Available online at www.sciencedirect.com

SCIENCE DIRECT®

Ecological Modelling 178 (2004) 389–397

ECOLOGICAL MODELLING

www.elsevier.com/locate/ecolmodel

An accurate comparison of methods for quantifying variable importance in artificial neural networks using simulated data

Julian D. Olden [a,*], Michael K. Joy [b], Russell G. Death [b]

[a] *Graduate Degree Program in Ecology, Department of Biology, Colorado State University, Fort Collins, CO, 80523, USA*
[b] *Institute of Natural Resources—Ecology, Massey University, Private Bag 11 222, Palmerston North, New Zealand*

Received 21 July 2003; received in revised form 23 February 2004; accepted 15 March 2004

Abstract

Artificial neural networks (ANNs) are receiving greater attention in the ecological sciences as a powerful statistical modeling technique; however, they have also been labeled a "black box" because they are believed to provide little explanatory insight into the contributions of the independent variables in the prediction process. A recent paper published in *Ecological Modelling* [Review and comparison of methods to study the contribution of variables in artificial neural network models, Ecol. Model. 160 (2003) 249–264] addressed this concern by providing a comprehensive comparison of eight different methodologies for estimating variable importance in neural networks that are commonly used in ecology. Unfortunately, comparisons of the different methodologies were based on an empirical dataset, which precludes the ability to establish generalizations regarding the *true accuracy and precision* of the different approaches because the *true importance of the variables* is unknown. Here, we provide a more appropriate comparison of the different methodologies by using Monte Carlo simulations with data exhibiting defined (and consequently known) numeric relationships. Our results show that a Connection Weight Approach that uses raw input-hidden and hidden-output connection weights in the neural network provides the best methodology for accurately quantifying variable importance and should be favored over the other approaches commonly used in the ecological literature. Average similarity between true and estimated ranked variable importance using this approach was 0.92, whereas, similarity coefficients ranged between 0.28 and 0.74 for the other approaches. Furthermore, the Connection Weight Approach was the only method that consistently identified the correct ranked importance of all predictor variables, whereas, the other methods either only identified the first few important variables in the network or no variables at all. The most notably result was that Garson's Algorithm was the poorest performing approach, yet is the most commonly used in the ecological literature. In conclusion, this study provides a robust comparison of different methodologies for assessing variable importance in neural networks that can be generalized to other data and from which valid recommendations can be made for future studies.

Keywords: Statistical models; Explanatory power; Connection weights; Garson's algorithm; Sensitivity analysis

1. Introduction

The ability of the human brain to perform complex tasks, such as pattern recognition, has motivated a large body of research exploring the computational capabilities of highly connected networks of relatively simple elements called artificial neural networks (ANNs). Although ANNs were initially developed to better understand how the mammalian brain functions, in recent years researchers in a variety of scientific disciplines have become more interested in the potential mathematical utility of neural network

* Corresponding author. Tel.: +1-970-491-2329; fax: +1-970-491-0649.
E-mail address: olden@lamar.colostate.edu (J.D. Olden).

0304-3800/$ – see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.ecolmodel.2004.03.013

algorithms for addressing an array of problems. In the ecological sciences, for example, ANNs have shown great promise for tackling complex pattern recognition problems (Lek and Guégan, 2000; Recknagel, 2003) such as the modeling of biological variables as a function of multiple descriptors of the environment (e.g. Lek et al., 1996b; Özesmi and Özesmi, 1999; Olden and Jackson, 2001; Joy and Death, 2004; Vander Zanden et al., 2004; are just a few examples). In these studies, and many others in the literature, the predictive advantages of ANNs over traditional modeling approaches continue to be illustrated.

Beyond the predictive realm of ANNs, focus has now switched to the development of methods for quantifying the explanatory contributions of the predictor variables in the network. This was, in part, prompted by the fact that ANNs were coined a "black box" approach to modeling ecological data. Our study focuses on the recent paper of Gevrey et al. (2003) published in *Ecological Modelling* that conducted a comprehensive comparison of eight different methodologies that have been widely used in the ecological literature. Although we commend the authors for providing such an important analysis, we feel the analytical approach they used resulted in an invalid comparison of the different methodologies. Consequently, we argue that the results of their paper fail to aid researchers in properly choosing among these different methods when using ANNs in their future studies. Our viewpoint is supported two-fold.

