

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/222549413>

An approach for determining relative input parameter importance and significance in artificial neural networks

Article in *Ecological Modelling* · June 2007

DOI: 10.1016/j.ecolmodel.2007.01.009

CITATIONS

31

READS

263

3 authors, including:



Stanley J. Kemp

University of Baltimore

13 PUBLICATIONS 358 CITATIONS


[SEE PROFILE](#)



Patricia A Zaradic

13 PUBLICATIONS 702 CITATIONS

[SEE PROFILE](#)

available at www.sciencedirect.comjournal homepage: www.elsevier.com/locate/ecolmodel

An approach for determining relative input parameter importance and significance in artificial neural networks

Stanley J. Kemp^{a,*}, Patricia Zaradic^b, Frank Hansen^a

^a Department of Biology, Leidy Laboratories, University of Pennsylvania, 415 S. University Ave., Philadelphia, PA 19104, United States

^b Stroud Water Research Center, 970 Spencer Road, Avondale, PA 19311, United States

ARTICLE INFO

Article history:

Received 31 August 2005

Received in revised form

10 January 2007

Accepted 19 January 2007

Published on line 28 February 2007

Keywords:

ANN

Simulation

Artificial neural network

Parameter importance

Virtual ecology

ABSTRACT

Artificial neural network (ANN) models are powerful statistical tools which are increasingly used in modeling complex ecological systems. For interpretation of ANN models, a means of evaluating how systemic parameters contribute to model output is essential. Developing a robust, systematic method for interpreting ANN models is the subject of much current research. We propose a method using sequential randomization of input parameters to determine the relative proportion to which each input variable contributes to the predictive ability of the ANN model (termed the holdback input randomization method or HIPR method). Validity of the method was assessed using a simulated data set in which the relationship between input parameters and output parameters were completely known. Simulated data sets were generated with known linear, nonlinear, and collinear relationships. The HIPR method was performed repetitively on ANN models trained on these data sets. The method was successful in predicting rank order of importance on all data sets, performing as well as or better than the recently proposed connectivity weight method. One main advantage of using this method relative to others is that results can be obtained without making assumptions regarding the architecture of the ANN model used. These results also serve to illustrate the consistency and information content of ANN models in general, and highlight their potential use in exploring ecological relationships. The HIPR method is a robust, simple, general procedure for interpreting complex ecological systems as captured by ANN models.

© 2007 Elsevier B.V. All rights reserved.

1. Introduction

One of the primary goals of ecology is to establish correlative and putative links between ecological parameters. The use of artificial neural networks (ANNs) comprises a statistical modeling tool which has been shown to be highly successful in investigating these links and have had much recent application to ecological systems (Guégan et al., 1998; Basheer and Hajmeer, 2000; Lek and Guégan, 2000; Vander Zanden et al., 2004). ANNs have been shown to be highly competent

approximators of many complex functions, and have shown advantages over general linear models in predictive ability (Hastie et al., 2001). Usage of ANN models has increased such that ANNs are now found in many statistical packages (e.g. S-Plus and R, Venables and Ripley, 2002; JMP 5.1-SAS Inst., 2003). However, a major limitation of modeling ecological data with ANNs is that the relative contribution of each systemic input parameter to model output generally remains unknown (Olden and Jackson, 2002; Gevrey et al., 2003; Olden et al., 2004). Although predictive accuracy is a major benefit of ANN

* Corresponding author. Tel.: +1 610 777 1792.

E-mail address: skemp@sas.upenn.edu (S.J. Kemp).

0304-3800/\$ – see front matter © 2007 Elsevier B.V. All rights reserved.

doi:10.1016/j.ecolmodel.2007.01.009

models, without information regarding the relative importance of parameters in the system, their utility is limited. Development of a method to couple input factors to meaningful output in ANN models is thus of critical importance. Accurate knowledge of the relative importance of model input parameters in ANN models in producing output would be useful in guiding conservation policy, monitoring and sampling strategies, and formulation of testable scientific hypotheses. ANNs may be useful in identifying ecological correlations, but without a way to identify the contributions of input parameters to ANN model output, it is impossible to identify putative relationships which could be investigated through further study.

Several authors have reviewed the effectiveness of various methods of estimating the relative importance of input parameters in ANN models (Gevrey et al., 2003; Olden et al., 2004). Olden et al. (2004) compared a number of methods to determine the importance of parameters in ANNs using a simulation-based approach. The approaches reviewed by these authors ranged from calculation of indices, to sensitivity analysis, to variations on stepwise parametric analysis. In the simulation-based approach, results from the various methods were compared against known relationships in the simulated data set to test for method effectiveness. Their finding suggests that their newly proposed connection weight method outperforms all other methods in determining relative parameter importance in ANNs. The connectivity weight method explicitly utilizes fit parameters for the processing of data through the ANN model to determine relative parameter importance.

We present an alternative method of determining parameter importance for ANNs and demonstrate that it performs as well as the connection weight method. Moreover, the proposed method operates through the manipulation of input data and thus, no knowledge about fit parameters is required. The proposed method can be used in conjunction with commercially available ANN simulators, whereas other proposed methods cannot be used because the architecture of the ANN models developed is either unknown or ambiguous. The Hold-back Input Randomization Method (HIPR-Method) randomizes systemic input parameter data sequentially, and compares the resulting change in the accuracy of the ANNs prediction of the output from unaltered input data. Through repetition, confidence intervals can be placed on the relative importance of the input parameters. A freely available software package for implementing this method using multivariate data sets is released in conjunction with this paper.

