Sonar Sensitivity Analysis Using a Neural Network Acoustic Model Emulator

Megan U. Hazen
Warren L.J. Fox
Chris J. Eggen
University of Washington, Applied Physics Laboratory
1013 NE 40th St., Seattle, WA 98105
Email: megan@apl.washington.edu

Robert J. Marks II
Mohamed A. El-Sharkawi
University of Washington, Dept. of Electrical Engineering
Box 352500, Seattle, WA 98105

***Abstract-* A technique is reviewed for training artificial neural networks to emulate the complicated input-output relationships of an acoustic model. This neural network acoustic model emulator is intended for use in a sonar controller, which may require a large number of forward model runs to determine the optimal sonar setting in a given environment. The neural network can supply sonar performance predictions to high enough fidelity for use in a controller, but with a much reduced computational burden compared to the original acoustic model. Among the challenges of developing control guidelines for highly variable littoral areas is the difficulty in understanding the sensitivity of acoustic response to small changes in environmental or sonar control parameters. An effective sensitivity analysis tool would allow users or automatic control algorithms to place a control emphasis on those parameters that have the greatest effect on sonar response. Additionally, an improved understanding of acoustic sensitivity may lead to improvements in model and controller development. In this paper, the neural networks, originally developed for automatic sonar controllers, are used to explore the sensitivity of the system. Given a properly trained neural network, sensitivity measures can be directly calculated. The neural networks can also be used to visualize the effect of changing environmental and control parameters. A variety of ways in which the neural network structures can be used to examine the sensitivity of the sonar system will be presented.**

I. INTRODUCTION

Naval sonar systems continue to evolve and become more capable, while at the same time becoming more complex to operate. With the emergence of littoral areas as the prime regions of interest, characterized by underwater acoustic environments that change quickly in both the temporal and spatial domains, the combination of environmental and sonar system complexities has placed an increasing burden on sonar operators. In order to aid operators, it has become necessary to develop automatic sonar controllers that can automatically optimize sonar line-ups for a specified environment.

The general design of such an automatic sonar controller is an optimization loop that requires large numbers of evaluations of sonar behavior. Traditionally, these evaluations are done using a physically based sonar model. Traditional models offer high fidelity assessments, but require large amounts of time for each evaluation, and are impractical for effective controllers. It is desirable, therefore, to develop a model that can be evaluated very quickly, but which approaches the fidelity of a traditional model. We have developed a surrogate model using artificial feed-forward neural networks (ANNs). This model responds in a small percentage of the time required for a traditional model, and provides the user with data that is sufficiently accurate for use in an automatic controller.

In addition to the basic model evaluations, it can be beneficial to be able to determine the sensitivity of the acoustic system. This type of analysis has many varied applications, which will be discussed in this paper. Currently, sensitivity analysis can be done by taking analytical derivatives on a physical model, or by evaluating the model many times to plot parametric curves. However, it may be impossible to take analytical derivatives of complex non-linear models, and it is computationally expensive to compute multiple model runs. We propose extending our current ANNs to compute the sensitivity in a variety of ways.

In this paper we will present an overview of the neural network development and performance. After a brief discussion of previous work, we will describe a variety of applications for sensitivity analysis, along with potential methods for calculating the requisite values.

II. NEURAL NETWORK ACOUSTIC MODEL EMULATION

A neural network acoustic model emulator was developed to meet the goal of speeding up acoustic model evaluations while maintaining sufficient accuracy for a controller to determine optimal sonar settings. This contribution is important because the controller architecture is an optimization loop whose objective function depends on acoustic model evaluations. This requires many model runs for a given environment in order to assess which sonar mode is best. The traditional physically based models are too computationally intensive to allow a controller to run in real time. A general discussion of this type of regression neural network may be found in [1].

The training begins by statistically characterizing a given operational area. Probability distribution functions are generated for the various environmental parameters that are required as input to a conventional acoustic model. Then, a large number of environmental and system parameter "realizations" is generated, on which the conventional acoustic model is run. We have used up to 40,000 model runs in this training ensemble.

