

# Neural Networks

## *Theoretical Foundations and Analysis*

---

Edited by

**Clifford Lau**

Office of Naval Research  
Arlington, Virginia



A Selected Reprint Volume  
IEEE Neural Networks Council, *Sponsor*

The Institute of Electrical and Electronics Engineers, Inc. New York

# A Performance Comparison of Trained Multilayer Perceptrons and Trained Classification Trees

LES ATLAS, MEMBER, IEEE, RONALD COLE, YESHWANT MUTHUSAMY, STUDENT MEMBER, IEEE, ALAN LIPPMAN, JEROME CONNOR, DONG PARK, MOHAMED EL-SHARKAWI, SENIOR MEMBER, IEEE, AND ROBERT J. MARKS II, SENIOR MEMBER, IEEE

*Multilayer Perceptrons and trained classification trees are two very different techniques which have recently become popular. Given enough data and time, both methods are capable of performing arbitrary nonlinear classification. We first consider the important differences between multilayer Perceptrons and classification trees and conclude that there is not enough theoretical basis for the clear-cut superiority of one technique over the other. For this reason, we performed a number of empirical tests on three real-world problems in power system load forecasting, power system security prediction, and speaker-independent vowel recognition. In all cases, even for piecewise-linear trees, the multilayer Perceptron performed as well as or better than the trained classification trees.*

## I. INTRODUCTION

We use and compare two types of regression and classification systems. A regression system generates an output  $Y$  for an input  $X$ , where both  $X$  and  $Y$  are continuous and perhaps multidimensional. A classification system generates an output class  $C$  for an input  $X$ , where  $X$  is continuous and multidimensional and  $C$  is a member of a finite alphabet.

The use of trained classification and regression systems has been studied by many researchers in the past (see, for example, [1]–[4]). However, there has been a recent surge of interest in trainable systems such as artificial neural net-

works (ANNs). In particular it has been shown that the multilayer Perceptron (MLP) can be trained by example to solve the nonlinearly separable exclusive-OR problem [5], and this architecture has been linked to previous neural-like processors [6], [7]. Less known to the engineering community is the statistical technique of classification and regression trees (CART) which was developed during the years 1973 through 1984 [8], [9].

CART, like the MLP, can be trained to solve the exclusive-OR problem, the solution it provides is extremely easy to interpret, and both CART and MLPs are able to approximate arbitrary nonlinear decision boundaries. Although there have been no links made between CART and biological neural networks, the possible applications and paradigms used for MLP and CART are very similar.

The authors of this paper represent diverse interests in problems which have the commonality of being important and potentially well suited for trainable classifiers. The *load forecasting* problem, which is partially a regression problem, uses past load trends to predict the critical needs of future power generation. The *power security* problem uses the classifier as an interpolator of previously known states of the system. The *vowel recognition* problem is representative of the difficulties in automatic speech recognition caused by variability across speakers and phonetic context.

In each problem area, large amounts of real data were used for training and disjoint data sets were used for testing. We were careful to ensure that the experimental conditions were identical for the MLP and CART. We concentrated only on performance as measured in error on the test set and did no formal studies of training or testing time. (CART was, in general, quite a bit faster in training and testing.)

In all cases, even with various sizes of training sets, the multilayer Perceptron performed as well as or better than the trained classification trees. We also believe that integration of many of CART's well-designed attributes into MLP architectures could only improve the already promising performance of MLPs.

Manuscript received September 1, 1989; revised March 16, 1990. This work was supported by a National Science Foundation Presidential Young Investigator Award for L. Atlas and by separate grants from the National Science Foundation, Washington Technology Center, U.S. West Advanced Technologies, and Puget Sound Power and Light Co.

L. Atlas, J. Connor, D. Park, M. El-Sharkawi and R. J. Marks II are with the Dept. of Electrical Engineering, University of Washington, Seattle, WA 98195, USA.