2. Materials and methods

2.1. ANN details and data preparation

The ANN type used for this analysis was a fully connected, feed forward network with a single input layer consisting of four input neurons, a single hidden layer with five hidden neurons and a single output layer with one output neuron. In addition, the input and hidden layer had a bias neuron feeding a continuously 'on' signal to each neuron in those layers. The sigmoidal activation function was of the form $f = 1/(1 + e^{-w})$,

and the transfer function for the output layer was the identity function. The type of ANN used in this experiment, and the error backpropagation training method employed are well described in Lek et al. (1996) and Olden and Jackson (2002). Training of the ANN model continued until one of the following criteria were met: (1) if the mean squared error (MSE) of the validation data set continued to increase for five straight epochs or (2) the MSE of the validation set was reduced by less than 0.1% for five continuous epochs. Learning rate η was 0.1.

The data set was divided randomly into three parts, a training, a validation, and a test data set, following a 3:1:1 ratio. After all data points in the training set were presented once in random order (termed an epoch), each data point of the validation set was also presented once in random order, in order to evaluate the general accuracy of the ANN model. The mean squared error (MSE) of the target values of the validation set versus the ANN model output values was calculated. During repeated epochs, the MSE of the training set declined monotonically as well as the validation data set. After a certain degree of training, however, the MSE of the validation data set will generally begin to increase. In this case, the ANN model is overtrained, as it is adjusting to specific characteristics of the training data set and not the overall patterns present in the data. Training was halted when the MSE of the ANN output using the validation data set began to increase. The test data set was then used to calculate the general effectiveness of the ANN model by calculating R^2 and MSE of the ANN model output versus the target values.

2.1.1. HIPR method

Our proposed method is a refinement of an approach presented by Scardi and Harding (1999), referred to as the 'perturbation' method by Olden et al. (2004), to achieve understanding of relative importance of input variables by systematically altering input data patterns. The HIPR method basically consists of an experimental test of an ANN model using data sets which have been altered in a way to deduce the relative importance of the input parameters in predicting the outcome. The procedure to perform the HIPR method can be summarized as follows:

- (1) Optimization of ANN model (see above)
- (2) Using the test data set to determine relative input parameter importance:
 - a. sequentially feeding each data point in the test data set to the ANN but replacing the values of one input parameter by uniformly distributed random values in the interval (0.1, 0.9), the range over which the net was originally trained,
 - b. calculating the mean squared error of the ANN when the randomized test set has been presented, and
 - c. repeating the procedure for each input parameter, each time substituting the original values with uniformly distributed random values.

The MSE values of the data set with a particular randomized parameter in relation to the MSE of the original data set reflect the relative importance of that input parameter for the prediction of the ANN. If a parameter does not contribute to the predictive ability of the ANN, the MSE of the randomized

data set will be close to the MSE of the original data set. If a parameter contributes strongly to the predictive ability of the ANN, the MSE of the data set in which this parameter is randomized will be greater than the MSE of the original data set.

Some similarities exist between this method and the stepwise methods reviewed by Olden et al. (2004). However, several critical differences exist. The HIPR approach does not involve the removal of parameters from inclusion in the network, does not hold parameters at mean values over all data sets, and operates conceptually on repetitions of the ANN training process. These differences are important to the improvement of the HIPR method over the other stepwise methods, and account for the basic ability of this method to elucidate relative parameter importance in ANN models.

2.1.2. Repetitions

For the HIPR method, the original values of a data set are replaced by random values. Thus, the actual outcome is a stochastic event based on sampling error at the data partitioning phase of the data and repetitions are needed. For each repetition, a new ANN is trained and new partitions of the original data set are used. Upon completion of a number of repetitions, relative parameter importance can be determined by ranking the mean MSE produced by parameter, and examining overlap between parameters to determine significance of the differences (a lack of overlap between one distribution and 95% of another distribution indicating a significant difference at a level of $\alpha = 0.05$). Because the distribution of variance measures such as MSE is not defined, the assumption of normality of distribution necessary to implement commonly used statistical tests such as ANOVA is not met (Sokal and Rohlf, 1981). Using the method of overlap of data to determine statistical significance is conservative relative to commonly used tests such as ANOVA, and therefore represents a viable alternative which avoids any statistical pitfalls related to assumptions.

2.1.3. Connection weight method

The connection weight method (Olden et al., 2004) sums the product of the weight of the connection from input neuron to the hidden neurons with the weight of the connection from the hidden neurons to the output neurons for all input parameters. The larger the sum of connection weights, the greater the importance of the parameter that is associated with this input neuron. The relative importance of input parameter i is determined through the following formula:

$$\text{Imp}(i) = \sum_{x=1}^n (\text{CW}_{ih(x)} \text{CW}_{ho(x)})$$

where $\text{Imp}(i)$ is the relative importance of parameter i ; n the total number of hidden nodes; x the index number of hidden node; $\text{CW}_{ih(x)}$ the connectivity weight between input parameter i and hidden node x ; $\text{CW}_{ho(x)}$ is the connectivity weight between hidden node x and the output node.