First, method comparisons by Gevrey et al. (2003) were based on an empirical dataset (relating the density of brown trout redds to a set of local habitat variables), which precludes the ability to establish any generalizations regarding the *true accuracy and precision* of the different approaches because the *true correlation structure* of the empirical data is unknown and, therefore, the *true importance of the variables* is unknown. The fact that this dataset has been analyzed in a number of other studies (e.g. Lek et al., 1996a, b) does not change the fact that the true correlative characteristics of the entire population (i.e. all brown trout redds and their associated habitat conditions in the entire river basin) remain unknown. As such, this data still only represents a sample from the entire population. Second, the authors erroneously assessed the stability of the methods by examining the variation in estimated variable importance based on the construction of multiple neural networks (10 in total) with the same empirical data. This procedure, however, does not assess method stability; rather, it assesses the differences among variable contributions arising solely from different initial connection weights used in network construction, that is, an issue of optimization stability (see Olden and Jackson (2002b) for further discussion). For the above reasons, we argue that the study of Gevrey et al. (2003) does not provide a valid comparison of the methods for assessing variable contributions in ANNs.

Prompted by the important methodological weaknesses of Gevrey et al. (2003), we provide a more appropriate comparison of nine different methodologies for assessing variable contributions in ANNs using simulated data exhibiting defined numeric relationships between a response variable and a set of predictor variables. Although rarely used in ecology, use of simulated data in comparative studies is preferable because the properties of the data are known (properties which cannot be determined, but only estimated from field data) and, thus, true differences among methods can be accurately assessed. Examples of simulation studies in ecology that compared different statistical methodologies include, but are not limited to, time-series analysis (Berryman and Turchin, 2001), ordination (Minchin, 1987; Jackson, 1993), regression analysis (Olden and Jackson, 2000), types I and II error analysis (Peres-Neto and Olden, 2001) and modeling species distributions (Olden and Jackson, 2001, 2002a) and ecosystem properties (Paruelo and Tomasel, 1997). The results from this study will demonstrate the true accuracy and precision of commonly used approaches for assessing variable contributions in ANNs, and, therefore, provides a robust comparison of the methodologies that can be generalized to other data and from which valid recommendations can be made for future studies.

2. Methods

2.1. Artificial neural networks

For the sake of brevity we refrain from discussing the details of the neural network methodology and instead refer the reader to the text of Bishop (1995) and the papers of Lek et al. (1996b) and Olden and Jackson (2001) for more comprehensive treatments. It is suf-

ficient to say that the methods described in this paper refer to the classic family of one hidden layer, feed-forward neural network trained by the backpropagation algorithm (Rumelhart et al., 1986). These neural networks are commonly used in ecological studies because they are suggested to be universal approximators of any continuous function (Hornik and White, 1989). The architecture of this network consists of single input, hidden and output layers, with each layer containing one or more neurons, in addition to bias neurons connected to the hidden and output layers. We determined the optimal number of neurons in the hidden layer of the network by comparing the performance of different cross-validated networks, with 1–25 hidden neurons, and choose the number that produced the greatest network performance (based on the entire Statistical Population: see Section 2.2). This resulted in a network with five input neurons (one for each of the predictor variables), five hidden neurons, and a single output neuron. Learning rate (η) and momentum (α) parameters (varying as a function of network error) were included during network training to ensure a high probability of global network convergence and a maximum of 1000 iterations for the backpropagation algorithm to determine the optimal axon weights.

2.2. Simulated data with known correlation structure

We compare nine methodologies for assessing variable contributions in artificial neural networks using a Monte Carlo simulation experiment with data exhibiting defined numeric relationships between a response variable and a set of predictor variables. By using simulated data with known properties, we can accurately investigate and compare the different approaches under deterministic conditions (known-error scenarios) and, therefore, provide a robust comparison of their performance. We generated a Statistical Population contained 10,000 observations for five predictor variables (x_1–x_5), each of which exhibit linear and incrementally lower (and positive) correlation with the response variable y ($r_{y\bullet x1} = 0.8, r_{y\bullet x2} = 0.6, r_{y\bullet x3} = 0.4, r_{y\bullet x4} = 0.2, r_{y\bullet x5} = 0.0$). Correlations between all predictor variables were set to 0.20 to represent a low level of multicollinearity that is likely typical of most ecological datasets. Preliminary analyses revealed that the level of multicollinearity in the simulated data had no effect on the results, and that the relative performance of the methods varied little when using non-linear simulated data (unpublished results). However, we feel the latter issue requires further investigation.

