Not to be cited without prior reference to the authors CM2004/R:05

Theme Session R: New developments in fisheries acoustics; applications in bottom trawl surveys and multi-frequency species identification

Using artificial neural networks to combine acoustic and trawl data in the Barents and North Seas

Suzanna Neville[1], Vidar Hjellvik[2], Steven Mackinson[1], and Jeroen van der Kooij[1]

Abstract

Groundfish have a wide and variable distribution making the use of trawling alone a highly inadequate sampling method. Trawl data provide species identification and numbers over a very small area and habitat type while acoustic data provide a wider coverage of the ecosystem, but fail to identify species. Both methods provide essential information required for assessing fish stock abundance and distribution, but no systematic method of combining these data has yet been identified. Acoustic and trawl data were collected for both the North and Barents Seas and their relationships investigated using artificial neural networks (ANNs). ANNs are information processing systems that were inspired by the structure of the brain. They can incorporate multiple variables and explore complex interactions between variables in far greater depth than many traditional statistical techniques. This modelling tool has had many successes in environmental modelling and is increasingly being applied in situations where the underlying relationships are poorly known. Network architectures, optimisation of connection weights (learning), model validation, and the reasons for the difference in performance between the North Sea and Barents Sea models are discussed.

Keywords: Trawl surveys; acoustics; groundfish; artificial neural networks; stock assessment; non-linear relationships

[1] The Centre for Environment, Fisheries and Aquaculture Science (CEFAS), Lowestoft Laboratory, Pakefield Road, Lowestoft, Suffolk, NR33 0HT, UK. [tel: +44 (0)1502 527729; fax: +44 (0)1502 524546, e-mail: s.neville@cefas.co.uk] www.cefas.co.uk
[2] Institute of Marine Research, P.O. Box 1870 Nordnes, N-5817 Bergen, Norway [tel: +47 55 23 85 00, fax: +47 55 23 85 31, e-mail: vidarh@imr.no] www.imr.no

1. Introduction

Trawl and acoustic data each provide essential information for the accurate assessment of fish stocks. The identification, size, and age of fish are provided by the trawl data, and the more precise location of fish and greater coverage of the ecosystem are provided by the acoustic data. These data cover different sections of the water column and have different geographical coverage making the use of either acoustic or trawl data alone highly limiting, but the combination of the two highly attractive. However, a method for combining these data in a systematic way has yet to be identified.

Artificial neural networks are an appropriate choice for the combination of acoustic and trawl data because the underlying relationships between the variables do not need to be known or programmed, but they are 'learned' iteratively by the network being supplied with a number of training examples from the data set (Lek *et al*, 1996; Basheer and Hajmeer, 2000; Maier and Dandy, 2000). Their success in such a vast array of predicting, classification, and control problems can be attributed to their capability to model extremely complex functions in heterogeneous data sets, with large numbers of variables, between which the relationships can be non-linear. Studies involving environmental data very rarely involve variables with linear relationships, although this linearity is commonly assumed by traditional statistical methods (James and McCulloch, 1990).

Neural networks have been successfully applied to a variety of different problem types in a diverse range of subject areas including; engineering, medicine, geology, physics, and environmental science. In fisheries science there have also been a wide range of successful applications; forecasting fish stock recruitment (Chen and Ware, 1999), fish population counting (Newbury *et al.*, 1995) and predicting daily fish food consumption in fish farming (Ruohonen, 1999), fish species recognition during trawl processing (Storbeck and Daan, 2001), and also species identification of acoustic data (Haralabous and Georgakarakos, 1996; Simmonds *et al.* 1996), prediction of stock abundance in capelin (Huse and Gjøsæter, 1999), and for analysing catch per unit effort/abundance relationships in tuna fisheries (Gaertner and Dreyfus-Leon, 2004).

In this study, both acoustic and bottom trawl data are combined using artificial neural networks (ANNs). Ultimately the index produced by this combination could be used to improve the fish stock assessment process by providing more realistic estimates of stock location and size. This is because the area covered by the assessment surveys would be much greater and the number of data points much higher resulting in a reduced variance. The type of networks used are known as regression networks where the network learns the pattern between the input variables and a continuous output variable. The combination of acoustic and trawl data in two different locations, the North Sea and the Barents Sea, are compared and discussed.

2. Materials and Methods

The two different locations used for this study were the North Sea between UK, France and south west Norway (56°N, 3°E), and the Barents Sea off the north coast of Norway (72°N, 25°E). Trawl, acoustic, and environmental data were collected during bottom trawl surveys in 2000, 2001, and 2002 for the North Sea and in 1997 –2002 inclusive for the Barents Sea.

North Sea data

In the North Sea, data were collected during the International Bottom Trawl (IBTS) surveys carried out on board R.V. Cirolana in August of each year. Acoustic data were collected using a Simrad EK500 scientific echosounder with a hull mounted 38 kHz split-beam transducer, and post-processed using Sonardata Echoview®. Acoustic measurements were logged continuously during all three surveys, but only data collected during daylight hours were used for this study, in line with the IBTS protocol for daytime trawling (ICES, 1999). Weather conditions were variable during the surveys, and data collected during the most adverse conditions were discarded from the analyses. Standard IBTS protocols were observed (ICES, 1999). CTD data were collected during each trawl. Data from all three surveys were prepared using the same method described below and then combined. A total of 211 trawl stations were used in the analysis.

Barents Sea data

In the Barents Sea, data were collected on R.V Johan Hjort and G.O. Sars in February each year. Acoustic data were collected using a Simrad EK500 scientific echosounder with a 38 kHz transducer, and post-processed using Bergen Echo Integrator (BEI). Acoustic measurements were logged continuously during all surveys, and only daylight data used. Environmental data such as temperature and depth were collected throughout the surveys. Data from all surveys were prepared using the same method and then combined. A total of 2973 trawl stations were used in the analysis.