R. Cole and Y. Muthusamy are with the Dept. of Computer Science and Engineering, Oregon Graduate Center, Beaverton, Oregon 97006-1999 USA.

A. Lippman is with the School of Oceanography, University of Washington, Seattle, WA 98195, USA.

IEEE Log Number 9037385.

II. BACKGROUND

A. Multilayer Perceptrons

The name “artificial neural networks” has in some communities become almost synonymous with MLPs trained by backpropagation. Our power studies made use of this standard algorithm [5] and our vowel studies made use of a conjugate gradient version [10] of backpropagation. In all cases the training data consisted of ordered pairs  $\{(X, Y)\}$  for regression, or  $\{(X, C)\}$  for classification. The input to the network is  $X$  and the output is, after training, hopefully very close to  $Y$  or  $C$ .

When MLPs are used for regression, the output  $Y$  can take on real values between 0 and 1. This normalized scale was used as the prediction value in the power forecasting problem. For MLP classifiers the output is formed by taking the (0, 1) range of the output neurons and either thresholding or finding a peak. For example, in the vowel study we chose the maximum of the 12 output neurons to indicate the vowel class.

B. Classification and Regression Trees (CART)

CART has already proven to be useful in diverse applications such as radar signal classification, medical diagnosis, and mass spectra classification. Given a set of training examples  $\{(X, C)\}$ , a binary tree is constructed by sequentially partitioning the  $p$ -dimensional input space, which may consist of quantitative and/or qualitative data, into  $p$ -dimensional polygons. The trained classification tree divides the domain of the data into nonoverlapping regions, each of which is assigned a class label  $C$ . For regression, the estimated function is piecewise constant over these regions.

The first split of the data space is made to obtain the best global separation of the classes. The next step in CART is to consider the partitioned training examples as two completely unrelated sets—those examples on the left of the selected hyperplane, and those on the right. CART then proceeds as in the first step, treating each subject of the training examples independently. A question that had long plagued the use of such sequential schemes was: when should the splitting stop? CART implements a novel, and very clever approach; splits continue until every training example is separated from every other, then a pruning criterion is used to sequentially remove less important splits.

The CART system was trained using two separate computer routines. One was the CART program from California Statistical Software; the other was a routine we designed ourselves. We produced our own routine to ensure a careful and independent test of the CART concepts described in [9].

C. Relative Expectations of MLP and CART

The nonlinearly separable exclusive-or problem is an example of one that both MLP and CART can solve with zero error. In Fig. 1(a)—a trained MLP solution to this problem—the values along the arrows represent trained multiplicative weights and the values. In Fig. 1(b)—the very simple trained CART solution— $y$  and  $n$  represent yes or no answers to the trained threshold and the values in the circles represent the output  $Y$ . It is interesting that CART did not train correctly for equal numbers of the four different input cases and that

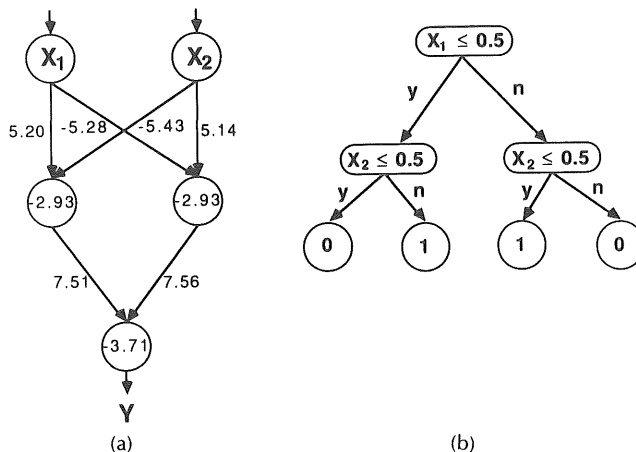


Fig. 1. (a) A multilayer perceptron (MLP) and (b) a classification tree. Both were trained to perfectly solve the exclusive-OR problem.

one extra example of one of the input cases was sufficient to break the symmetry and allow CART to train correctly. (Note the similarity to the well-known requirement of randomness and different initial weights for the MLP.)

