessary to fit nine such models, one to each of the nine 1200-sample training data sets.

$$P_k(\mathbf{t}|k) \propto \exp \left\{ -\frac{(\mathbf{t} - \mathbf{m}_k)^T \mathbf{C}_k^{-1} (\mathbf{t} - \mathbf{m}_k)}{2} \right\} \quad (1)$$

A key motivation for adopting a generative modelling approach was the requirement to be able to apply models to a subset of the data, i.e. a subset of elements of the vector \mathbf{t} . This requires integrating out the unobserved sample points from the joint distribution defined above so as to compute the marginal distribution over the observed samples only. For a multivariate Gaussian, this marginal computation is straightforward – if we observe a subset of points on the TL curve $\hat{\mathbf{t}}$, then the class-

conditional marginal is simply a modified Gaussian distribution of the same form as Equation 1 with $\hat{\mathbf{t}}$ replacing \mathbf{t} , $\hat{\mathbf{m}}_k$ replacing \mathbf{m}_k as the subset of the mean vector corresponding to the observed points, and $\hat{\mathbf{C}}_k$ replacing \mathbf{C}_k as the covariance matrix realised by retaining only the rows and columns corresponding to the observed points. As such, any Gaussian model derived from the full data may be readily applied to any subset of it. However, the covariance \mathbf{C} that must be estimated for each full class-conditional model is a 500×500 matrix, which equates to a huge number of parameters to be specified in the context of 1200 data samples. In order to avoid over-fitting, it was essential to adopt a more concise parametric representation. A basic approach would be to utilise a simple isotropic model ($\mathbf{C} = \sigma^2 \mathbf{I}$, one parameter), or more flexibly, a diagonal covariance (500 parameters). An even more flexible option is the PPCA parameterisation ($500 \cdot q + 1$ parameters). This allows the model complexity to be tuned through adjustment of the parameter q , the approach adopted here. The PPCA parameterisation controls complexity of the covariance model via a restricted-rank decomposition:

$$\mathbf{C} = \mathbf{W}\mathbf{W}^T + \sigma^2 \mathbf{I} \quad (2)$$

where \mathbf{W} is a $500 \times q$ matrix, with q free to be chosen. This captures the covariance structure in a q -dimensional subspace precisely, while very roughly approximating the remainder (in the $500 - q$ dimensional orthogonal subspace) by an isotropic (spherical) model with scalar variance σ^2 in all directions. The model was fitted to the data following the standard maximum likelihood procedure, and a convenient feature of the PPCA parameterisation is that there is an analytic solution: the columns of \mathbf{W} are proportional to the eigenvectors of the measured sample covariance matrix. In terms of “tuning” the PPCA model, it is necessary to set the q parameter appropriately. This was achieved using 25-fold cross-validation. The process was to vary q over a chosen range, build the nine class-conditional models with the given q , and assess each value by measuring the classification error on the held-out data. In this case, a single shared value of q is assumed for all class-conditional models. In principle, we would expect to improve the overall classification accuracy by allowing q to be tuned per class, but as this is computationally expensive (the search space grows exponentially with the number of models), optimisation of multiple q values was not undertaken. Note that $q = 0$, giving $\mathbf{C} = \sigma^2 \mathbf{I}$, is the simplest (isotropic covariance) Gaussian model as noted above. Implicitly, new data will be classified according to the nearest class mean \mathbf{m}_k . With $q = 2$ we have a slightly more complex model, loosely analogous to working in the space of the previous PCA visualisations. The outcome of the process described above was that the lowest mis-classification rate (2.34%) was given for $q = 9$, hence this represented the best model, based upon the training set.

4 RESULTS

The confusion matrices for the PPCA model (with $q = 9$) are shown in Figure 5 for both the training dataset and the independent test dataset. The error rates are shown above the matrices, i.e. the classification accuracy was 97.66% for the training dataset and 93.19% for the test dataset.