2.1 Data Preparation

Acoustic Data

Initial processing of acoustic data involved removing spurious data associated with rough weather, vessel turning, interference from other instruments, and excessive surface aeration, to ensure only high quality data were included. This was followed by sub-dividing data into ‘on station’ and ‘interstation’ sections (Fig. 1). On station sections correspond to data collected during a tow where both acoustic and trawl data are available, and interstation data to the section of cruise track between trawl stations where only acoustic data are recorded.

The line respresenting the seabed on the echogram was corrected by making adjustments where wrecks or fish aggregations had been included in the bottom signal, or where the seabed line had been recorded as below the actual seabed. A backstep was introduced of between 0.1 m in favourable weather conditions, to 0.5 m in those less favourable. This enabled the area very close to the seabed to be excluded which is vital where conditions have been rough causing the bottom line to contain a number of significant inaccuracies. If included, they would interfere with any relationships in the data, but by exclusion, there is a very real possibility that the fish within this area would not be represented.

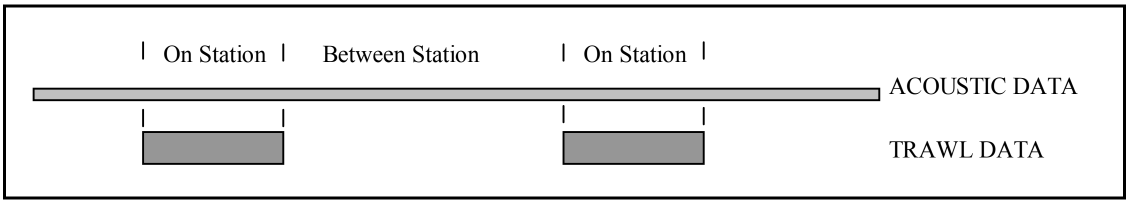

Fig. 1: Diagram representing on station and interstation sections of the cruise track; on station sections have both acoustic and trawl data present, and interstation have only acoustic data as no trawling has taken place.

The water column for all sections of the cruise track was divided into depth layers, which were started from the seabed and continued upwards (Fig.2). For the North Sea layers of 1 m were created from the seabed to 10 m above, and then 10 m layers were created up to the surface. For the Barents Sea only layers of 10 m were created due to the greater depth in this location. The total number of layers depended on the depth in different areas of the cruise track.

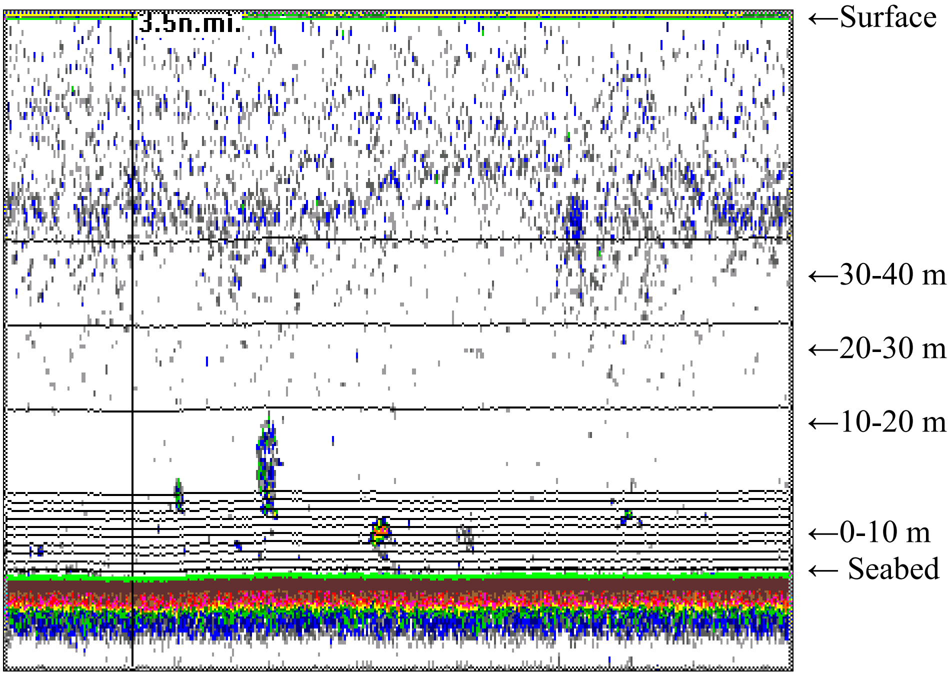

Fig. 2: Echogram showing the bottom referenced layers created at different depths. The ten 1 m layers in the first 10 m above the seabed were only created for the North Sea data. For the Barents Sea data one single layer was created to respresent the first 10 m above the seabed.

Following postprocessing, data were exported into spreadsheets using threshold of -70 dB to ensure that any values smaller than this caused by plankton and noise did not mask the signals from fish. The exported spreadsheets provide a value of nautical area scattering coefficient (NASC, $m^2\ n.mi^{-2}$) for each individual layer within each EDSU:

$$NASC = 4\pi(1852)^2\, s_a \qquad (1)$$

where s_a is the area backscattering coefficient and 1852 is metres per nautical mile.

Trawl Data

Data were recorded for time, position, gear geometry, and environmental conditions for each trawl. This included mean depth and mean water temperature close to the bottom. The catch data were divided up into demersal, pelagic, bottom fish and others. The catch data for each haul included number caught, root mean square length, mean target strength and mean weight for each species caught. Only species making up greater than 5% by weight of the haul were included in this study.

