A Note on a Closed-Form Training Algorithm for the Neocognitron

David Lovell, Ah Chung Tsoi & Tom Downs
Intelligent Machines Laboratory, Department of Electrical Engineering
University of Queensland, Queensland 4072, Australia

Abstract

In this note, a difficulty with the application of Hildebrandt's closed-form training algorithm for the neocognitron is reported. In applying this algorithm we have observed that S-cells frequently fail to respond to features that they have been trained to extract. We present results which indicate that this training vector rejection in an important factor in the overall classification performance of the neocognitron trained using Hildebrandt's procedure.

1 Introduction

Recently, Hildebrandt [1] presented a closed-form training algorithm for the feature extracting S-cells used in Fukushima's neocognitron (NC). Part of the function of this algorithm is to tune the *selectivity parameter* (r) of the S-cells in a quantitative manner. Hildebrandt argued that his algorithm allows determination of an optimal value of r because it ensures that the acceptance regions are as large as they can be without overlapping. This made the closed-form algorithm very interesting since the only other methods for choosing r, that we are aware of, are based upon heuristics.

It is the purpose of this letter to describe our experience in applying Hildebrandt's algorithm which, we have found, leads to performance that is inferior to that obtained from heuristic techniques. We offer some suggestions as to why this might occur.

It should be noted that the following discussion applies only to Hildebrandt's closed-form training algorithm for the neocognitron. It does not apply to the general thresholded linear correlation classifier model, and its closed-form training algorithm, which were also presented in [1].

2 Comparative Test Results

We measured the classification performance of a NC that had been trained using the methods described by Fukushima [2, 3] and compared it to an identically structured NC trained using Hildebrandt's closed-form training algorithm [1]. In the following, these networks are referred to as NC-F and NC-H, respectively.

In order to comply with Hildebrandt's definition of an optimal thresholded linear correlator, the selectivity parameters of the S-cells in the NC-H network were chosen to maximize generalization whilst maintaining a discrimination index of unity[1].

The network structure and the training patterns that we used are described in [2]. The values of the selectivity parameter used in the NC-F network are given in [3], and the Gaussian kernels (used to weight the S-cell and C-cell inputs in both the NC-F and NC-H networks) were calculated using the methods described in that paper.

[1] we use the terms *generalization* and *discrimination* as defined in [1].

Network	% Correctly Classified	% Misclassified	% Unclassified
NC-F	52.50%	26.25%	21.25%
NC-H	3.25%	10.00%	86.75%

Table 1: Classification performance of a neocognitron trained using Fukushima's method (NC-F) and Hildebrandt's closed-form training algorithm (NC-H).

A set of 400 distorted digits[2] was used to estimate the performance of the two networks to within 5% accuracy with 95% confidence [4]. The results shown in Table 1 have been partitioned into three categories: trials in which the input was correctly classified, trials in which the input was misclassified and trials in which the input was left unclassified (*i.e.* there was no input activity in any of the final stage units).

3 Discussion

Table 1 indicates that a network trained using the closed-form training algorithm is unable to match the classification performance of Fukushima's neocognitron. But the implication of Hildebrandt's paper is that an S-cell in a network trained with the closed-form algorithm should be able to generalize (*ie.* respond to feature vectors similar to those it was trained with) to the greatest extent possible *without* responding to features that are extracted by other S-cells in the same layer.

One contributing factor to the results in Table 1 is that, in using the neocognitron model, we have imposed a vector representation upon features which are not necessarily "Euclidean" in nature [5]. A geometrical approach (as taken by Hildebrandt) to partitioning some high dimensional feature space may not accord with the nature of the features present in that space.

A more tangible difficulty with the closed-form training algorithm is that it does not *guarantee* that an S-cell will respond to a feature that it was trained to extract. We have frequently observed this phenomenon in networks trained with the closed-form algorithm and we use a geometric argument to explain its occurrence.