The Monte Carlo experiment consisted of randomly sampling 50 observations from the Statistical Population in order to provide a reasonable degree of generality (note that results were similar for larger sample sizes: unpublished results), constructing a neural network using the selected data, and calculating the explanatory importance of the five predictor variables using each of the different methodologies (see Section 2.3). This procedure was repeated 500 times and for each trial the predictor variables were ranked based on their estimated explanatory importance in the neural network. The degree of similarity between the estimated ranked importance of the variables (based on each of the nine methods) and the true ranked importance of the variables (i.e. $x_1 = 1$, $x_2 = 2$, $x_3 = 3$, $x_4 = 4$, $x_5 = 5$ based on the true Statistical Population) was assessed using Gower's coefficient of similarity for multi-state descriptors (Legendre and Legendre, 1998; p. 259). Because we applied each of the variable contribution methods to the same neural network constructed with the same random data the influence of network performance (i.e. its ability to model the data) is controlled for when making method comparisons.

2.3. Methods for quantifying variable importance in ANNs

Here we provide a brief description of the different methodologies for assessing variable contributions in neural networks examined in this study. These methods are used widely in the ecological literature and include those presented by Gevrey et al. (2003) and the Connection Weight Approach of Olden and Jackson (2002b). We refer the reader to these papers for more details.

2.3.1. Connection weights

Calculates the product of the raw input-hidden and hidden-output connection weights between each input neuron and output neuron and sums the products

across all hidden neurons (see Olden and Jackson, 2002b).

2.3.2. Garson's algorithm

Partitions hidden-output connection weights into components associated with each input neuron using absolute values of connection weights (see Garson, 1991). Called the 'Weights' method by Gevrey et al. (2003).

2.3.3. Partial derivatives

Computes the partial derivatives of the ANN output with respect to the input neurons (see Dimopoulos et al., 1995, 1999).

2.3.4. Input perturbation

Assesses the change in the mean square error of the network by adding a small amount of white noise to each input neuron (following Gevrey et al. (2003) we set the white noise to 50% of each input neuron), in turn, while holding all other input neurons at their observed values. The resulting change in mean square error for each Input Perturbation illustrates the relative importance of the predictor variables (see Scardi and Harding, 1999). Called the 'Perturb' method by Gevrey et al. (2003).

2.3.5. Sensitivity analysis

Involves varying each input variable across 12 data values delimiting 11 equal intervals over its entire range and holding all other variables constant at their minimum, 1st quartile, median, 3rd quartile and maximum (see Lek et al., 1996). The median predicted response value across the five summary statistics is calculated and the relative importance of each input variable is illustrated by the magnitude of its range of predicted response values (i.e. maximum–minimum). Called the 'Profile' method by Gevrey et al. (2003).

2.3.6. Forward stepwise addition

Assesses the change in the mean square error of the network by sequentially adding input neurons to the neural network (rebuilding the neural network at each step). The resulting change in mean square error for each variable addition illustrates the relative importance of the predictor variables (see Gevrey et al., 2003).

2.3.7. Backward stepwise elimination

Assesses the change in the mean square error of the network by sequentially removing input neurons from the neural network (rebuilding the neural network at each step). The resulting change in mean square error for each variable removal illustrates the relative importance of the predictor variables (see Gevrey et al., 2003).

2.3.8. Improved stepwise selection 1

Assesses the change in the mean square error of the network by sequentially removing input neurons (and their associated weights) from the neural network. The resulting change in mean square error for each variable removal illustrates the relative importance of the predictor variables (see Gevrey et al., 2003).

2.3.9. Improved stepwise selection 2

Assesses the change in the mean square error of the network by sequentially setting input neurons to their mean value. The resulting change in mean square error for holding each variable at its mean value illustrates the relative importance of the predictor variables (see Gevrey et al., 2003).