To investigate the relationship between the fish caught in the trawl and the acoustic marks recorded on the echogram, comparable units need to be calculated. The *observed* NASC corresponds to the data exported from Echoview/BEI that was directly measured, and the *equivalent* NASC (ENASC) was calculated using measured quantities from the trawl data; catch numbers, the root mean square (rms) length (cm), and the swept area (n.mi.^{2}). A reference quantity, b, was also used to calculate the mean target strength for an individual fish:

$$TS = 20\log L + b \tag{2}$$

where b is the species specific target strength constant (ICES, 2000) and L is the rms length for a particular species within a haul. The derived target strength was then used to calculate the expected acoustic energy from an individual fish:

$$\sigma_{sp} = 4\pi 10^{(TS/10)} \tag{3}$$

The derived value of spherical scattering cross section (m^2) was multiplied by the number of each species for each haul to produce the ENASC of the total catch at each station:

$$\text{ENASC of total catch for a station} = \sigma_{sp} \text{ x number of fish} \tag{4}$$

To make the ENASC comparable with the observed NASC it must be divided by the swept area to ensure that both are standardised densities:

$$\text{Swept area} = \text{tow length (n.mi) x door spread (n.mi)} \tag{5}$$

$$\text{Standardised ENASC } (m^2 \text{ n.mi.}^{-2}) = \text{ENASC} / \text{swept area} \qquad (6)$$

The ENASC represents the equivalent acoustic energy that would have been produced by the fish caught in the trawl, which could then be compared against the *observed* acoustic energy for the corresponding section of cruise track.

2.2 Artificial Neural Networks

Multi layer perceptrons (MLPs) trained using backpropagation (BP) are the simplest and most commonly used network type primarily due to their versatility and speed (Basheer and Hajmeer, 2000). They consist of three layers: input, output and hidden (Fig. 3). Occassionally two hidden layers can be used for more complex problems. The input layer contains the independent variables, the output layer the dependent variables, and the hidden layer contains the hidden nodes which capture the non-linearity in the data. The hidden nodes do not interact with the outside of the network and are therefore 'hidden' within the network structure. The input layer has no activation function as it supplies the hidden layer with the input values. The number of nodes in this layer equals the number of variables supplied to the network. Most MLPs are feedforward, fully connected networks where all the neurons in each layer emit a signal to all neurons in the next layer, with the exception of the output layer.

Learning is supervised where both the input and the output variables are supplied to the network, in contrast to unsupervised where only the inputs are supplied and the network learns the stucture of the input values to produce a consistent output. Weights are applied to the connections between nodes and are initially random. They are adusted throughout the iterative training procedure undertaken using training algorithms to optimise the predictions made by the network in the output layer. Learning is complete when further modifications to the weights no longer improves the predictions, and at this point the weights are then fixed. Data that have not been used in training are then supplied to the input layer, run through the network with fixed weights and predictions are made. Values in this previously unseen data that are similar to those in the training data should have a similar output when run through a trained network.

Training Algorithms

Back propagation is the most widely used algorithm for network training and there are numerous examples in numerous fields including fisheries science (Engelhard *et al*., 2003; Maier and Dandy, 2000; Gaertner and Dreyfus-Leon, 2004; Mastrorillo *et al.*, 1997; Ruohonen, 1999). The name back propagation refers to the propagation of error back through each layer of the network (Rumelhart *et al.*, 1986). Neurons in the hidden layer evaluate the intensity of the signal given by the input layer and this signal is quantified by a sigmoid function (logistic or hyperbolic).

The algorithms extensively tested in this study were back propagation (BP), conjugate gradient descent (CGD), and Levenberg-Marquardt (LM). Algorithms can be local or global. Global methods include simulated annealing (Kirkpatrick *et al*., 1983) and genetic algorithms (Huse *et al.*, 1999). BP, CGD, and LM are all local methods of which there are two different categories: first-order methods (BP) and second-order methods (CGD, LM). BP is based on a linear method (gradient descent) whereas CGD and LM are based on quadratic methods. Second-order algorithms more commonly get stuck in local minima, but are much faster overall. For these reasons the networks in this study were trained first using backpropagation, followed by either CGD or LM.

Resampling

Network performance is dependent on the random weights initially assigned to the connections. This can lead to networks with identical architecture performing quite differently depending on the initial weights assigned. For this reason it was necessary to run several samples and then use the mean when assessing design decisions to ensure that the network architecture selected was robust, and not just a good network due to a particular set of initial weights.

Network Architecture

Adjustments to network architecture lead to an extensive range of model complexities. The number of nodes in the hidden layer, learning algorithms, and the type of activation function, must all be selected correctly to optimise the network architecture. Many of these choices have to be made by trial and error and this is a very difficult and time consuming and task in the model building process (Maier and Dandy, 2000).

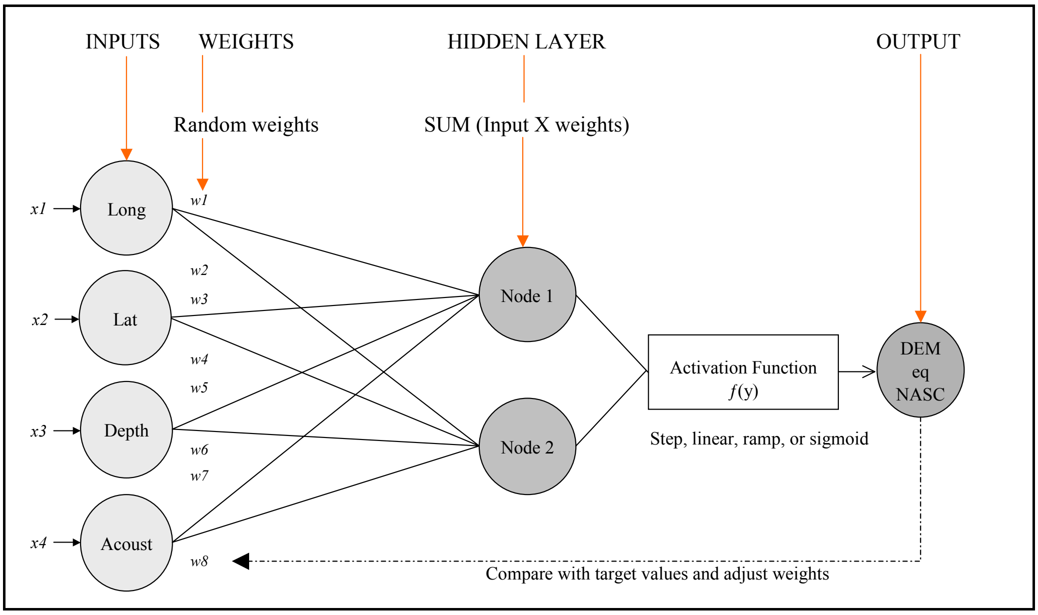

Fig. 3: The basic structure of an artificial neural network showing the input, hidden and output layers; connections between the nodes and layers; and the weights and activation functions within the network.