A neural network architecture is then established where the model inputs are placed on the input layer of neurons, and the acoustic model output products are placed on the output layer, with typically three hidden layers. The neural networks then are trained using a version of the back-propagation algorithm. Our modifications to the back-propagation algorithm are detailed in [2,3].

The trained neural networks were then used to calculate the objective function for the controller optimization. The neural networks respond in a small fraction of the time necessary for the traditional acoustic model. In testing, ANNs emulated the physically based model with enough fidelity for the controller to choose appropriate sonar settings.

III. PREVIOUS SENSITIVITY WORK

There is little published work in the area of neural network sensitivity analysis, but some papers have proposed successful applications of the technology. Most commonly, the authors use perturbation analysis to look at the relative impact of different input parameters on the output values [4, 5]. This allows them to prune the input vector size, which increases the efficiency of ANN training and performance. Sometimes this technique also leads to improved fidelity of the output.

In one paper it is suggested that sensitivity information can be used to improve the back-propagation algorithm. Specifically, the algorithm is adjusted to take into account not only the error in the output value, but also the error in the sensitivity. The analysis of this work suggests that minimizing the latter error can significantly reduce the number of training vectors required to cover the model space, thereby reducing training time significantly. This work is detailed in [6].

IV. SENSITIVITY APPLICATIONS

A. *Neural Network Development*

It has been shown that sensitivity information can depict the relative impact that different input parameters have on neural network output. By pruning the input vector, that is, removing those parameters that have the smallest effect on the output, we can decrease training and evaluation time. In some cases, reducing the size of the input vector can also provide higher fidelity output of the neural network (this is assumed to be because extraneous input values are not correlated with the outputs and can 'confuse' the ANNs).

We have done some work with perturbation analysis, which can be applied to optimizing the structure of the neural network input vector [2]. We begin this analysis by choosing an operating point. We then vary each input parameter in turn, holding all other input parameters constant. The input is modified by –2.5% and 2.5% of its entire range.

A sample output of this perturbation analysis is seen in Fig. 1. In this figure the bar graph shows the relative sensitivities of the different input parameters (DE corresponds to the vertical steering angle of the sonar, the other parameters are self-explanatory). In the four columns of figures to the right, each row corresponds to the input parameter in the leftmost bar graph. The first and third columns show sonar performance in a vertical slice of the ocean when the input is perturbed by –2.5% and 2.5% respectively. The second and fourth columns show the difference between the perturbed output and the initial output. In this example, the two shallowest sound speeds have the highest sensitivities, i.e., the performance maps change the most when these sounds speeds are changed a relatively small amount.

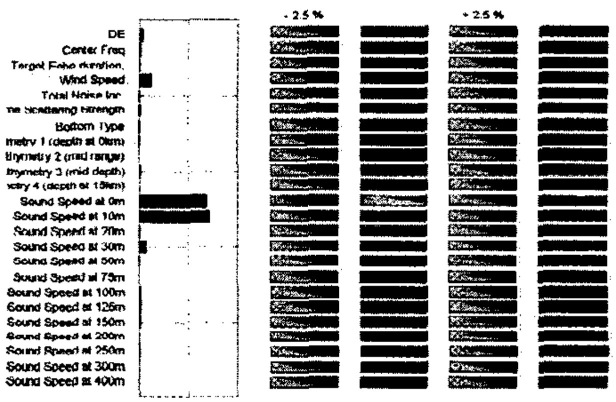

Fig. 1. Sample perturbation analysis for a sonar performance ANN.