All Monte Carlo simulations and neural network analyses were conducted using computer macros in the MatLab® programing language written by the authors.

3. Results

The results from the Monte Carlo simulation experiments provide an accurate comparison of commonly used methodologies for quantifying variable contributions in artificial neural networks. The Connection Weight Approach was found to exhibit the best overall performance compared to the other approaches in terms of its accuracy (i.e. the degree of similarity between true and estimated variable ranks) and precision (i.e. the degree of variation in accuracy) of estimating the true importance of all the variables in the neural network (Fig. 1). Average similarity between true and estimated variable ranks using the Connection Weight Approach was 92%. Partial Derivatives, Input Perturbation, Sensitivity Analysis and Improved Stepwise Selections 1 and 2 approaches showed moderate performance (ca. 70% average similarity), whereas, Forward Stepwise Addition, Backward Stepwise

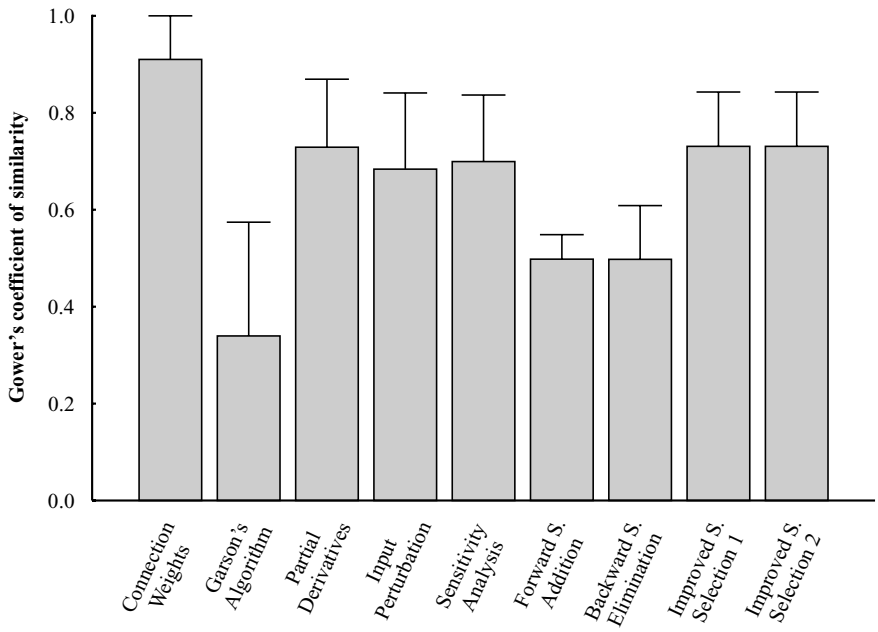

Fig. 1. Gower's coefficient of similarity between the true ranked importance and estimated ranked importance of the variables (based on each of the nine methodologies) in the neural network based on 500 Monte Carlo simulations. Bars represent 1 standard deviation.

Elimination and Garson's Algorithm performed very poorly (<50% average similarity). The performances of all approaches were found to increase slightly with larger sample sizes, although the magnitude of differences among the methods remained unchanged with the Connection Weight Approach always exhibiting the greatest performance (unpublished results).

Next, we compared the methods' ability to correctly estimate the true ranked importance of each individual variable in the neural network. In agreement with the above results, the Connection Weight Approach showed the best performance for estimating the correct ranked importance (Fig. 2). Our results showed that mean estimated ranked variable importance based on the connection weights exhibits very strong concordance with the actual ranked importance across all variables. Furthermore, this approach exhibits the highest precision compared to the other methods, as indicated by the small variation around its mean. In contrast, Partial Derivatives, Input Perturbation, Sensitivity Analysis and Improved Stepwise Selections 1 and 2 approaches were only successful at identifying the true importance of the two most influential variables (i.e. x_1 and x_2). Estimates for x_3–x_5, on the other hand, were highly variable and resulted in poor resolution to differentiate among their importance. Similarly, Forward Stepwise Addition and Backward Stepwise Elimination were only successful at estimating the ranked importance of the most influential variable (i.e. x_1). The poorest performance was exhibited by Garson's Algorithm, which was found to rarely identify the correct ranked importance of any variable, and exhibited the greatest variation in its estimates.