Building a model

The data used for network training in this study were the on station data. The relationships to be mapped were those between the acoustic data and the trawl data that were collected simultaneously. Following network training, the interstation acoustic data were then supplied to the network, and the trawl data for the section of the cruise track where no trawling occurred could be estimated by the model.

There were two main components of the modelling process. Firstly, network training involving the optimisation of the network parameters described above. Secondly, model selection and testing.

Network Training

The number of hidden nodes for each network determines the number of connection weights (free parameters). The greater the number of connection weights, the better the network is able to represent the function to be approximated, but the less likely the model will be able to generalise. This will result in the model failing when previously unseen data sets are applied. Although the aim of training is to reduce the error as far as possible, reducing the error too much leads to the network learning the noise in the training data, rather than the underlying relationships. When this occurs the networks has been overtrained. Taking the time to optimise the network design is critical in producing a reliable model when new data are supplied. Optimality is defined as the smallest network that captures the relationships in the training data (Maier and Dandy, 2000).

Sensitivity Analysis

Sensitivity analysis is performed to assess the contribution of each variables to the overall network performance. A ratio is calculated of the error with the variable unavailable, to the error with it available. Significant variables have a high ratio which indicates the network performs far less well without them present, and poor variables have a low ratio. The variables for each of the two locations were ranked in order of ratio value so that the contributions of each variable could be compared.

Model Selection

To assess the initial performance of each of the networks two parameters were used:

- The standard Pearson-R correlation coefficient between the observed and predicted output values;
- The error:standard deviation ratio to ensure the estimates have a lower prediction error standard deviation than the standard deviation of the training data.

Model Testing

To test the models that had performed well using the model selection criteria, previously unseen data were supplied to the networks, and statistics of the predicted data were compared to those statistics of the observed data. By comparing the mean, minimum, maximum, range, and standard deviation of these data, it was possible to assess whether the models had been overtrained, or whether they were able to generalise well when supplied with new data. Some of the models predicted output values that were unlikely to occur, such as negative or extremely high values. Although acoustic data are characterised by few extremely high values which can be justified by true biological phenomena, some of the models predicted values beyond these.

3. Results

The networks were able to map the pattern between the input and the output values much more consistently for the Barents Sea data than for the North Sea data. This was demonstrated by the consistency of the network performance compared to the difficulties experienced when training the networks using the North Sea data. Two multilayer perceptrons were trained, both with four input variables; latitude, longitude, depth, and acoustics (Fig. 4). The dependent variable in the output layer was demersal NASC.

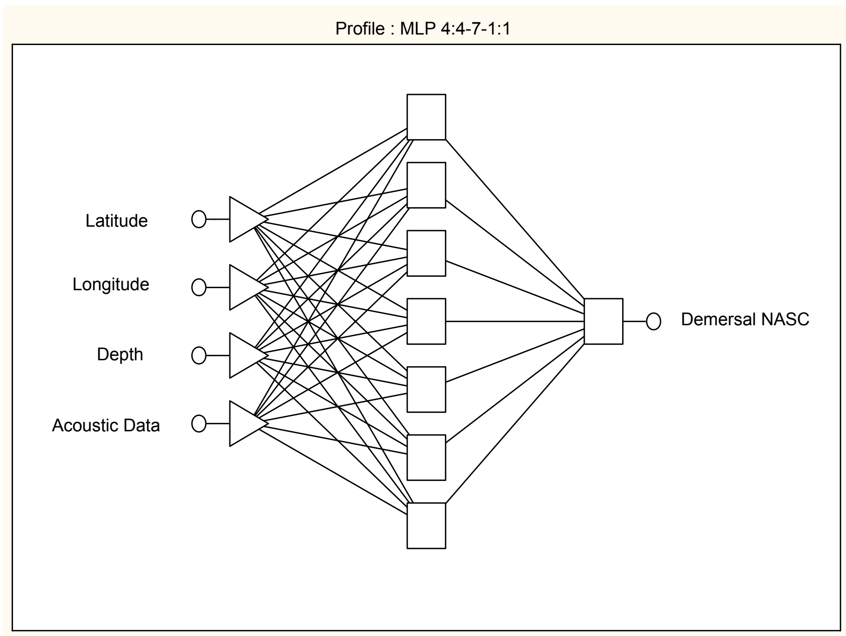

Fig. 4: Network architecture for the model used in this study. The input layer consisted of 4 variables in both the Barents Sea and the North Sea; the number of nodes in the hidden layer was determined by testing the performance of the model using a range of node number; and the output layer consists of the dependent variable demersal NASC.

Optimising the number of nodes in the hidden layer

Testing the network using varying numbers of nodes within the hidden layer forms a vital part of model optimisation. For each node number 10 samples were tested, and the mean value of the correlation coefficient between observed and predicted output values, and the ratio between the error and training data standard deviations were calculated (Fig. 5).

The correlations between the observed and predicted values for the North Sea are consistently lower than those for the Barents Sea indicating a poorer ability for the networks to learn the patterns. The values for the SD ratio are significantly higher for the North Sea, indicating again that the networks have not been able to map the inputs to the outputs as well as for the Barents Sea data.

Figure 5a is a graphical respresentation of the results for the North Sea networks. As the node number increases there is virtually no improvement in network performance, and the standard deviation bars remain relatively large. This results in a very difficult decision regarding the optimal number of nodes to use in the hidden layer of the network. A network with 4 nodes was used as keeping the nodes number low prevents the network overlearning the data.