2.1.4. Simulated data sets

The HIPR method was tested through application to simulated data sets. In the simulated population data sets, all

functional relationships were known, and results from the HIPR method were expected to reflect a reasonable interpretation of relative importance of input variables. This has been termed the virtual ecology approach, the merits of which have been described by Grimm et al. (1999). The use of standard simulation-based approaches to assess efficacy of models is also widely used in statistics (Hastie et al., 2001). A number of data sets were created where the output variable y showed different mathematical dependency to the input parameters x_1 through x_4 . Parameters x_1 through x_4 were generated as uniformly distributed random numbers over the range 1–10,000, and were then fed into four fundamentally different functional relationships in the production of an output. For each correlation, a set of 10,000 patterns was created. All data sets were rescaled to the interval (0.1, 0.9) before being presented to the ANN model. The following four functional relationships were used:

- Multiple linear relationship

$$y = a_1x_1 + a_2x_2 + a_3x_3 + a_4x_4 \quad \text{with } a_1 = 0; a_2 = 0.2; \\ a_3 = 0.4; a_4 = 0.8.$$

- Multiple linear relationship with random variation

$$y = a_1x_1 + a_2x_2 + a_3x_3 + a_4x_4 + \text{random}(10,000)$$

with $a_1 = 0; a_2 = 0.2; a_3 = 0.4; a_4 = 0.8$, and $\text{random}(10,000)$ represents a uniformly distributed random number from 1 to 10,000.

- Logarithmic relationship

$$y = \log(e^{x_3^{0.001}} + e^{x_4^{0.01}}).$$

- Collinear relationship

$$y = x_3x_4.$$

2.1.5. Simulations

For each of the four data sets, a series of 99 repetitions of data set partitioning and ANN optimization were enacted. For each repetition, the HIPR method and the connectivity weight method were implemented to determine relative parameter importance. These results were then used to determine the absolute and relative accuracy of both of these methods.

All simulations were implemented using the C++ programming language in the Borland Builder 6 environment. An executable version of this simulation is available at <http://www.bio.upenn.edu/faculty/dunham/hipr/PennNN.zip>.

2.2. Mangrove impoundment case study

In order to demonstrate the application of the HIPR method to actual data, a data set collected in mangrove marsh areas by one of us (Kemp, unpublished dissertation) in East central Florida was used as a case study for the HIPR method. The data set consisted of physical characteristics of sampling stations in mangrove marsh areas under different management strategy for mosquito control. Details of sampling strategy and methodology are presented in Appendix A.

To examine the relative effects of impoundment on water quality factors, physical characteristics of different stations

were incorporated into a model with the management strategy of the impoundment as an output variable. Model effectiveness was evaluated as the ability to predict management strategy of a particular station in a randomly determined validation data set. A total of 203 patterns from nine sites were available for analysis. Thus, as compared to the computer simulated data, in this ecological data set the amount of information was much more restricted. Variables included in the model were as follows: year of sampling (year), Julian date of sampling (JD), surface temperature at station (Temp), surface dissolved oxygen in parts per million (O2), surface salinity in parts per thousand (Sal), surface sulfide level in parts per million (Sulf), per meter stratification in temperature at station (T.strat), stratification in dissolved oxygen at station (O2.strat), per meter stratification in sulfide level at station (Sulf.strat), and per meter stratification in salinity at station (Sal.strat). Model output was management strategy (categorical: open, managed, or closed). The data was otherwise prepared in an identical manner as the simulation data, and the details of the ANN models trained were also identical. A total of 157 repetitions of the method were usable for the analysis.

3. Results

Performance of the HIPR and connectivity weight methods was compared by determining relative parameter importance in ANNs trained on the simulated data sets, and comparing these values with known parameter importance in the functional relationship used to construct the data set. The two techniques used different criteria to accomplish this. The HIPR method used the mean squared error resulting from the randomization of each parameter, and the connectivity weight method quantified the total of the connection weights from input to output for each parameter. Following are the results of the two methods for each of the four data sets.

3.1. Multiple linear relationship

ANNs trained on the data set constructed according to a multiple linear relationship had an average R^2 value of 0.997 (S.D. = 0.000375; $N=99$) for ANN model output versus target (actual) values. Of the four simulated data sets, the known relative parameter importance was the most interpretable for this relationship. The relative importance of input parameters was determined using the connectivity weight method and the HIPR method ($N=99$ separate resamplings of data set and resulting trained ANN models). Fig. 1 shows the relative importance estimates for the input parameters in the ANNs using both of these methods. Both the HIPR (Fig. 1A) and the connectivity weight (Fig. 1B) method were successful in the identification of parameter importance ($x_4 > x_3 > x_2 > x_1$). Neither method ascribed significant importance to x_1 , a variable with no functional relationship to the output. However, the HIPR method appeared to ascribe relative importance in a nonlinear rather than linear fashion. The connectivity weight method resulted in more outliers, and greater apparent standard deviations versus the HIPR method, although variations between these methods were not directly comparable.

