Enhancing Model Based Acoustic Localisation Using Quantum Annealing

Robert Wezeman[1], Tariq Bontekoe[2], Sander von Benda-Beckmann[1], Frank Phillipson[1,3]

[1]The Netherlands Organisation for Applied Scientific Research
Anna van Buerenplein 1, 2595 DA, Den Haag, The Netherlands

[2]University of Groningen
Broerstraat 5, 9712 CP, Groningen, THE NETHERLANDS

[3]Maastricht University
Minderbroedersberg 4-6, 6211 LK, Maastricht, THE NETHERLANDS

Robert.wezeman@tno.nl

ABSTRACT

In model based acoustic localisation (MBAL) the locations of underwater objects are estimated by comparing sensor measurements with model predictions. To obtain high quality predictions, computationally demanding propagation models need to be run for a large set of environmental parameters. The computational resources that are available onboard are typically not sufficient to perform accurate MBAL estimations on real-time data. In this work, we develop a quantum algorithm that uses quantum annealing to enhance underwater acoustic localisation.

A restricted Boltzmann machine (RBM) is trained to predict the probability distribution of the candidate location of an underwater target. Advantage of this approach is that part of the computation can be moved to offline-training. Moreover, once the model is trained, the probability distribution can be sampled more efficiently using a quantum annealer. Potentially, this could enable real-time accurate target estimations to be made onboard. The RBM is applied to a simplified multi-sensor horizontal localisation problem where we assume a constant and linear acoustic propagation. Using simulated annealing we show that the RBM is able to learn probability distributions that resemble target locations. First results show that the training and sampling can be done using quantum annealing hardware by D-Wave Systems for a limited size example.

Our contribution is the first work that explores how quantum algorithms can be applied for more efficient information processing for acoustic underwater localisation.

1.0 INTRODUCTION

Underwater sensing is a key capability for navies. Various platforms, such as frigates, submarines, autonomous vehicles and underwater networks rely on passive acoustic monitoring to detect threats, or to navigate safely underwater. Passive localisation of objects in an underwater environment is a much more challenging task when compared to localisation on land [1]. Electromagnetic waves that are typically being used in air cannot be used in underwater settings as these waves dampen out too fast. Instead, acoustic waves are being used for underwater sensing. Underwater propagation of acoustic waves is far from trivial due to reflections off the bottom or surface and variable sound speed [2]. The exact propagation depends on many parameters such as depth, temperature, pressure and salinity of the underwater environment [3].

Model based acoustic localisation (MBAL) is the technique of localising an object in underwater environments by comparing acoustic sensor measurements to model output predictions of different candidate locations. Measurement errors, as well as uncertainties and variability in environmental parameters, for example the sound speed profile, can result in uncertain or biased locations [5]. To quantify these uncertainties one would like to run these acoustic propagation models a large number of times for a large set of different environmental parameters to then get an indication of the likelihood of different possible locations. Ideally, these calculations, including sampling from different optimal environmental parameters, should be performed on real-time data onboard of an underwater receiver.

Various approaches have been developed to efficiently solve the model-based localisation problem with recent approaches investigating machine learning techniques such as convolutional neural networks [6] and Boltzmann machines [7]. A review on machine learning approaches applied to the domain of underwater acoustics, including acoustic localisation, is given by Bianco et al. [2].

At present, quantum computing, is still in its early stage of development and has not yet attained the size and quality necessary to address real-world problems of significant complexity. Nonetheless, it is anticipated that as quantum computers continue to mature, they will gain more traction and will be able to deliver benefits in various domains. In [8], an overview of quantum computing in the specific area of radar and sonar information processing is presented, including model based acoustic localisation as one of the most promising use cases.

In this work a restricted Boltzmann machine (RBM) [9], a stochastic graphical network, is trained to predict the probability distribution of the candidate location of an underwater target. Advantage of this approach for acoustic localisation is that part of the computation can be moved to offline-training. Moreover, with an efficient sampling strategy, this potentially could enable real-time accurate target estimations being made onboard once the model is trained.

One of the known disadvantages of RBMs is that evaluating their state, and hence training, is a time consuming task. Evaluating an RBM requires drawing samples from a Boltzmann distribution, typically done using Gibbs sampling which requires long equilibration time [9]. Adiabatic quantum computing (AQC) was proposed to be a more efficient way of sampling from Boltzmann distributions [10]. Numerous papers have since explored the possibility of using AQC for training RBMs with various applications. For a more comprehensive examination of other research conducted on quantum-RBM, we refer to Dixit et al. [11].