CART trains on the exclusive-or very easily since a piecewise-linear partition in the input space is a perfect solution. In general, the MLP will construct classification regions with smooth boundaries, whereas CART will construct regions with “sharp” corners (each region being, as described previously, an intersection of half planes.) We would thus expect MLP to have an advantage when classification boundaries tend to be smooth and CART to have an advantage when they are sharper.

Other important differences between MLP and CART include:

1) For an MLP the number of hidden units can be selected to avoid overfitting or underfitting the data. CART fits the complexity by using an automatic pruning technique to adjust the size of the tree. The selection of the number of hidden units or the tree size was implemented in our experiments by using data from a second training set (independent of the first).

2) An MLP becomes a classifier through an ad hoc application of thresholds or peak-picking to the output value(s). Great care has gone into the CART splitting rules while the usual MLP approach is rather arbitrary.

3) A trained MLP represents an approximate solution to an optimization problem. The solution may depend on initial choice of weights and on the optimization technique used. For complex MLPs many of the units are independently and simultaneously adjusting their weights to best minimize output error.

4) MLP is a distributed topology where a single point in the input space can have an effect across all units or analogously, one weight, acting alone, will have minimal effect on the outputs. CART is very different in that each split value can be mapped onto one segment in the input space. The behavior of CART makes it much more useful for data interpretation. A trained tree may be useful for understanding the structure of the data. The usefulness of MLPs for data interpretation is much less clear.

The above points, when taken in combination, do not make a clear case for either MLP or CART to be superior for

the best performance as a trained classifier. We thus believe that the empirical studies of the next sections, with their consistent performance trends, will indicate which of the comparative aspects are the most significant.

III. LOAD FORECASTING

A. The Problem

The ability to predict electric power system loads from an hour to several days in the future can help a utility operator to efficiently schedule and utilize power generation. This ability to forecast loads can also provide information that can be used to strategically trade energy with other generating systems. In order for these forecasts to be useful to an operator, they must be accurate and computationally efficient.

B. Methods

Hourly temperature and load data for the Seattle/Tacoma area were provided for us by the Puget Sound Power and Light Company. Forecasting for weekdays is a more critical problem for the power industry than for weekends and we selected the hourly data for all Tuesdays through Fridays in the interval of November 1, 1988 through January 31, 1989. These data consisted of 1368 hourly measurements from a total of 57 days.

These data were presented to both the MLP and the CART systems as a 6-dimensional input with a single, real-valued output. The MLP required that all values be normalized to the range (0, 1). These same normalized values were used with the CART technique. Our training and testing process consisted of training the classifiers on 523 days of the data and testing on the 4 days left over at the end of January 1989. Our training set consisted of 1272 hourly measurements and our test set contained 96 different hourly readings.

Several techniques of input and output pairing were tried; after some investigation we found that a good choice of data organization for our trainable classifier was

$$(X, Y) = (k, L_{k-2}, L_{k-1}, T_{k-2}, T_{k-1}, T_k, L_k)$$

where  $k$  was the hour (1-24) of the day and  $L_i$  and  $T_j$  signified the load and temperature at the  $i$ th and  $j$ th hour, respectively. The input thus consists of the hour, two previous load and temperature readings, and the current temperature. The actual current temperature was used during training and the predicted temperature was used during testing, thus representing the actual technique of relying upon weather reports. The output part is the predicted load  $L_k$ .

The MLP we used in these experiments had 6 inputs (plus the trained constant bias term), 10 units in one hidden layer, and one output. This topology was chosen by making use of data outside the training and test sets.

C. Results

We used an  $l_1$  norm for the calculation of error rates and found that both techniques worked quite well. The average error rate was 1.39% for the MLP and 2.86% for CART. Although this difference (given the number of testing points) is not statistically significant, it is worth noting that the trained MLP offers performance at least as good as the current techniques used by the Puget Sound Power and

Light Company and is currently being verified for application to future load prediction.