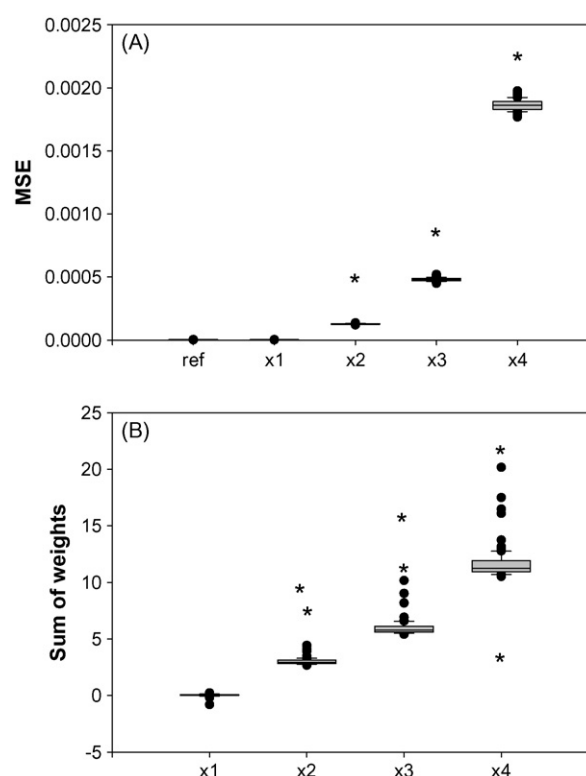


Fig. 1 – Multiple linear correlation. (A) HIPR method: Relative importance of input parameters x_1 – x_4 as indicated by MSE of test data set and modified test data sets. (B) Relative importance of parameters (x_1 – x_4) as indicated by the connectivity weight method (see text for details) (box plot: upper and lower boundaries of grey areas represent 75th and 25th percentile, whiskers above and below plot are 90th and 10th percentile; points are outliers. An asterisk indicates a significant difference from either the unaltered test data set (HIPR method) or zero (connectivity weight method) as shown by lack of overlap.

Variation around the mean MSE value obtained by the HIPR method was narrower than the confidence intervals around the sum of connectivity weights produced by the connectivity weight method. For parameters x_2 , x_3 , and x_4 , the HIPR method found MSE values produced by randomizing each one of these parameters to be significantly higher than MSE values resulting from unaltered data set. This indicates a loss of information used by the ANN model, and therefore, parameters x_2 , x_3 , and x_4 all significantly contribute to the predictive ability of the ANN model. Likewise, the connectivity weight method showed that connectivity weight values for parameters x_2 , x_3 , and x_4 were all significantly greater than zero. Correctly, neither method showed x_1 to be of significant importance in predicting the output. In summation, both models performed similarly in terms of determining parameter importance and significance. Differences between the two methods included the exponential difference in importance indicated by the HIPR method as opposed to the connectivity weight method, and the between-run consistency in the mean squared error in the HIPR method versus the connectivity weight method.

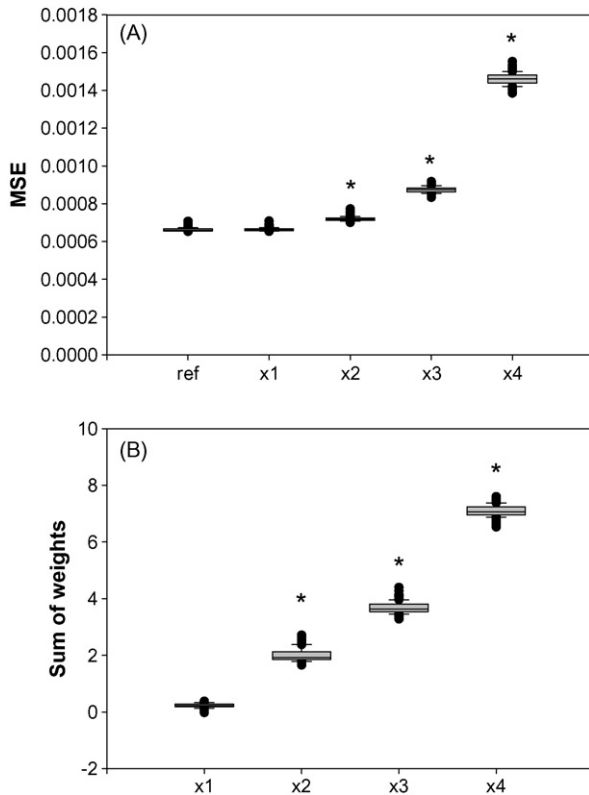


Fig. 2 – Multiple linear correlation with uniform random variation. (A) HIPR method: Relative importance of input parameters x_1 – x_4 as indicated by MSE of test data set and modified test data sets. (B) Relative importance of parameters (x_1 – x_4) as indicated by the connectivity weight method (see text for details) (box plot: see Fig. 1 for details).

3.2. Multiple linear relationship with random variation

The addition of uniformly distributed random variation to the multiple linear relationship substantially reduced the ability of ANNs to predict target values. The mean R^2 value for output versus target (actual) values from ANNs trained on this type of data set was 0.449 (S.D. = 0.0092; $N = 99$). The relative importance of input parameters as determined through these methods is shown in Fig. 2. Since R^2 values produced by ANN models for these data sets were lower than for other simulated data sets, it is important to note a particular aspect of the results. ANN models vary in predictive ability according to the relationship between variables in the system. When there is little correlation between input and output parameters, the relative parameter importance measured by HIRP only reflects the partial contribution of the parameter to the predictive ability of that particular ANN model.

This data set involved a duplicate of the underlying multiple linear relationship of the previous data set, masked by an addition of uniformly distributed random values. Results mirrored those derived from the multiple linear relationship data set, with HIRP again appearing to exaggerate dependence of the model upon the most important parameters (Fig. 2). Both methods correctly identified the relative importance of the

input parameters and their significance, despite the reduced predictive ability of the ANNs. The differences in outliers and confidence interval width for the two methods was not as great for ANN models derived from this data set as for the multiple linear relationship without noise, although there appeared to be more variation around the mean per parameter for the connectivity weight method (Fig. 2B).