The results for the Barents Sea data are significantly clearer than for the North Sea. The performance parameters show a rapid increase in model performance initially as the number of nodes increases from 1 to 6, and then becoming more stable after this number (Fig. 5b). The optimal node number chosen for the final models for the Barents Sea data was 7. The standard deviation bars are relatively small in the Barents Sea networks indicating a greater consistency in the training of the networks. After this number there is little improvement in model performance, but with a sharp increase in the likelihood of the model becoming overtrained due to a high number of nodes in the hidden layer. This decision is also justified by the relatively small standard deviation indicating that the network will perform consistently.

a) North Sea

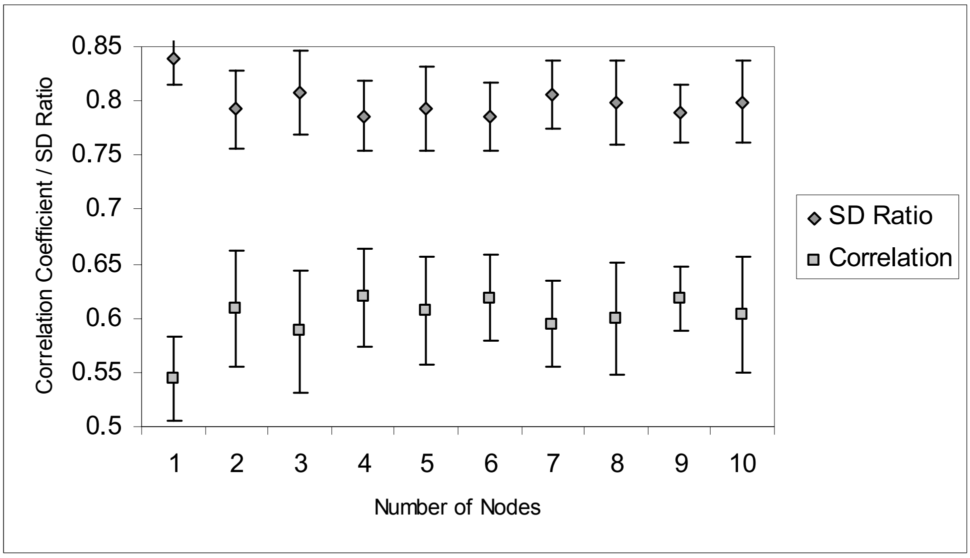

b) Barents Sea

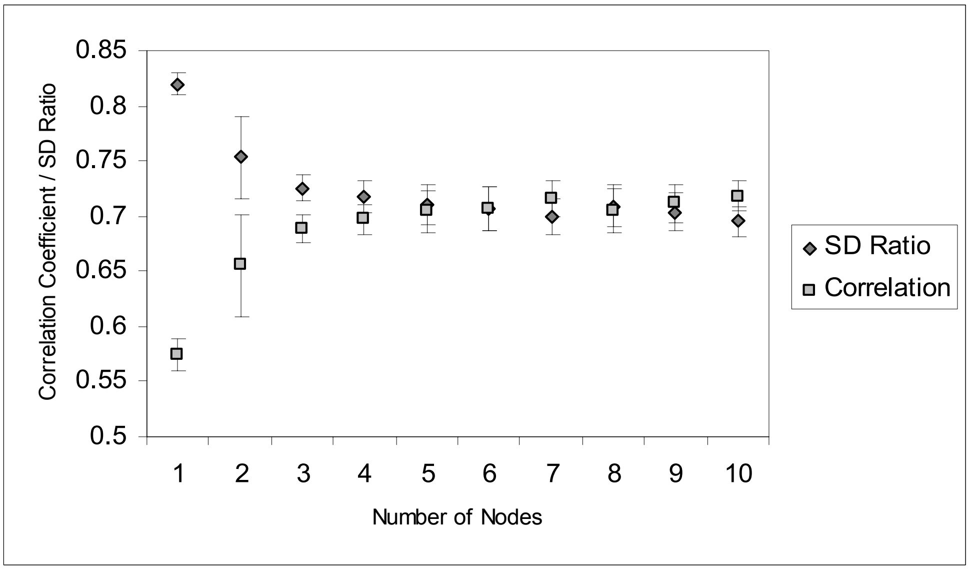

Fig. 5: Changes in the correlation coefficient between the observed and predicted output values; and the ratio (SD Ratio) between the error standard deviation (predicted output minus observed output) and the training data standard deviation (observed output values) when the number of nodes in the hidden layer are adjusted.

Sensitivity Analysis of Input Variables

The results of sensitivity analysis must be interpreted with caution, as redundancies and interdependencies between variables can lead to quite different results on another network applied to the same data set. For this reason, sensitivity results were calculated for 100 models for each location to check for the consistency of the ranking for each variable (Fig. 6). Each model has the same input and output variables, but different initial weights. There are ten samples for each node number (0-10).

The acoustic variable used in network training was consistently ranked higher for the Barents Sea models than for the North Sea data (Fig. 6). The acoustic data for the Barents Sea was ranked either first or second in 80 % of the models (Fig. 6b), with the acoustic data for the North Sea ranked first or second in 0 %, and lowest in 80 % (Fig. 6a). This indicated that the acoustic data contributed more to the overall performance of the models for the Barents Sea than for the North Sea. Depth proved to be an important variable in both locations, with longitude and latitude being ranked lowest. During initial analyses using all available input variables, different combinations were tested, and those ranked the lowest were removed to assess whether this had any positive impact on the models. All variables were removed apart from acoustics, depth, latitude, and longitude for the Barents Sea. For the North Sea, acoustics was ranked consistently low and would have been removed also if it was not the primary variable in the analysis.