As a first step towards the quantum application of acoustic localisation, an RBM is applied to a simplified multi-sensor horizontal localisation problem, where we assume a constant sound speed profile resulting in straight paths. Simulated annealing performance is studied extensively as it offers a natural starting point to further pursue a quantum annealing RBM, for which preliminary results are shown.

2.0 PROBLEM FORMULATION

2.1 Assumptions

The full MBAL problem, as described in the introduction, depends on a large number of environmental parameters and is therefore a far too complex problem as a starting point for this novel solution approach. Instead, we make the following two assumptions to simplify the problem. First. we assume that acoustic waves propagate in the (x,y)-plane. Hence ignoring reflections with surface and bottom. Next, we assume that acoustic waves propagate linearly with a constant velocity.

Given these assumptions, consider the following problem scenario. We look at a submarine that is located at the origin (x, y) = (0, 0). A target, that emits acoustic waves in the direction of the submarine, is located at (x_{source}, y_{source}). The submarine has 3 sensors in a line formation, each located at a distance 45m apart from its neighbour(s), at the locations (0, 45), (0, 0) and (0, -45). Sensors measure bearing angle and time of arrival of incoming acoustic waves. For each pair of sensors we can compute the time difference of arrival (TDOA).

The localisation problem consists of locating the target based on the sensor measurements: three bearing angles, and two independent TDOA values. At this stage we assume perfect measurements.

The values that can be represented on a quantum computer are limited by the available quantum resources (number of qubits). Hence the input and output parameter ranges need to be discretised into bins. When the available quantum resources increase over the course of time, more and finer bins can be used, leading to higher resolution predictions. The range for the angular coordinate θ is divided in equidistant sized bins. The radial coordinate is divided into exponential increasing bins, to obtain a higher resolution for sources that are nearby compared to sources that are still far away. If only candidate locations within a certain search region are considered it is possible to obtain a higher resolution, i.e., a larger number of grid points per area.

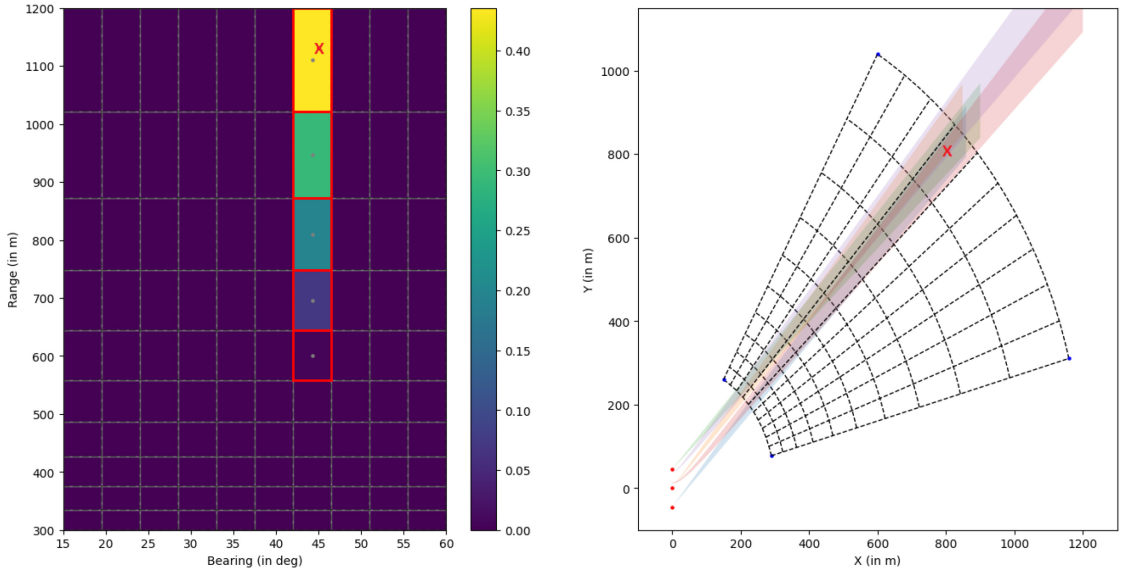

Figure 1: The (bearing, range)-grid and its mapping to the (x,y)-plane. The three sensors on the submarine are shown with red dots and a target at (x,y)=(800m, 800m) is indicated with a red cross. Colours in the left figure indicate the probability that an output bin corresponds to the given input shown on the right.