3.3. Logarithmic relationship

The nonlinear relationship expressed in these data sets resulted in highly consistent ANN models with a mean R^2 value for model output versus target (actual) values of 0.997 (S.D. = 0.0014; $N = 99$). The relative importance of the parameters in ANNs trained on this relationship is shown in Fig. 3. Neither method identified parameters x_1 or x_2 as significant, which was expected because these parameters were not included in the nonlinear relationship. However, while parameter x_3 was explicitly included in the nonlinear relationship, neither method identified this parameter as being significant. The only important parameter identified by both methods was x_4 , which was multiplied by a factor an order of magnitude greater than x_3 (0.01 versus 0.001). The nonlinear character of the relationship coupled with the difference in the magnitude of the exponents producing this data set most likely accounted

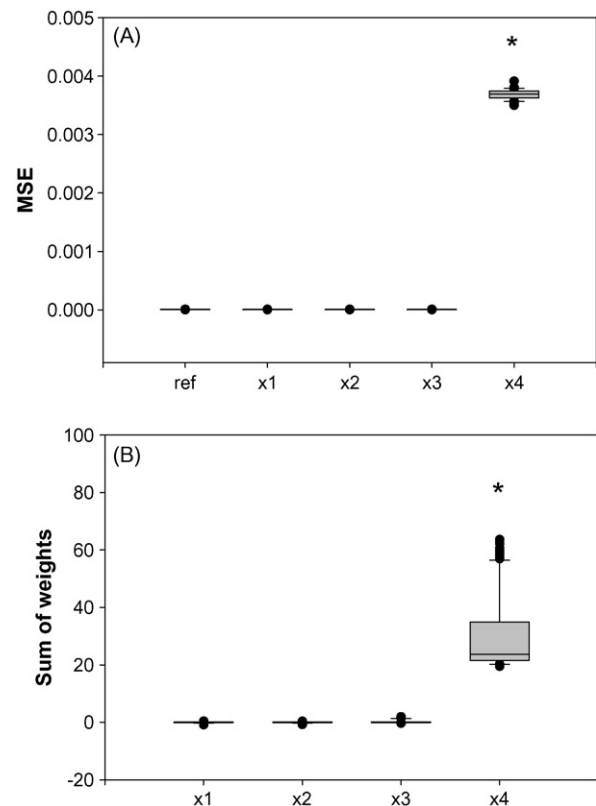


Fig. 3 – Logarithmic (nonlinear) correlation. (A) HIRP method: Relative importance of input parameters x_1 – x_4 as indicated by MSE of test data set and modified test data sets. (B) Relative importance of parameters (x_1 – x_4) as indicated by the connectivity weight method (see text for details) (box plot: see Fig. 1 for details).

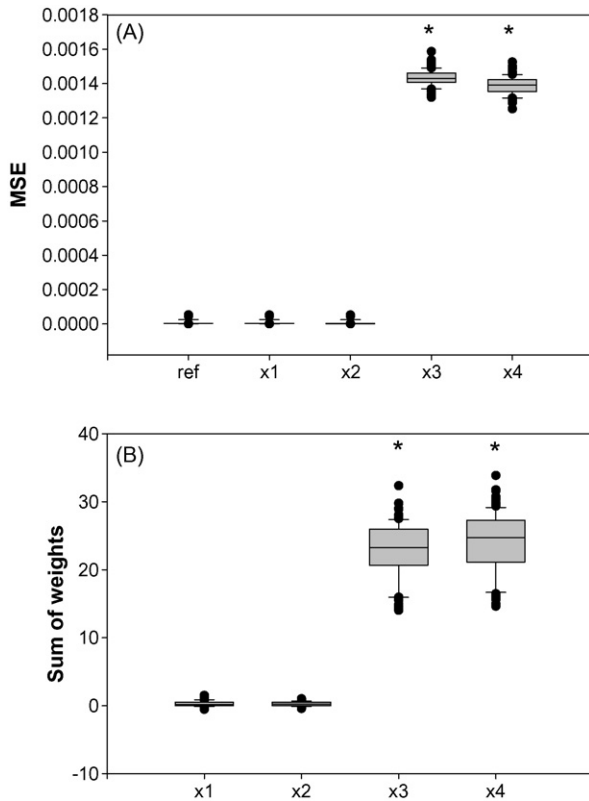


Fig. 4 – Multiple co-linear relationship. (A) HIPR method: Relative importance of input parameters x_1 – x_4 as indicated by MSE of test data set and modified test data sets. (B) Relative importance of parameters (x_1 – x_4) as indicated by the connectivity weight method (see text for details) (box plot: see Fig. 1 for details).

for the lack of significance of x_3 . In the connectivity weight method there was extensive variation around the mean connectivity weight sum for parameter x_4 (Fig. 3B), while there was very little variation around the mean MSE for the same parameter produced by the HIPR method (Fig. 3A).

3.4. Co-linear relationship

ANNs produced from data sets with this relationship had a relatively high R^2 value for output versus target (actual) values of 0.994 (S.D.=0.012; $N=99$). Because this relationship involved a simple multiplication of parameters x_3 and x_4 with no additional factorization, the identification of the functional equivalence of these two parameters was expected. Fig. 4 shows the results of the two methods. In both methods, x_3 and x_4 were identified as being equal and significant, and x_1 and x_2 were identified as not significant. As in previous examples, within parameter between-run variation was much greater in the results from the connectivity weight method than from the HIPR method (Fig. 4).

The HIPR method proposed in this paper and the connectivity weight method proposed by Olden et al. (2004) performed equally well in predicting relative input parameter importance and significance over four data sets constructed according to

different functional relationships (linear, linear with noise, logarithmic, and co-linear). Main differences between the two methods appeared to be a greater between-run consistency in results produced by the HIPR method versus the connectivity weight method, dispersion around the mean values obtained in the connectivity weight method, and the tendency of the HIPR method to emphasize relative dependence on the most important parameters.