a) North Sea

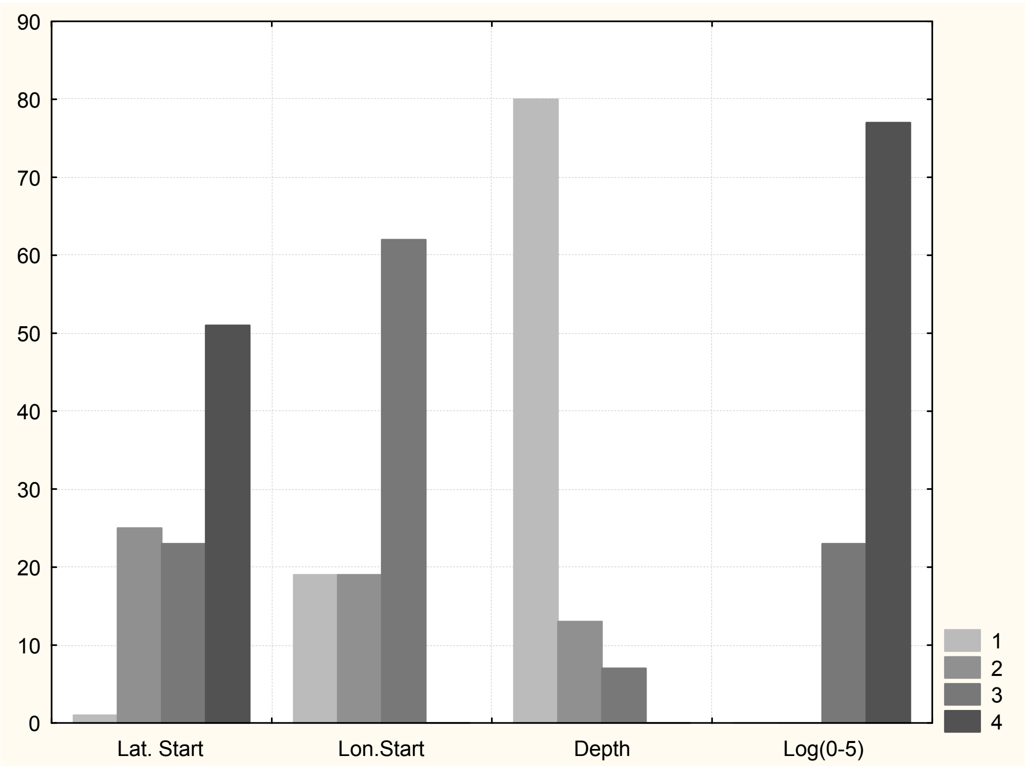

b) Barents Sea

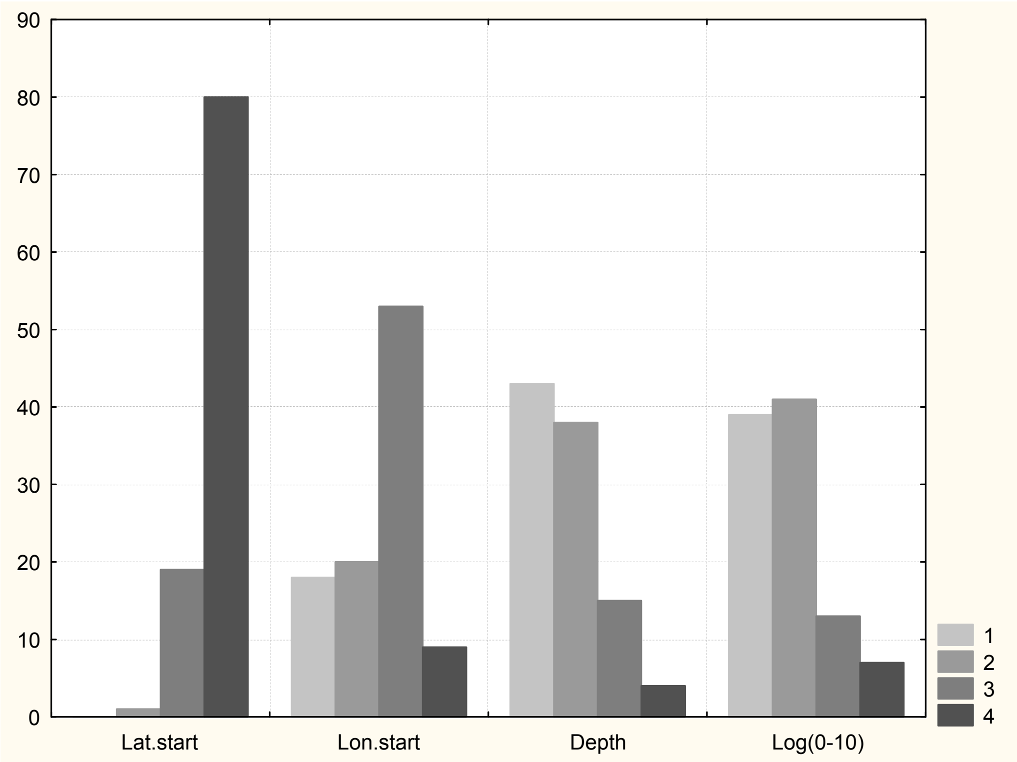

Fig. 6: The ranking of variables for the Barents Sea data indicating the relative contributions of each variable to the network training process for 100 different networks for each location (1=highest contribution, 4=lowest contribution).