We have developed another type of sensitivity analysis that may be useful in this application. In this analysis the output is the 'total' sensitivity of the neural network to each input parameter. These values are obtained by summing the ANN weights connected to each input parameter and provide a measure of the total possible impact of the given input parameter. The total sensitivity of the input i may be found as shown in (1.1).

$$S_i = \sum_{j=1}^{J} O_j(L) \qquad (1.1)$$

where

$$O_j(l) = \sum_{k=1}^{K} O_k(l-1) * w_{kj}$$

$$1 <= l <= L$$

These equations hold for a neural network with L-1 hidden layers. The starting output values $O_j(0) = 1$ in the case of $i=j$; 0 otherwise. The resulting sensitivity can also be normalized by dividing by the total number of weights in the neural network.

Fig. 2 shows a comparison of perturbation analysis to the total sensitivities for a sample reverberation network (i.e., reverberation is the acoustic model product on the output layer of the neural network as opposed to overall performance). The total sensitivity analysis highlights slightly different input elements because it does not take into account the non-linear functions contained within the neural network. The advantage to the total sensitivity analysis is that it does not depend on the operating point.

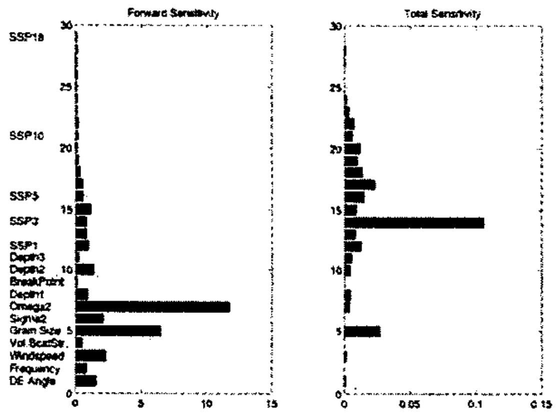

Fig. 2. Comparison of perturbation analysis sensitivities (on the left) to total sensitivities for a reverberation ANN.

B. *Efficient Controller Decisions*

One of the problems with sonar control is that it is usually a joint optimization problem where more than one parameter can be modified. In some cases it is expensive to alter one or more of the parameters. Sensitivity information (like that provided by perturbation analysis) would allow the controller to make intelligent decisions about the necessity of altering one parameter or another. If the sensitivity analysis showed, for example, that the acoustic behavior was highly sensitive to transmitter depth the controller may suggest altering the depth despite the relative difficulty of this action. Alternatively, if the system proved to be insensitive to transmitter depth the controller may opt to work with the current depth.

C. *Improved Optimization Algorithms*

The current optimization algorithms used for sonar control are based on brute force searches, or on stochastically based optimization techniques. It is known that, for many systems, steepest-descent type algorithms may converge more quickly than the former types of optimization. The problem with steepest descent algorithms is that they require gradient information, which is difficult to obtain. This results in another application for sensitivity analysis, as it provides gradients at the operating point. It is possible to use this gradient information to implement steepest descent algorithms that use neural networks for their model evaluations.

Previous work has supplied a method for calculating the gradients of a neural network for a given operating point. One way of doing this is to use the previously described perturbation analysis. Another way, which provides more accurate gradient information, may be found in [7].

The latter method calculates the gradient directly, using only the weights in the neural networks. For each point in the input space we can calculate the gradient $\rho_{kj}(0)$ of the kth output neuron with respect to the jth input neuron as shown in (1.2).

$$\rho_{kj} = \frac{\partial a_k(L)}{\partial a_j(0)} \tag{1.2}$$

where

$$\rho_{ij} = \sum_{m=1}^{N_{l+1}} \rho_{im}(l+1) f'(u_m(l+1)) w_{mj}(l+1)$$

$$0 <= l <= L-1$$

In these equations $a_k(l)$ is the activation value of the kth neuron on the lth layer. The output of a neuron is denoted by u, and the weight connecting the jth neuron of the lth hidden layer to the mth neuron of the $(l+1)$th hidden layer is denoted at w_{mj}. The initial values are given at $\rho_{ij}(L) = 1$ if $i=j$, 0 otherwise.

This method of calculating the sensitivities depends on the operating point of the system. However, it provides a more accurate value for the gradient at a given point than the secant estimation of perturbation analysis.