4. Discussion

Our study provides a robust comparison of the performance of nine different methodologies for assessing variable contributions in artificial neural networks by using simulated data with known correlative properties. Unlike the use of empirical data in comparative studies, the true performances and stabilities of the different methodologies can be accurately assessed using simulated data because the true importance of the variables is known. Consequently, the results from this study exhibit generality and provide researchers with a valid framework for the selection of appropriate

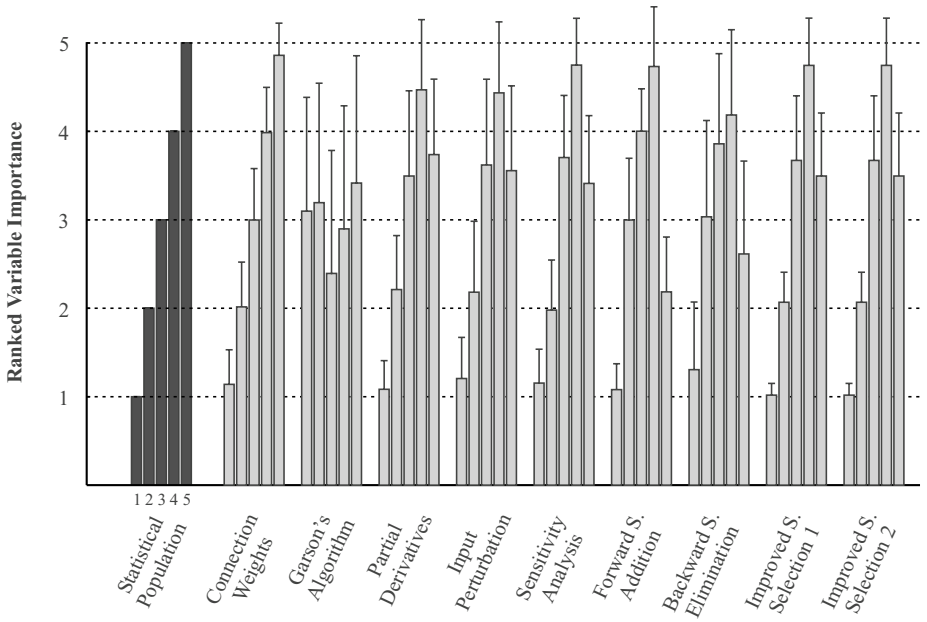

Fig. 2. Mean (bar) and standard deviation (whiskers) of true (Statistical Population) and estimated ranked importance of predictor variables in the neural network (1, most important; 5, least important) according to each of the nine methodologies based on 500 Monte Carlo simulations. Bars in each group represent the ranked importance for variables x_1–x_5 (from left to right); for clarify only labeled for the Statistical Population. Mean values more similar to true rank and those values will small standard deviations indicate methods that consistently estimated the true ranked importance of the particular predictor variable in the neural network.

methodologies in future studies using neural network modeling; a result that is not possible based on the study of Gevrey et al. (2003).

We found that the Connection Weight Approach provides the best overall methodology for accurately quantifying variable importance and should be favored over the other approaches examined in this study. This approach successfully identified the true importance of all the variables in the neural network, including variables that exhibit both strong and weak correlations with the response variable. Moreover, a randomization test for the Connection Weight Approach has been recently developed (see Olden and Jackson, 2002b), which provides a tool for pruning null-connection weights and neurons from the network and determining the statistical significance of variable contributions (see Olden (2000); Olden and Jackson (2001) for examples of its application). This is a distinct advantage of the Connection Weight Approach over the other methods examined in this study.

Plate 1. Illustration of the critical difference between the Connection Weight Approach (Olden and Jackson, 2002b) and Garson's Algorithm (Garson, 1991) that results in their differential ability to correctly identify variable importance in neural networks. The input-hidden and hidden-output connection weights are those of a neural network from the simulation study. The inability of Garson's Algorithm to correctly estimate true variable importance can be simply illustrated for input variable 4, which was incorrectly ranked the most important variable in the network. The Connection Weight product matrix shows that although input neuron 4 positively influences the output neuron via hidden neurons B, C, and D, it also negatively influences the output neuron via hidden neurons A and E. Therefore, because Garson's Algorithm uses absolute connection weights in its calculations it fails to account for the contrasting influences of input neuron 4 through different hidden neurons, resulting in an incorrect estimation of variable importance. In contrast, the Connection Weight Approach uses raw connection weights, which accounts for the direction of the input–hidden–output relationship and results in the correct identification of variable contribution. The presented formulas represent the calculation of variable importance for predictor variable X (where $X = 1$–5) using the weights connecting each of the input neurons Z (where $Z = 1$–5) to each of the hidden neurons Y (where $Y =$ A–E) to the single output neuron.