3.5. Mangrove impoundment case study

MSE values obtained in the HIPR analysis for each input parameter, ranked from lowest to highest, are shown in Fig. 5. Randomization of the signal in the oxygen parameter produced the largest MSE, and as such had the largest effect on the ability of the ANN models to predict management strategy from physical data. Using the method of overlap between parameters to determine statistical significance, only MSE values derived from randomizing oxygen, salinity, and stratification in salinity were significantly different from the baseline (validation data set). While the mean MSE values related to randomization of stratification in sulfur were similar to those from stratification in salinity, variation was too great to distinguish this variable from the validation data set. All other parameters had sufficient amounts of overlap with the validation data set to make them statistically indistinguishable from those produced by the validation data set. This illustrates the relative importance of parameters in ANN models related to management strategies in mangrove impoundments, as well as the significance of input parameters in determining the output. Because the ANN models were trained to differentiate between management strategies, it can be assumed that the parameters which were most important to the ANN models are those most consistently related to management strategy. The average R^2 value obtained over all 157 model runs of ANN models was 0.31.




Fig. 5 – Relative levels of average mean square error produced in HIPR method for multivariate water quality data set from impounded mangrove. Abbreviations are detailed in the text under Section 2. Error bars represent 95% confidence interval. An asterisk denotes a significant difference between MSE values at the $\alpha = 0.05$ level using the overlap method outlined in text.

4. Discussion

The HIPR method introduced here performed at a high level in identifying relative importance values of input parameters, for all types of data sets studied. In all cases, parameters which were functionally important in determining output were identified as significant, while those which were not important were not incorrectly assigned importance. When compared with the recently proposed connectivity weight method recently proposed by [Olden et al. \(2004\)](#), both methods performed acceptably. This research also represents an independent confirmation of the accuracy of the connectivity weight method. Some advantages of the connectivity weight method lie in its intuitive simplicity. However, we feel there are a number of distinct advantages to the HIPR method relating to its functional generality and its use of nonarbitrary criteria in determining relative parameter importance (MSE). Some of these are well shown by the ecological example provided.

The analysis of the mangrove ecosystem data demonstrates the practical use of the HIPR method in interpreting ecological data. This technique allows the combination of the great predictive power of ANN models ([Hastie et al., 2001](#)), with the ability to rank the relative importance of different parameters within the ANN model based on their contribution to predictive power. Also, a procedure can be employed to determine their significance. This combination of predictive power of ANN models, ranking of importance of input parameters, and determination of parameter significance is a primary benefit of the HIPR method to ecologists analyzing multivariate data sets. The example provided here illustrates both these benefits and some of the caveats of using this method. This data set, somewhat limited and noisy as compared to the abundant data provided by computer simulations, is much more along the lines of the kind of data set likely to be available on ecological systems, in contrast to the abundant data provided by computer simulations. Employed here, the simulation data evaluate the effectiveness of the HIPR technique, while the ecological data represent its application to a more realistic data set. The significant traits used by ANN models to distinguish between management strategies in mangrove impoundments as shown by the HIPR method were salinity and stratification in salinity, and most important, oxygen. Other factors introduced as input variables did not significantly affect the predictive ability of the ANN models. Order of importance, as well as variation in importance between different ANN model constructions, is easily observed in this example. When examining limited ecological data, the choice of a model validation data set could conceivably affect the result of the ANN model. In the present example, this is shown in the high degree of variation seen between model runs for stratification in sulfide ([Fig. 5](#)). While the mean MSEs produced by this parameter were similar to a significant factor, stratification in salinity, variation in the importance of this parameter between runs as indicated by variation in MSE precluded it from being significant. Restated, the importance of this factor depended on the particular choice of the validation set. The repeated resampling of the data implicit in the HIPR method prevents a spurious attribution of importance to a variable owing to chance factors alone. Finally, this example also shows

a caveat of interpreting results from this and any other ANN or multivariate model. This is mainly that the overall importance of these parameters in the system determined by the HIPR method must be constrained by the predictive power of the ANN models used in the first place. The R^2 values obtained for this particular natural system indicate that much of the water quality differences between impoundment sites is not readily explained by the ANN models produced. Other factors not included in the model, such as site-specific differences, are also undoubtedly contributing to the variation between sampling stations. The greater the predictive power of the ANN model, the greater weight which should be assigned to factors in the model, and vice versa. However, the example shown demonstrates how the HIPR method can be used and the results interpreted to achieve a better understanding of ecological systems through ANN models.

The criteria by which HIPR and the connectivity weight methods judge the relative importance of parameters marks a key difference in approach between the two methods. Whereas the HIPR method utilizes MSE produced by randomizing input parameters in ANN models, the connectivity weight method uses the connectivity weights of the ANN model. This may explain some of the differences observed in the performance of the two methods. Because the connectivity weight method relies on the unitless connectivity weights produced during the training of the ANN model, the variation in connectivity weight values between runs is sensitive to the initial random values of the weights relative to the HIPR method ([Figs. 1, 3 and 4](#)). In the connectivity weight method, the between-run variation in the sums of connectivity weights depends on their initial random values in the model before training, and thus the variation in these quantities observed between runs is arbitrary. Consequently, the arbitrary nature of the connectivity weight values clouds interpretation of the results between sampling runs. The HIPR method relies on MSE, which is fully comparable between sampling runs. Greater consistency in results was observed for the HIPR method and this indicates its greater accuracy and repeatability in determining relative parameter importance. In data sets where the number of available data patterns is limited (as is often the case with ecological data sets), the need to repeat sampling of validation and test data sets and train concurrent ANN models becomes much greater to avoid the problems of sampling error ([Lek et al., 1996](#)). In such a situation, the standardized MSE criterion decidedly makes the HIPR method the tool of preference.