Figure 5 - Confusion matrix for the $q=9$ model against the training data (left) and the test data (right)

Given the observations made on the TL data in Section 2, the performance shown in Figure 5 significantly exceeded expectations. The (limited) evaluation of the LVGP model showed an even higher accuracy of 96.22% (using 66 selected basis functions). However, this was a preliminary result as the time was not available to conduct as robust an evaluation process as for the PPCA algorithm.

Having trained the classifier and proved its accuracy against the independent test dataset, the remaining questions were on the performance with only a subset of TL data made available, and the practical impact of misclassification. A number of simulations were performed where the classifier was presented with sections of TL curves representative of data collection over practically reasonable timescales (~45 and 90 minutes) at a vehicle speed of 2 knots, i.e. 2.5 and 5km sections respectively. Simulations included examination of a single variant (TL curve) with fixed and variable starting ranges and successively more data points, and averaging results over multiple variants and starting ranges to give representative average performance, as well as calculating the overall performance metric (combining confusion and cost as described in Section 1). Figure 6 portrays how class probability for a single test variant within three classes varies over range from a given starting range. The starting range was different in each case.




Figure 6 - Class probability averaged over all test variants for multiple starting ranges

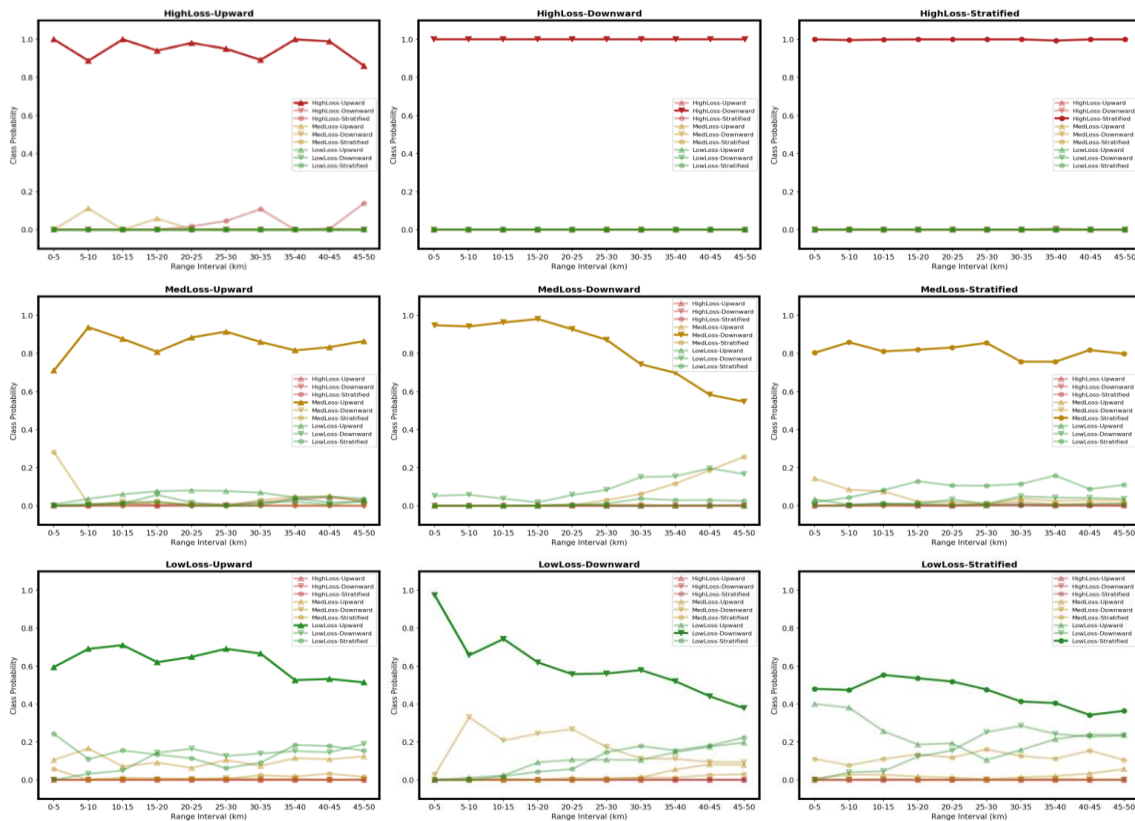


Figure 7 - Class probability averaged over all test variants for multiple starting ranges