D. *Environmental Exploration*

One of the problems inherent to using computationally expensive physically based models is that it is difficult to complete sufficient model runs to visualize data across its possible spectrum. For example, a user may wish to look at the effects of changing transmit pulse center frequency across the entire range of possible center frequencies. The user may also want to look at these effects at many different operating points. This type of analysis would require many model runs, and would be expensive to perform with a slow acting model. It would also be difficult to visualize the resulting data.

We can use sensitivity analysis to address this problem. Perturbation analysis can provide the user with the relative sensitivity of the output to a given input at many operating points. Additionally, this analysis can be performed extremely quickly. It is possible to average the results on perturbation analysis over multiple operating points, thereby obtaining a more general measure of sensitivity. Access to the results of sensitivity analysis may allow the users to better understand environmental behavior.

E. *Technology Development*

Ideas for improving technology arise from a realization that the current technology is insufficient. In the sonar control problem, many issues stem from measurements of environmental parameters. It is possible to use sensitivity information to determine which of the environmental parameters need to be measured to higher precision, or with greater accuracy. Having this information can help direct research efforts to those areas where improvement will be of the greatest utility.

V. FUTURE WORK

This paper proposes many potential applications for sensitivity analysis. It also details methods in which the requisite sensitivities can be calculated. However, more work is required to fully develop the applications.

The methods for calculating sensitivities have been developed, but no complete comparison has been made between them. Such a comparison would allow for a better understanding of the relative strengths and weaknesses of different sensitivity calculations. This, in turn, would allow for better application of the technology. It is possible that such work would also reveal ways to improve the sensitivity calculations proposed above.

It will be important to determine the best way to interpret the outputs from each type of analysis. One case in which this is difficult is the situation where two input values have different units, and possibly vastly different scales. It is difficult to know how a change in one set of units compares to a change in another set of units. This issue is addressed somewhat by looking at the sensitivities as percentages of the entire input and output range. However, there could be improvement with a novel presentation of the data.

Finally, this paper suggests that sensitivity information could be used by the optimization algorithm to improve controller behavior. Specifically, using a steepest descent algorithm may provide a faster convergence than the current brute force or stochastic searches. The sensitivity information can be used to implement the improved algorithm, but a rigorous study of this idea is still necessary.

ACKNOWLEDGMENTS

This work is sponsored by the Office of Naval Research, contract number N00014-01-G-0460.

REFERENCES

[1] R.D. Reed and R.J. Marks II, Neural Smithing: Supervised Learning in Feedforward Artificial Neural Networks, MIT Press, 1999.

[2] J.B. Jung, "Neural network ensonification emulation: training and application," PhD. Dissertation, University of Washington, Seattle, WA, 2001.

[3] W.L.J. Fox, M.U. Hazen, C.J. Eggen, R.J. Marks, and M.A. El-Sharkawi, "Environmentally adaptive sonar control in a tactical setting", in press.

[4] A.G. Bahbah and A.A. Girgis, "Input feature selection for real-time transient stability assessment for artificial neural network (ANN) using ANN sensitivity analysis", *Power Industry Computer Applications, Proceedings of the 21st IEEE International Conference*, pp. 295-300, 1999.

[5] A.H Rambhia, R. Glenny, and J.-N. Hwang, "Critical input data channels selection for progressive work exercise test by neural network sensitivity analysis", *Proceedings IEEE International Conference on Acoustics, Speech, and Signal Processing*, Volume: 2 , pp. 1097-1100, 1999.

[6] C. Panayiotou, C. Cassandra, and W.-B. Gong, "Model abstraction for discrete even systems using neural networks and sensitivity information", *IEEE Proceedings of the 2000 Winter Simulation Conference*, pp.335-341.

[7] J.-N. Hwang, J.J. Choi, S. Oh, and R.J. Marks II, "Query-based learning applied to partially trained multilayer perceptrons", *IEEE Transactions on Neural Networks*, Vol. 2, pp. 131-136, January 1991.