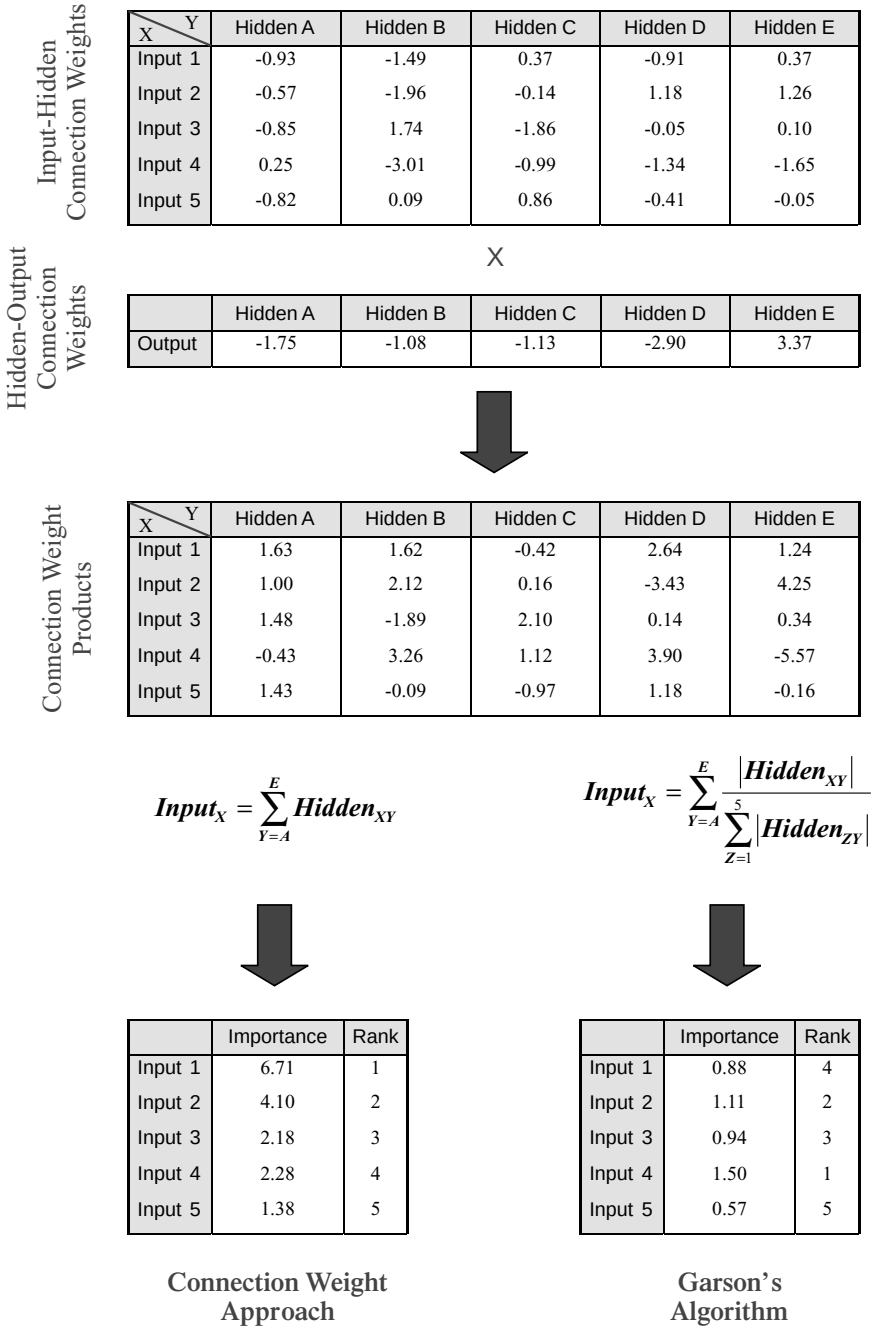
Input-Hidden Connection Weights
X
Y
Hidden A
Hidden B
Hidden C
Hidden D
Hidden E
Input 1 -0.93 -1.49 0.37 -0.91 0.37
Input 2 -0.57 -1.96 -0.14 1.18 1.26
Input 3 -0.85 1.74 -1.86 -0.05 0.10
Input 4 0.25 -3.01 -0.99 -1.34 -1.65
Input 5 -0.82 0.09 0.86 -0.41 -0.05
X
Hidden-Output Connection Weights
Hidden A
Hidden B
Hidden C
Hidden D
Hidden E
Output -1.75 -1.08 -1.13 -2.90 3.37
Connection Weight Products
X
Y
Hidden A
Hidden B
Hidden C
Hidden D
Hidden E
Input 1 1.63 1.62 -0.42 2.64 1.24
Input 2 1.00 2.12 0.16 -3.43 4.25
Input 3 1.48 -1.89 2.10 0.14 0.34
Input 4 -0.43 3.26 1.12 3.90 -5.57
Input 5 1.43 -0.09 -0.97 1.18 -0.16
$Input_X = \sum_{Y=A}^{E} Hidden_{XY}$
$Input_X = \sum_{Y=A}^{E} \frac{|Hidden_{XY}|}{\sum_{Z=1}^{5} |Hidden_{ZY}|}$
Importance
Rank
Input 1 6.71 1
Input 2 4.10 2
Input 3 2.18 3
Input 4 2.28 4
Input 5 1.38 5
Importance
Rank
Input 1 0.88 4
Input 2 1.11 2
Input 3 0.94 3
Input 4 1.50 1
Input 5 0.57 5
Connection Weight Approach
Garson's Algorithm