Our aim is to train a model, that is able to predict the output bin containing the target given the measured bearing and TDOA bin values. This is however not a straightforward one-to-one mapping. As can be seen from Figure 1, the overlap of the five input bins spans multiple output bins. So instead of predicting a single output bin, the model is trained to predict a probability distribution over multiple output bins.

2.2 Restricted Boltzmann Machine

An RBM is a graphical model consisting of two types of units: visible and hidden. The visible units represent the information from the environment, the hidden units are features that are to be estimated. All nodes of Boltzmann machine are binary valued and stochastic. Undirected edges exist between visible and hidden units and have trainable weights.

To solve the localisation problem a supervised RBM is trained to learn the probability distributions for different target locations. The model is trained and sampled using two different methods: (1) simulated annealing; or (2) quantum annealing, similar to Neumann et al.[12].

3.0 RESULTS

In Figure 2, the probability distribution, as predicted by the RBM is shown, and compared with the true probability distribution for a given target. The model predicts the angular distribution particularly well. The accuracy in the radial direction could be improved, which is a common challenge for acoustic localisation methods.

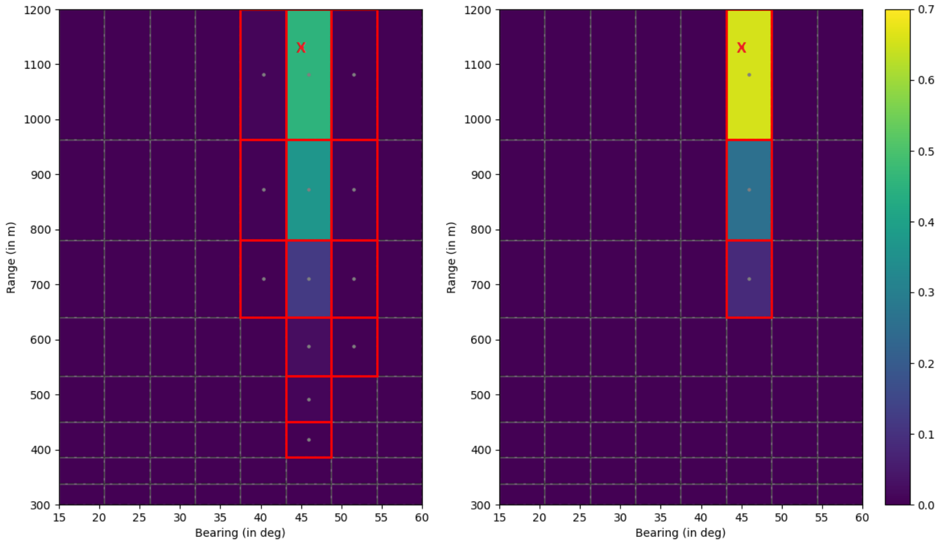

Figure 2: The found probability distribution by the model (left) compared to the true probability distribution (right). The model has 75 hidden units and is trained for 5000 epochs using simulated annealing. Bins that have nonzero probability are indicated with red.

As a measure of quality for the RBM, three metrics are defined. (1) **Distribution distance:** The squared difference between the model probability distribution and the true probability. (2) **Valid output:** The percentage of output samples that can be used for localisation. (3) **Wasserstein distance:** The amount of work that is required to transform the model probability distribution into the true probability distribution.

Each metric is computed as an average over 1000 different target locations as shown in Figure 3. The quantum model is only trained for 700 epochs due to limited available resources. The size of the considered model was chosen to be close to the maximum embeddable size on current D-Wave hardware. All three metrics, for both training methods, show that that the model is learning in the direction of the desired probability distribution. It can be shown that the error that remains is almost purely in the radial direction.

4.0 CONCLUSIONS

Our study explores how quantum algorithms can be applied for more efficient information processing for acoustic underwater localisation. Our results demonstrates some success in training an RBM model for a simplified localisation problem. The main benefit of using an RBM for acoustic localisation is that the training of the model can be done offline and only the sampling has to be done onboard. Quantum annealing is proposed to enable more efficient training, but also able to improve during the sampling stage.

Using simulated annealing we show that the RBM is able to learn the angular dependence of the probability distribution particularly well. Challenges remain in the radial dependence, which is a known and common issue within the field of acoustic localisation, and is driven by the limited resolution imposed by the available quantum resources. As quantum hardware size continues to advance, output resolution can be improved by increasing the number of TDOA bins or by incorporating additional hidden layers into the RBM.