Figure 2 shows a detail of the comparative forecasting performance for three days. The daily periodicity in hourly loads was followed quite well by both techniques, and the MLP performed somewhat better than CART around the peaks in load.

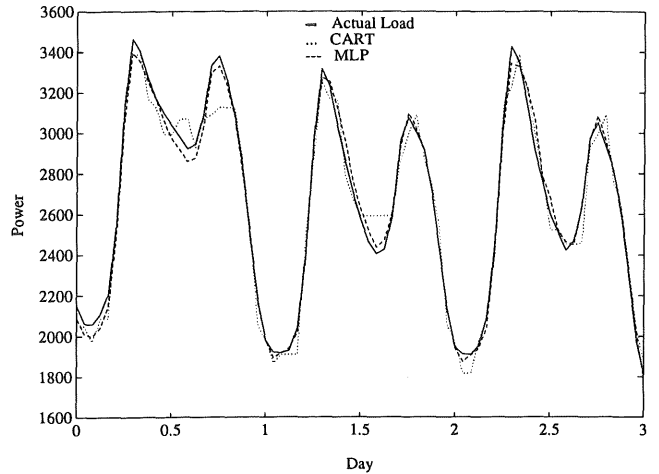


Fig. 2. A comparison of the performance of an MLP and CART (with linear combinations) in predicting three days of hourly power loads in megawatts.

IV. POWER SYSTEM SECURITY

The assessment of security in a power system is an ongoing problem for the efficient and reliable generation of electric power. Static security addresses whether, after a disturbance such as a line break or other rapid load change, the system will reach a steady-state operating condition that does not violate any operating constraint and cause a "brownout" or "blackout."

The most efficient generation of power is achieved when the power system is operating near its insecurity boundary. In fact, the ideal case for efficiency would be the full knowledge of the absolute boundaries of the secure regions. The complexity of the power systems makes this full knowledge impossible. Load flow algorithms, which are based on iterative solutions of nonlinearly constrained equations, are conventionally used to slowly and accurately determine points of security or insecurity. In real systems the trajectories through the regions are not predictable in fine detail. Also, these changes can happen too fast to compute new results from the accurate load-flow equations.

We thus propose to use the sparsely known solutions of the load flow equations as a training set. The test set consists of points of unknown security. The error of the test set can then be computed by comparing the result of the trained classifier to load flow equation solutions.

Our technique for converting this problem to a problem for a trainable classifier involves defining a training set  $\{(X, C)\}$  where  $X$  is composed of real power, reactive power, and apparent power at another bus. This three-dimensional input vector is paired with the corresponding security status ( $C = 1$  for secure and  $C = 0$  for insecure). Since the system was small, we were able to generate a large number of data points for training and testing. In fact, well over 20 000

superior performance of the MLP technique, all of which we are currently investigating. One advantage may stem from the ability of MLP to easily find correlations between large numbers of variables. Although it is possible for CART to form arbitrary nonlinear decision boundaries, the efficiency of the recursive splitting process may be inferior to MLP's nonlinear fit. Another relative disadvantage of CART may be caused by the successive nature of node growth. For example, if the first split that is made for a problem turns out to be suboptimal given the successive splits, it becomes very inefficient to change the first split to be more suitable.

We feel that the careful statistics used in CART could also be advantageously applied to MLP. The superior performance of MLP is not yet indicative of best performance and it may turn out that careful application of statistics may allow further advancements in the MLP technique. Other input representations also might produce better performance for CART than for MLP.

New developments have been made in trained statistical classifiers since the development of CART. More recent techniques, such as projection pursuit [14], may prove as good as or superior to MLP. This continued interplay between MLP techniques and advanced statistics is a key part of our ongoing research.

#### ACKNOWLEDGMENT

The authors wish to thank Professor R. D. Martin of the University of Washington Department of Statistics and Professors Aggoune, Damborg, and Hwang of the University of Washington Department of Electrical Engineering for their helpful discussions. David Cohn and Carlos Rivera assisted with many of the experiments.