Because the HIPR method is a quantification of MSE residuals obtained when system parameters are randomized, the method should perform well over a diverse range of ANN applications. If more than one output is present in the model, then relative input parameter importance will be determined by the MSE summed over all output nodes when that particular parameter is randomized. For input or output nodes consisting of binary data, randomization would proceed as described for the method. Randomization would produce many values in between the "0"s and "1"s of the binary data set, but it would still meet the main requirements of the HIPR method, i.e. the uncoupling of the data interrelationships found in that particular parameter. Following, it would then be possible to examine the loss in predictive ability of the ANN

resulting from randomization of that parameter. In the case of classification type data, a similar coding of categorical data into a continuous output would allow for the quantification of MSE values occurring when each of the parameters is randomized. For example, let a hypothetical data set consist of a multivariate parameter space characterizing three different states (output). In order to train the model, these states would be coded in the data patterns as outputs in the form of a continuous variable, with the states represented as equidistant points on the scale. If all variables are to be rescaled over the range (0.1, 0.9), then the three different output states would be represented as 0.1, 0.5, and 0.9. The MSEs of the ANN model for target versus observed values could then be calculated as in other scenarios, and the relative importance of the input parameters in accurately classifying the output patterns could be determined. The use of MSEs to determine relative parameter importance allows for much flexibility in the application of the HIPR method to various types of data sets and ANN models.

A key issue regarding the interpretation of HIPR results is the situation where the ANN used has relatively low predictive power. We can expect this situation in many ecological data sets. As these systems are often very complex, it is highly unlikely all relevant system variables will be included within a single model. Interpretation of the results must always be done within the context of the predictive power of the ANN. A high importance attributed to a particular input variable only represents a relative partial contribution of predictive power of the ANN model, and any extrapolation back to its true importance in the system must be constrained by it. However, even in situations where the predictive power of the ANN is low, the use of the HIPR method can provide useful insight into important components of the system. Similar to a multivariate analysis (i.e. discriminant analysis), the ANN training procedure identifies patterns of variation which provide the strongest predictive power. If other significant, but lesser axes of variation are orthogonal to the primary variation represented by the ANN, they will not be represented in the resulting ANN model. Therefore, the importance of any parameter identified by HIPR in a modeled system must be constrained by the overall predictive power of the ANN model from which it is derived.

One interesting difference between the two methods lies in the exponential appearance of relative parameter importance produced by the HIPR method. This may be due to the fundamental difference in criteria used by this method to evaluate relative parameter importance. Because this procedure uses MSE, linear relationships in parameter importance of input parameters appear to be exponential when the HIPR method is performed. The absolute order does not change, and it would be possible to determine the actual linear relationships between the input parameters by quantifying absolute errors produced rather than the sum of squared error. While absolute residuals could be used in our method, using MSE instead may serve to accentuate differences between measured parameter importances and thus provide better resolution as to relative parameter importance. Moreover, MSE is the primary criterion used to optimize ANN models and may therefore be more representative of the relative importance of an input parameter to network predictive ability.

The repeated resampling of test/validation data sets from the original data sets in the HIPR method and the multiple optimizations of ANN models is a superior approach for determining relative importance estimates, especially when data are limited. First, the repeated nature of the process allows for the placement of confidence intervals on the relative importance of input parameters, and allows for the objective determination of significance of input parameters. In this paper, this is done using a determination of the amount of overlap between two distributions. Repeated resampling of test data sets reduces the potential for sampling error caused by choice of a validation or test data set (Lek et al., 1996). While sampling error is not likely when extremely large, functionally homogenous data sets are used (such as in this study), it is likely to be a factor when used for actual ecological data sets, where data availability is limited. Repeated random selection of validation data sets enhances the probability that actual patterns in the data being modeled with ANNs are accurately represented in the HIPR method.

The HIPR method is also capable of assessing the relative importance of input parameters in any conceivable configuration of ANN. While the architecture used by most researchers is very simple (e.g. a feed forward network with a single layer of hidden nodes), a great number of complex architectures of ANN are possible. In many cases, these more complicated architectures outperform ANNs with simpler architectures (Hastie et al., 2001). The HIPR method requires fewer assumptions regarding the architecture of the network due to the criterion of relative contribution of input parameters in predicting target values. This may be a key advantage in using this method over the connectivity weight method in situations where summation of connectivity weights would not be a logical way to assess relative parameter importance, or if the details of the derived ANN are unavailable (as in some statistical packages). The HIPR method can easily be implemented in standard ANN simulators by creating multiple data sets in which one parameter at a time is replaced by uniformly distributed random values. This can be accomplished via computer program, or even by using a spreadsheet program.

The HIPR approach has some other important differences with the ‘perturbation’ method (Scardi and Harding, 1999) evaluated by Olden et al. (2004), which should be noted here. The randomization of data contained in different input parameters constrains the values so they do not exceed the numerical scope of the trained network, i.e. over the range [0.1, 0.9]. Because the ANN is trained over a given range of values, the overextension of values fed to the network through the addition of white noise may result in an unpredictable or irrelevant outcome. The HIPR method controls for the potential for the randomized values to produce this undesired effect by constraining the values to a range ‘native’ to the ANN.