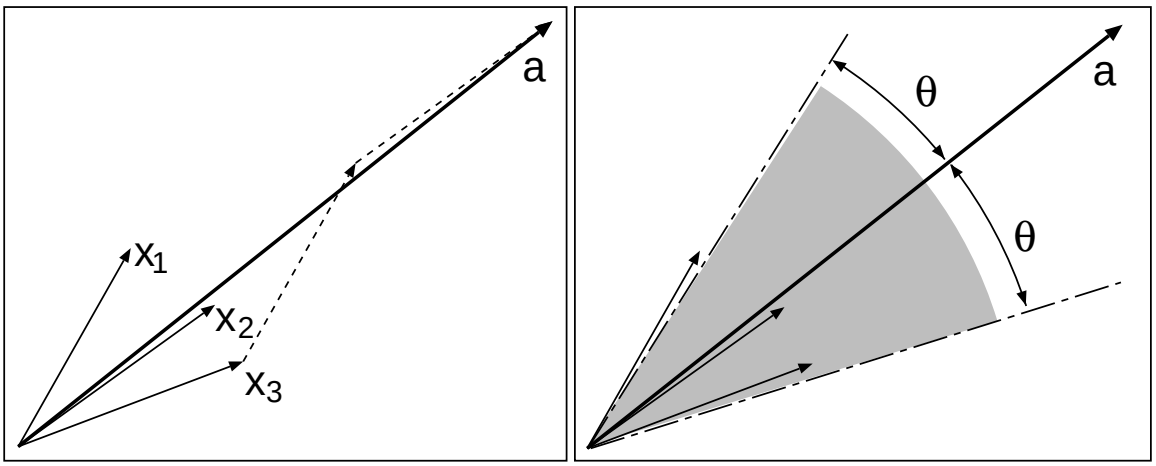

Figure 1: (a) A template is the sum of the training vectors that the S-cell is to respond to (b) failure to respond to a training feature.

The closed-form training algorithm has two stages. The first stage establishes the direction of the feature vector that the S-cell is to respond to. This template vector is the sum of all

[2]This test set of digits is available via email request to the first author at `dlovell@s1.elec.uq.oz.au`

features that the cell is assigned to extract. Figure 3(a) shows a template vector a generated in response to features x_1, x_2 and x_3.

The second stage of the closed-form training algorithm determines the *selectivity* of the S-cell being trained. The selectivity defines an angle θ in N-dimensional feature space. If the input vector lies within the angle θ of the template, then it will elicit a response from the S-cell.

The closed-form training algorithm makes it possible for the situation of Figure 3(b) to arise, where one of the features (x_1) used to determine the template a does not lie within the region of feature space that the S-cell will respond to (the shaded sector). Thus, an S-cell can easily fail to respond to one of its own training patterns. This inability to extract training patterns contributes substantially to the degradation in performance seen in Table 1.

In [1] Hildebrandt pointed out that the training methods employed by Fukushima are susceptible to the same problem of training vector rejection. But in our application of Fukushima's method, which involves hand-tuning of selectivity parameters, this phenomenon has not been observed. Since the selectivity parameters constitute the only difference between the NC-H and NC-F networks, we must conclude that the difference in performance evident in Table 1 arises from the fact that training vector rejection is much more prevalent when Hildebrandt's method is employed.

4 Conclusions

In this note, we have described a difficulty which arises in applying Hildebrandt's closed-form training algorithm to the neocognitron. We have observed that a neocognitron trained using Hildebrandt's algorithm does not work as well as a neocognitron trained by Fukushima's methods. We have explained how Hildebrandt's algorithm can cause feature extractors to become unresponsive to their training patterns, a phenomenon that is a major contributor to the degradation in classification performance that we have observed.

Acknowledgements

The first author acknowledges the support of an Australian Postgraduate Research Award. The second and third authors acknowledge support from the Australian Research Council.

References

[1] Hildebrandt T. H. (1991). Optimal Training of Thresholded Linear Correlation Classifiers. *IEEE Transactions on Neural Networks*, **2**(6), 577-588.

[2] Fukushima, K. (1988). Neocognitron: A Hierarchical Neural Network Capable of Visual Pattern Recognition. *Neural Networks*, **1**, 119-130.

[3] Fukushima, K., & Wake, N. (1991). Handwritten Alphanumeric Character Recognition by the Neocognitron. *IEEE Transactions on Neural Networks*, **2**(3), 355-365.

[4] Hogg, R. V., & Ledolter, J. (1987). *Engineering Statistics*, 163-166. Macmillan, New York.

[5] Goldfarb, L. (1985). A New Approach to Pattern Recognition. *Progress in Pattern Recognition 2*. Kanal, L. N., & Rosenfeld, A. (eds.). North-Holland.