We would also like to thank Milan Casey Brace of Puget Power and Light for providing the load forecasting data.

#### REFERENCES

- [1] N. J. Nilsson, *Learning Machines*. New York: McGraw-Hill, 1965.
- [2] A. G. Arkadev and E. M. Braverman, *Computers and Pattern Recognition*. Washington, D.C.: Thompson, 1966.
- [3] W. S. Meisel, *Computer-Oriented Approaches to Pattern Recognition*. New York: Academic Press, 1972.
- [4] R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*. New York: Wiley, 1973.
- [5] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning internal representations by error propagation," ch. 2 in *Parallel Distributed Processing*, D. E. Rumelhart, J. L. McClelland, and the PDP Research Group. Cambridge, MA: M.I.T. Press, 1986.
- [6] W. S. McCulloch and W. H. Pitts, "A logical calculus of the ideas imminent in nervous activity," *Bulletin of Mathematical Biophysics*, vol. 5, pp. 115-133, 1943.
- [7] E. Rosenblatt, *Principles of Neurodynamics*. New York: Spartan Books, 1962.
- [8] W. S. Meisel and D. A. Michalopoulos, "A partitioning algorithm with application in pattern classification and the optimization of decision trees," *IEEE Trans. Computers*, vol. C-22, pp. 93-103, 1973.
- [9] L. Breiman, J. H. Friedman, R. A. Olshen and C. J. Stone, *Classification and Regression Trees*. Belmont, CA: Wadsworth International, 1984.
- [10] P. E. Barnard and D. Casasent, "Image processing for image understanding with neural nets," *Proc. Intl. Joint Conf. on Neural Nets*. Washington, D. C., June 18-22, 1989.
- [11] W. Fisher, G. Doddington, and K. Goudie-Marshall, "The DARPA Speech Recognition Research Database: Specification and Status," *Proc. of the Darpa Speech Recognition Workshop*, pp. 993-1000, Feb. 1986.
- [12] L. Lamel, R. Kassel, and S. Seneff, "Speech database development: design and analysis of the acoustic-phonetic corpus," *Proc. of the DARPA Speech Recognition Workshop*, pp. 100-110, Feb. 1986.
- [13] Y. Muthusamy, R. Cole, and M. Slaney, "Vowel information in a single spectral slice: Cochleagrams versus spectral slice: Cochleagrams versus spectrograms," *Proc. ICASSP '90*, Apr. 3-6, 1990 (to appear).
- [14] J. H. Friedman and W. Stuetzle, "Projection pursuit regression," *J. Amer. Stat. Assoc.*, vol. 79, pp. 599-608, 1984.

total data points were available for the (disjoint) training and test sets.

### A. Results

We observed that for any choice of training data set size, the error rate for the MLP was always lower than the rate for the CART classifier. This performance difference is illustrated in Fig. 3. For 10 000 points of training data, the MLP

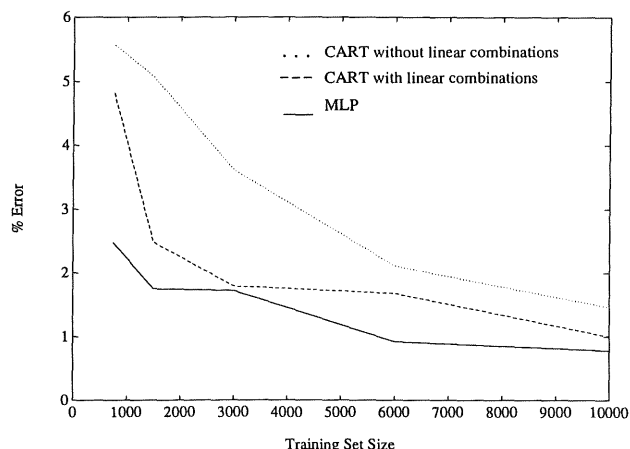


Fig. 3. The error rate in security prediction for the MLP and two versions of CART.

had an error rate of 0.78% and CART (using linear combinations) had an error rate of 1.46%. Although both of these results are impressive, the difference was statistically significant ( $p > 0.99$ ).