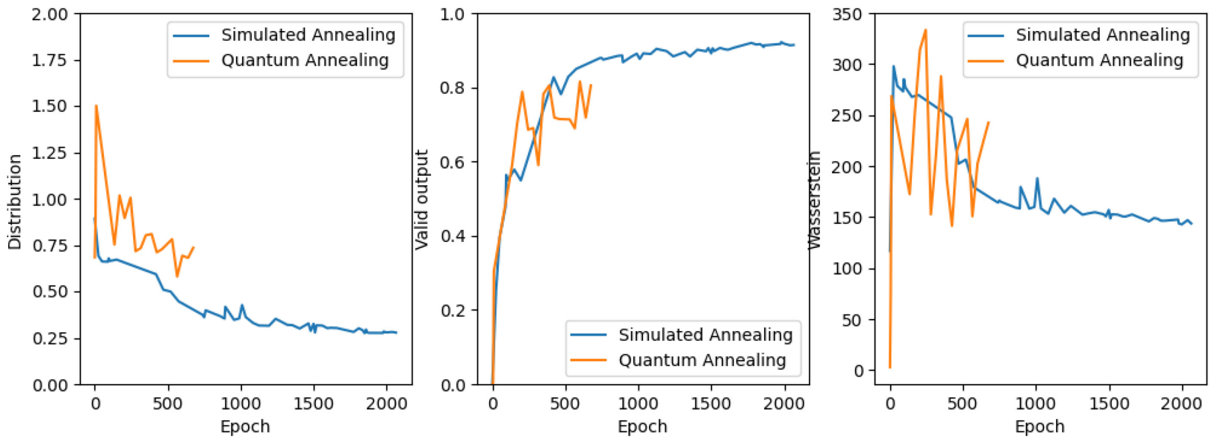

Figure 3: Performance metrics for quantum- and simulated annealing trained RBM.

The total number of training steps that were performed using quantum annealing is not sufficient to draw definitive conclusions regarding possible quantum improvements during training or sampling stage. More extensive research is needed to demonstrate the potential benefits of using quantum annealing for training and sampling an RBM. Missing in this study is a comparison with an RBM that is trained with conventional techniques. Adding such comparison will be particularly interesting when more quantum annealing training steps are performed.

5.0 REFERENCES

[1] Dullaart B. (2020), Model-based localization using vertical line arrays. Master's thesis, TU Delft, The Netherlands

[2] Bianco, M. J., Gerstoft, P., Traer, J., Ozanich, E., Roch, M. A., Gannot, S., & Deledalle, C.-A. (2019). Machine learning in acoustics: Theory and applications, The Journal of the Acoustical Society of America

[3] Abraham, D.A. (2019), Underwater acoustic signal processing: modelling, detection and estimation. Modern acoustics and signal processing. Springer

[4] Porter, M.B. (2011), The BELLHOP Manual and User's Guide: PRELIMINARY DRAFT

[5] Tolstoy, A. (1989), Sensitivity of matched field processing to sound-speed profile mismatch for vertical arrays in deep water pacific environment. The Journal of the Acoustical Society of America

[6] Liu, W., Yang, Y., Xu, M., Lü, L., Liu, Z., & Shi, Y. (2020), Source localization in the deep ocean using a convolutional neural network, The Journal of the Acoustical Society of America

[7] Xinwei, L. & Feng, Y. (2020). An underwater acoustic target recognition method based on restricted Boltzmann machine, Sensors, Switzerland

[8] Bontekoe, T.H., Neumann, N. M. P., Phillipson, F., & Wezeman, R. S. (2022), Quantum computing for radar and sonar information processing, Quantum Information Science, Sensing and Computation XIV, SPIE

[9] Hinton, G.E. (2012). A Practical Guide to Training Restricted Boltzmann Machines, Springer

[10] Adachi S. H. & Henderson, M. P. (2015). Application of Quantum Annealing to Training of Deep Neural Network, arXiv

[11] Dixit, V., Selvarajan, R., Alam, M. A., Humble, T. S., & Kais, S. (2021). Training restricted Boltzmann machines with a d-wave quantum annealer, Frontiers in Physics

[12] Neumann, N. M. P., Heer, P. B. U. L. & Phillipson, F. (2023). Quantum reinforcement learning, Quantum Information Processing