In Figure 6, test variant number 99 was arbitrarily selected from the test data for the three classes LowLoss-Downward, MedLoss-Stratified and LowLoss-Upward. In each case the bold curve corresponds to the true class. The LowLoss-Downward and MedLoss-Stratified show good classifier performance over a 5 km range interval, which becomes more certain of the correct classification as more data is gathered over time. The LowLoss-Upward plot is an example of a case where the model fails to classify correctly with a high level of certainty.

Figure 7 shows the class probabilities averaged over all test variants for multiple starting ranges, using intervals of 5 km. The accuracy for all High Loss classes is immediately evident, in particular for the Downward and Stratified classes which show close to 100% certainty across all starting ranges. Accuracy for the Medium Loss classes is also encouraging, with Upward and Stratified profile types specifically having a relatively constant and high probability of correct classification across all starting ranges. Accuracy for the Low Loss classes is generally lower. However, it is worth noting that the correct class is always assigned the highest probability even though the absolute probabilities are lower than for the Medium Loss and High Loss classes. The averaged confusion matrices and performance metrics were also evaluated during the simulations. The performance metric was always above 95%, indicating that even when misclassification occurred, the cost was low.

5 CONCLUSIONS

The principal classification model investigated (PPCA) showed very good performance against the test data, with a classification accuracy of over 93%. When presented with short intervals of data (2.5 and 5km), the classifier still performed well, with performance metrics (taking into account the practical impact of misclassification) over 95%. Preliminary results from the LVGP model were also suggestive of good performance. There may be scope for further improvement of the PPCA algorithm with additional training data. The classification accuracy exceeded expectations given the complexity of the data, and there is potential for exploitation across a number of marine acoustic application areas.

Acknowledgments

The work reported in this paper was funded by the Defence and Security Accelerator¹ (DASA) under DSTL contract DSTLX1000127682 and was conducted by SEA and the University of Bath.

6 REFERENCES

1. COPERNICUS, "Atlantic - European North West Shelf - Ocean Physics Analysis and Forecast" EU Copernicus Marine Service Information; doi 10.48670/moi-00054
2. Chen and Millero, UNESCO, http://resource.npl.co.uk/acoustics/techguides/soundseawater/underlying-phys.html#up_unesco
3. *A bottom-sediments province database and derived products*, Naval Oceanographic Office, DOI 10.1121/1.3654538
4. *APL-UW High-Frequency Ocean Environmental Acoustic Models Handbook*, APL-UW Technical Report 9407, Applied Physical Laboratory, U. Washington, October 1994
5. PyRAM, "Python adaptation of the Range-dependent Acoustic Model (RAM v1.5)", <https://github.com/marcuskd/pyram>
6. M D Collins, *A split-step Padé solution for the parabolic equation method*, J. Acoust. Soc. Am. 93, 1736–1742 (1993)
7. Michael E. Tipping, Christopher M. Bishop, *Probabilistic Principal Component Analysis*, Journal of the Royal Statistical Society Series B: Statistical Methodology, Volume 61, Issue 3, September 1999, Pages 611–622, DOI 10.1111/1467-9868.00196

¹ DASA finds and funds exploitable innovation to support UK defence and security quickly and effectively, and support UK prosperity. DASA's vision is for the UK to maintain its strategic advantage over its adversaries through the most innovative defence and security capabilities in the world. DASA is a cross-Government organisation, launched in December 2016 by the Secretary of State for Defence.