In order to gain insight into the reasons for differences in importance, we looked at classifier decisions for two-dimensional slices of the input space. While the CART boundary sometimes was a better match, certain pathological difficulties made CART more error-prone than the MLP. Our other studies also showed that there were worse interpolation characteristics for CART, especially for sparse data. Apparently, starting with nonlinear combinations of inputs, which is what the MLP does, reduced error better than the piecewise linear fit of CART.

## V. SPEAKER-INDEPENDENT VOWEL CLASSIFICATION

Speaker-independent classification of vowels excised from continuous speech is a most difficult task because of the many sources of variability that influence the physical realization of a given vowel. These sources of variability include the length of the speaker's vocal tract, phonetic context in which the vowel occurs, speech rate, and syllable stress.

To make the task even more difficult, the classifiers were presented only with information from a single spectral slice. The spectral slice, represented by 64 DFT coefficients (0–4 kHz), was taken from the center of the vowel, where the effects of coarticulation with surrounding phonemes are least apparent.

The training and test sets for the experiments consisted of featural descriptions  $X$  paired with an associated class  $C$  for each vowel sample. The 12 monophthongal vowels of English were used for the classes, as heard in the following

words: beat, bit, bet, bat, roses, the, but, boot, book, bought, cot, bird. The vowels were excised from a wide variety of phonetic contexts in utterances of the TIMIT database, a standard acoustic phonetic corpus of continuous speech, displaying a wide range of American dialectical variation [11], [12]. The training set consisted of 4104 vowels from 320 speakers. The test set consisted of 1644 vowels (137 occurrences of each vowel) from a different set of 100 speakers.

The MLP consisted of 64 inputs (the DFT coefficients, each normalized between zero and one), a single hidden layer of 40 units, and 12 output units (one for each vowel category). The networks were trained using backpropagation with conjugate gradient optimization [10]. The procedure for training and testing a network proceeded as follows: The network was trained on 100 iterations through the 4104 training vectors. The trained network was then evaluated on the training set and a different set of 1644 test vectors (the test set). The network was then trained for an additional 100 iterations and again evaluated on the training and test sets. This process was continued until the network had converged; convergence was observed as a consistent decrease or leveling off of the classification percentage on the test data over successive sets of iterations.

### A. Results

In order to better interpret the vowel classification results, we performed listening experiments on a subset of the vowels used in these experiments. The vowels were excised from their sentence context and presented in isolation. Five listeners first received training in the task by classifying 900 vowel tokens and receiving feedback about the correct answer on each trial. During testing, each listener classified 600 vowels from the test set (50 from each category) without feedback. The average classification performance on the test set was 51% correct, compared to chance performance of 8.3%. Details of this experiment are presented in [13]. When using the scaled spectral coefficients to train both techniques, the MLP correctly classified 47.4% of the test set while CART without linear combinations performed at only 38.2%.

One reason for the poor performance of CART without linear combinations may be that each coefficient (corresponding to energy in a narrow frequency band) contains little information when considered independently of the other coefficients. For example, reduced energy in the 1-kHz band may be difficult to detect if the energy in the 1.06-kHz band is increased by an appropriate amount. The CART classifier described in the preceding operates by making a series of inquiries about one frequency band at a time, an intuitively inappropriate approach.

We achieved our best CART results, 46.4%, on the test set by making use of arbitrary hyperplanes (linear combinations). This search-based approach gave results within 1% of the MLP results.

## VI. CONCLUSIONS

In all cases the performance of the MLP was, in terms of percent error, better than CART. However, the difference in performance between the two classifiers was only significant (at the  $p > 0.99$  level) for the power security problem.

There are several possible reasons for the sometimes