a) North Sea

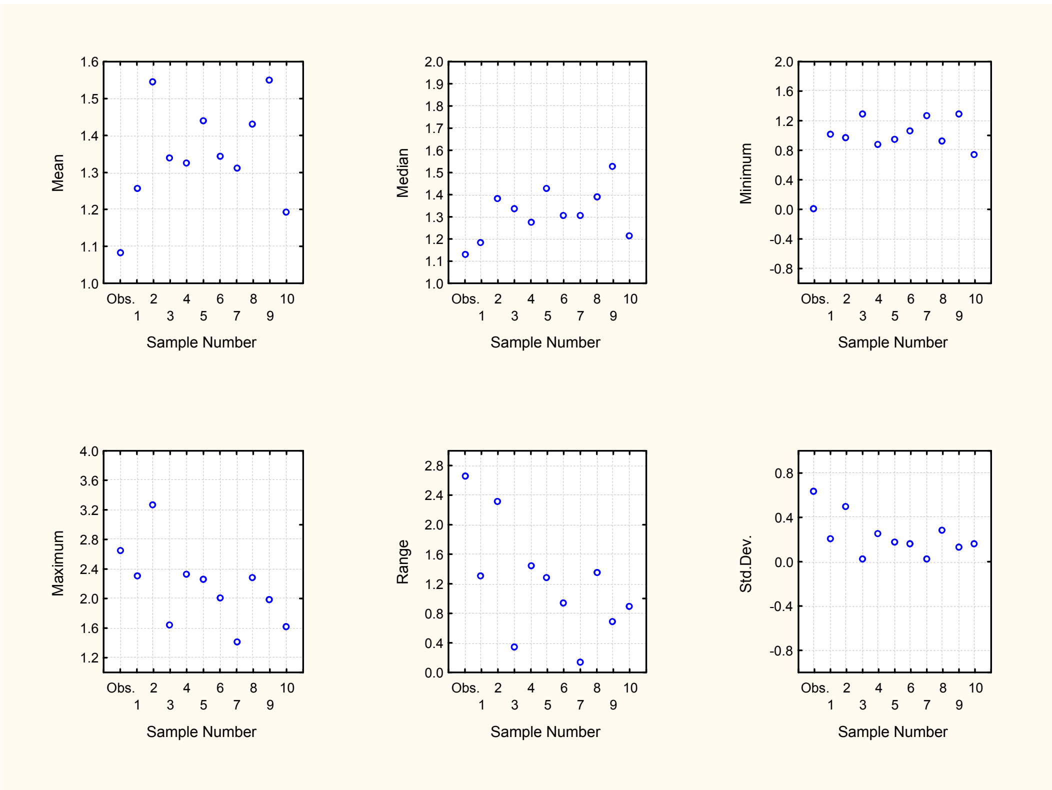

b) Barents Sea

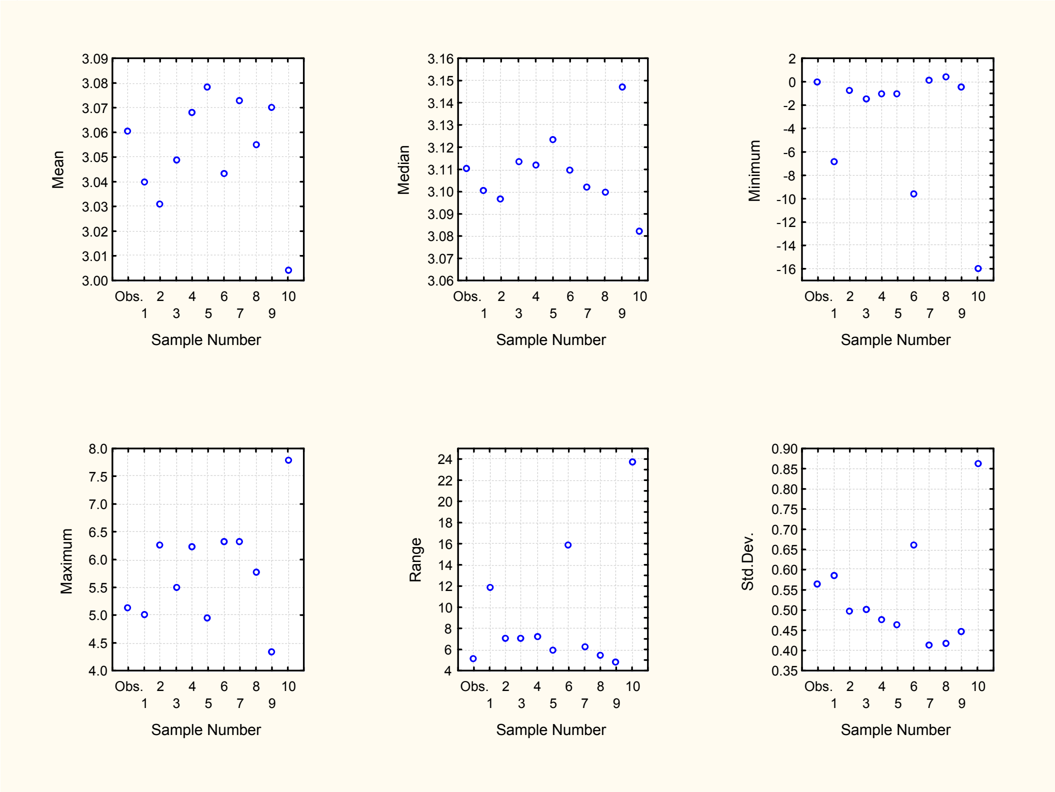

Fig. 7: Comparison between demersal NASC observed on station data with predicted data from the 10 sample models.

Model Testing

Following network architecture optimisation, a method to select the best performing model from the samples for each location had to be derived. The samples differ only in the initial weights applied which are randomly assigned and then adjusted throughout the iterative training process. The between station data were supplied to all the samples and the predicted values for demersal NASC were output. The statistics of these data were compared with those of the on station observed data to test whether the predicted values appeared reasonable. The differences between the statistics of the observed and model predictions for the Barents Sea data were much smaller than for those for the North Sea (Fig. 7). For the Barents Sea sample 8 was selected as the best performing model as the statistics of the predicted data more closely matched the observed data than for the other samples (Fig.7b). For the North Sea, the selection of the best performing model was much more difficult as the observed data point was often found to be completely outside the range of any of the data points of the predicted data (Fig.7a). This led to no model in particular being able to be confidently chosen from the samples.

4. Discussion

In combining acoustic and trawl data for two different locations, it has been demonstrated that there are clear differences in the success of extrapolating on station trawl data using interstation acoustic data. The combination using these methods worked much more successfully for the Barents Sea than for the North Sea. However, neither of the final models allowed the trawl data to be extrapolated adequately, with the models often under predicting the abundance of fish by being unable to predict the few particularly high values that are characteristic of fish stocks.