With regard to the development of ANN models as a useful tool for scientific investigations beyond predicting outcomes from a series of inputs to the models, this study has some important broader implications. The prevailing hypothesis that ANN models operate in an opaque, or even mysterious manner, and that these properties inhibit the gaining of valuable insights into the system is not supported by these results. ANN models derive their predictive power from actual

patterns of correlation present in the data. As shown in this paper, they do so in a manner which is consistent over many repetitions, and the ANN models developed attribute relative weights of importance proportional to the actual patterns of importance present in the system. The research presented here, along with that of other authors, should act to dispel the notion that ANNs are only useful in classifying patterns in data, and supports the idea that ANNs are useful in conjunction with more detailed, even mechanistic studies as a preliminary identifier of putative ecological correlations.

In summary, the HIPR method represents a robust, general tool ideal for examining relative parameter importance. This method may be especially useful in application to multivariate data sets which are of a size sufficient to train an ANN model but limited in extent, such as ecological data sets. This procedure provides the greatest likelihood of identifying relative contribution of parameters to ecological relationships using ANN models. This method should be seen as complementary to existing methods while at the same time having distinct advantages from the perspective of heuristics, generality, and interpretability.

Acknowledgements

Comments from Michael O'Connor, Arthur Dunham, James McNair, and an anonymous reviewer greatly improved the manuscript. Stanley Kemp would like to thank University of Pennsylvania and Ursinus College for support. Patricia Zaradic was supported as a D.H. Smith Conservation Fellow through the Nature Conservancy, and F.H. was partially supported by a cooperative agreement from the United States Department of Agriculture's National Wildlife Research Center (NWRC) (03-7408-0479-CA) to the University of Pennsylvania. We thank NWRC's Wildlife Disease Project: Surveillance, Monitoring, and Response for its continued support of this work.

Appendix A

Impounded mangrove data set: The sites at which data were collected represented a cross section of types of mangrove areas along the Indian River, Florida. Sites were visited regularly over the time period June 2000–August 2002. At each site, several (up to five if possible) sampling stations were established. It was not always possible to obtain five sampling stations at every site due to limitations in the amount of open, standing water and access. On every visit, these stations were examined for a number of variables, including temperature, dissolved oxygen, salinity, and sulfide levels. These variables represented those most likely to be affected by impoundment of the marsh. In addition, a bottom water sample was obtained using a drop bottle. Bottom water was examined for the same variables as surface waters for the purpose of obtaining levels of stratification. Temperature was measured using a 12 gauge copper-constantan calibrated thermocouple dipped in

thermally conductive silicone. Temperature in degrees Celsius was read to the nearest 0.1° with a handheld digital thermocouple reader (Omega model HH-25TC). Dissolved oxygen was measured using a Sentry analog dissolved oxygen meter and read to the nearest 0.1 ppm. Salinity was measured using a handheld VeeGee salinity refractometer to the nearest dimensionless unit (ppt). Finally, sulfide levels were measured using a LaMotte chemical sulfide test kit (model P-70) and were read to the nearest ppm. Stratification in these variables at each sampling station was determined by subtracting the surface reading from the bottom reading for each variable, and then standardized for depth of the station in meters to obtain stratification at the station as (units/m). Stratification was thought to be important in the system because a primary effect of impoundment is to prevent mixing of the water column.

REFERENCES

- Basheer, I.A., Hajmeer, M., 2000. Artificial neural networks: [fundamentals, computing, design, and application](#). *J. Microbiol. Methods* 43, 3–31.
- Gevrey, M., Dimopoulos, I., Lek, S., 2003. Review and comparison of methods to study the contribution of variables in artificial neural network models. *Ecol. Model.* 160, 249–264.
- Grimm, V., Wyszomirski, T., Ackman, D., Uchmanski, J., 1999. Individual-based modeling and ecological theory: synthesis of a workshop. *Ecol. Model.* 115, 275–282.
- Guégan, J.F., Lek, S., Oberdorff, T., 1998. Energy availability and habitat heterogeneity predict global riverine fish diversity. *Nature* 391, 284–382.
- Hastie, T., Tibshirani, R., Friedman, J., 2001. *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. Springer, New York.
- Lek, S., Delacoste, M., Baran, P., Dimopoulos, I., Lauga, J., Aulagnier, S., 1996. Application of neural networks to modeling nonlinear relationships in ecology. *Ecol. Model.* 90, 39–52.
- Lek, S., Guégan, J.F. (Eds.), 2000. *Artificial Neuronal Networks. Applications to Ecology and Evolution*. Springer-Verlag, New York, USA.
- Olden, J.D., Jackson, D.A., 2002. Illuminating the black box: a randomization approach for understanding variable contributions in artificial neural networks. *Ecol. Model.* 154, 135–150.
- Olden, J.D., Joy, M.K., Death, R.G., 2004. An accurate comparison of methods for quantifying variable importance in artificial neural networks using simulated data. *Ecol. Model.* 178, 389–397.
- Scardi, M., Harding Jr., L.W., 1999. Developing an empirical model of phytoplankton primary production: a neural network case study. *Ecol. Model.* 120, 213–223.
- Sokal, R.R., Rohlf, F.J., 1981. *Biometry*. W.H. Freeman and Co., New York.
- Vander Zanden, M.J., Olden, J.D., Thornse, J.H., Mandrak, N.E., 2004. Predicting the occurrences and impacts of smallmouth bass introductions in north temperate lakes. *Ecol. Appl.* 14, 132–148.
- Venables, W.N., Ripley, B.D., 2002. *Modern Applied Statistics with S*, fourth ed. Springer, New York.