In agreement with Gevrey et al. (2003), we found that the Partial Derivatives and Input Perturbation approaches performed relatively well, although our results show that they were only consistent in identifying the two most important variables in the network. The Improved Stepwise Selection approaches showed similar performances. The remaining approaches showed both poor accuracy and precision, and are not recommended in future studies. The most notably result of our study is that Garson's Algorithm was the poorest performing approach, yet is the most commonly used in the ecological literature. The inadequacy of Garson's Algorithm is not surprising given that it uses absolute connection weights to calculate variable contributions, and, therefore, does not account for counteracting connection weights linking input and output neurons, that is, opposite directions for incoming and outgoing weights from the hidden layer neurons (see Olden and Jackson (2002b) for further discussion). In contrast, all other approaches use raw connection weights to calculate variable importance. At first glance the Connection Weight Approach and Garson's Algorithm (the best and worst performers) might be expected to produce similar results because they both use input-hidden and hidden-output connection weights when calculating variable importance. In fact, this is not the case. Plate 1 provides a simple example illustrating their key difference and shows how Garson's Algorithm can result in incorrect estimates of variable importance. For this reason, we strongly urge researchers not to use Garson's Algorithm.

In conclusion, although the apparent complexity of artificial neural networks was originally believed to limit our ability to gain explanatory insight into the prediction process, recent advancements such as our study and Gevrey et al. (2003) have illustrated that this indeed not the case. By using a combination of quantitative approaches, such as the use of connection weights or Partial Derivatives, with visual approaches, such as neural interpretation diagrams (Özesmi and Özesmi, 1999) or sensitivity plots (Lek et al., 1996), researchers now have the ability to identify individual and interacting contributions of the predictor variables in neural networks. Although the interpretation of neural networks will obviously never be as straightforward as simple regression models, it is now apparent that a select number of robust methodologies are available for illuminating the "black box" that was once thought to typify artificial neural networks.

Acknowledgements

We thank Rasmus Ejrnaes for his helpful comments and Sovan Lek for his constructive feedback on the manuscript. Muriel Gevrey kindly providing the MatLab® code used in Gevrey et al. (2003), which was used to confirm the accuracy (and correct for Partial Derivatives and Sensitivity methods) the programs used in the analysis. We also thank Don Jackson for his comments on a previous version of the manuscript. Funding for J.D. Olden was provided through a graduate scholarship from the Natural Sciences and Engineering Research Council of Canada. Massey University Center for Ecosystem Modeling and Management graciously provided travel funds for J.D. Olden's visit.