The reasons for this contrast in model performance can be attributed to a number of different factors. The surveys in the Barents Sea have a far higher density of trawl stations with approximately 175 hauls per survey. This is compared to a maximum of 75 hauls per survey in the North Sea. A total of 211 data points were available for the North Sea, which must be divided up for the network training process. This does not allow the network to be trained using a large enough range of points. This problem is exacerbated by the heterogeneous nature of the North Sea; making it necessary to ensure enough data points are available to fully represent all areas. The current surveys do not collect a sufficient quantity data that satisfy these criteria.

In the southern North Sea depths as shallow as 30 m are common, resulting in significant noise in the acoustic data if the weather is rough. Further to the north between Scotland and Norway the depth increases to ~ 200 m. By treating all the data points in the survey in the same way, any difference in relationships due to these significant depth changes are not accounted for. The species composition is quite different in the northern and southern regions also, leading to differences in the behaviours of the fish. This may mask any relationship between the trawl and acoustic data if the two areas of the North Sea are treated similarly. In the Barents Sea, the depth is much greater leading to a smaller percentage of the water column containing noise due to bad weather or acoustic dead zones. The species compositions are more distinct leading to clearer relationships being found between what is caught in the net and what is captured by the echosounder (Beare *et al.,* CM2004).

The combination of acoustic and trawl data is based on the assumption that the fish caught in the trawl can be detected by the echosounder. There are a number of reasons why this may not always hold true. The net may not always be aligned directly behind the boat leading the fish being caught that are not passing underneath the ship. Also the different behaviours of the fish after the boat has passed over and when the net is fishing will seriously affect the relationship between these two data sets (Aglen *et al*. 1999, Michalsen *et al*., 1996).

Numerous different methods have been used to approach this problem but with little success (Beare *et al.*, CM2004; Mackinson *et al*., CM2004). Until these behavioural and survey design issues have been further investigated, it is difficult to see how a relationship between these data sets can be determined.

Acknowledgements

This work was funded under EU contract QLRT-CT-2000-02038 (CATEFA). We are grateful to the crews of the R.V. Cirolana, R.V Johan Hjort and G.O. Sars for their assistance during the fieldwork phase of this project.

References

Basheer, I.A. and Hajmeer, M., 2000. Artificial Neural Networks: fundamentals, computing, design and application. Journal of Microbiological Methods **43,** 3-31.

Chen, D.G. and Ware, D.M. 1999. A neural network model for forecasting fish stock recruitment. Canadian Journal of Fisheries and Aquatic Science **56**, 2385-2396.

Engelhard, G.H., Dieckmann, U., and Godø, O.R. 2003. Age at maturation predicted from routine scale measurements in Norwegian spring-spawning herring (*Clupea harengus*) using discriminant and neural network analyses. ICES Journal of Marine Science **60**, 2, 304-313.

Gaertner, D. and Dreyfus-Leon, M. 2004. Analysis of non-linear relationships between catch per unit effort and abundance in a tuna purse-seine fishery simulated with artificial neural networks. ICES Journal of Marine Science **61**, 812-820.

Haralabous, J and Georgakarakos, S. 1996. Artificial Neural Networks as a Tool for Species Identification of Fish Schools. ICES Journal of Marine Science **53**, 173-180.

Huse, G., Strand, E., and Giske, J. 1999. Implementing behaviour in individual-based models using neural networks and genetic algorithms. Evolutionary Ecology **13**, 469-483.

Huse, G. and Gjøsæter, H. 1999. A neural network approach for predicting stock abundance of Barents Sea capelin. Sarsia **84**, 457-454.

ICES, 1999. Manual of the International Bottom Trawl Surveys, Revision VI. ICES CM 1999/D:2, Addendum 2.

ICES, 2000. Report of the Planning Group for herring Surveys. ICES CM 2000/G:02, Appendix 6.

James, F.C. and McCulloch, C.E. 1990. Multivariate analysis in ecology and systematics: panacea or Pandora's box? Ann. Rev. Ecol. Syst. **21**, 129-166.

Kasabov, N.K., 1996. Foundations of Neural Networks, Fuzzy Systems, and Knowledge Engineering. MIT Press.

Kirkpatrick, S., Gilatt, C.D., and Vecchi, M.P. 1983. Optimisation by simulated annealing. Science **220**, 671-680.

Lek, S., Delacoste, M., Baran, P., Dimopoulos, I., Lauga, J., and Aulagnier, S., 1996. Application of neural networks to modelling nonlinear relationships in ecology. Ecological Modelling **90**, 39-52.

Maier, H.R and Dandy, G.C., 2000. Neural networks for the prediction and forecasting of water resources variables: a review of modelling issues and applications. *Environmental Modelling & Software* **15**, 101-124.

Mastrorillo, S., Lek, S., and Dauba, F. 1997. Predicting the abundance of minnow *Phoxinus phoxinus* (Cyprinidae) in the River Ariège (France) using artificial neural networks. Aquatic Living Resources **10** 169-176.

Michalsen, K., Godø, O.R., and Fernö, A. 1996. Diel variation in the catchability of gadoids and its influence on the reliability of abundance estimates. ICES Journal of Marine Science **53**, 389-395.

Newbury, P.F., Culverhouse, P.F., and Pilgrim, D.A. 1995. Automatic fish population counting by artificial neural network. Aquaculture **133**, 45-55.

Rumelhart, D.E., Hinton, G.E., and Williams, R.J. 1986. Learning representations by back-propagating errors. Nature **323**, 533-536.

Ruohonen, K. 1999. Modelling fish food consumption with artificial neural networks. Aquaculture Research **30**, 545-548.

Simmonds, E.J., Armstrong, F., Copland, P.J. 1996. Species identification using wideband backscatter with neural network and discriminant analysis. ICES Journal of Marine Science **53**, 189-195.

Storbeck, F. and Daan, B. 2001. Fish species recognition using computer vision and a neural network. Fisheries Research **51**, 11-15.