References

Berryman, A., Turchin, P., 2001. Identifying the density-dependent structure underlying ecological time-series. Oikos 92, 265–270.

Bishop, C.M., 1995. Neural Networks for Pattern Recognition. Clarendon Press, Oxford.

Dimopoulos, Y., Bourret, P., Lek, S., 1995. Use of some sensitivity criteria for choosing networks with good generalization. Neural Process Lett. 2, 1–4.

Dimopoulos, I., Chronopoulos, J., Chronopoulos-Sereli, A., Lek, S., 1999. Neural network models to study relationships between lead concentration in grasses and permanent urban descriptors in Athens city (Greece). Ecol. Model. 120, 157–165.

Garson, G.D., 1991. Interpreting neural-network connection weights. Artif. Intell. Expert 6, 47–51.

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.

Hornik, K., White, H., 1989. Multilayer feedforward networks are universal approximators. Neural Netw. 2, 359–366.

Jackson, D.A., 1993. Stopping rules in principal components analysis: a comparison of heuristical and statistical approaches. Ecology 74, 2204–2214.

Joy, M.K., Death, R.G., 2004. Neural network modeling of freshwater fish and macro-crustacean assemblages for biological assessment in New Zealand. In: Lek, S., Scardi, M., Jorgensen, S.E. (Eds.), Modeling Community Structure in Freshwater Ecosystems. Springer-Verlag, New York, USA, in press.

Legendre, P., Legendre, L. 1998. Numerical Ecology. Elsevier Scientific, Amsterdam.

Lek, S., Belaud, A., Baran, P., Dimopoulos, I., Delacoste, M., 1996a. Role of some environmental variables in trout abundance models using neural networks. Aquat. Living Resour. 9, 23–29.

Lek, S., Delacoste, M., Baran, P., Dimopoulos, I., Lauga, J., Aulagnier, S., 1996b. 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.

Minchin, P.R., 1987. An evaluation of the relative robustness of techniques for ecological ordination. Vegetato 69, 89–107.

Olden, J.D., 2000. An artificial neural network approach for studying phytoplankton succession. Hydrobiology 436, 131–143.

Olden, J.D., Jackson, D.A., 2000. Torturing data for the sake of generality: how valid are our regression models? Écoscience 7, 501–510.

Olden, J.D., Jackson, D.A., 2001. Fish-habitat relationships in lakes: gaining predictive and explanatory insight by using artificial neural networks. Trans. Am. Fish Soc. 130, 878–897.

Olden, J.D., Jackson, D.A., 2002a. A comparison of statistical approaches for modeling fish species distributions. Fresh Biol. 47, 1976–1995.

Olden, J.D., Jackson, D.A., 2002b. Illuminating the "black box": understanding variable contributions in artificial neural networks. Ecol. Model. 154, 135–150.

Özesmi, S.L., Özesmi, U., 1999. An artificial neural network approach to spatial habitat modeling with interspecific interaction. Ecol. Model. 116, 15–31.

Paruelo, J.M., Tomasel, F., 1997. Prediction of functional characteristics of ecosystems: a comparison of artificial neural networks and regression models. Ecol. Model. 98, 173–186.

Peres-Neto, P.R., Olden, J.D., 2001. Assessing the robustness of randomization tests: examples from behavioral studies. Anim. Behav. 61, 79–86.

Recknagel, F. (Ed.), 2003. Ecological Informatics: Understanding Ecology by Biologically-Inspired Computation. Springer-Verlag, New York, USA.

Rumelhart, D.E., Hinton, G.E., Williams, R.J., 1986. Learning representations by backpropagation errors. Nature 323, 533–536.

Scardi, M., Harding, L.W., 1999. Developing an empirical model of phytoplankton primary production: a neural network case study. Ecol. Model. 120, 213–223.

Vander Zanden, M.J., Olden, J.D., Thorne, J.H., Mandrak, N.E., 2004. Predicting occurrences and impacts of smallmouth bass introductions in north temperate lakes. Ecol. Appl. 14, 132–148.