In *The Neurobiology of Computation*, pp. 433-438, Eeckman and Bower (Eds), Kluwer, 1994

# COMPARISON OF GENDER RECOGNITION BY PDP AND RADIAL BASIS FUNCTION NETWORKS

**Shih-Cheng Yen, Paul Sajda, Leif Finkel**

*Department of Bioengineering and*

*Institute for Neurological Sciences*

*University of Pennsylvania, Philadelphia, PA 19104-6392*

## INTRODUCTION

Despite a long history of neurological, psychological, and computational efforts, no satisfactory explanation has been offered for the extraordinary ability of humans to recognize other human faces. However, a number of different network-based approaches (Turk and Pentland,1991; Brunelli and Poggio, 1993; Buhmann et al., 1989) have achieved surprisingly good ability to recognize faces, at least under certain restricted conditions. We decided to compare the solutions developed by different network architectures including PDP and radial basis function (RBF) networks to the problem of gender classification. Given a picture of a face, including external features such as hair, beard, jewelry, etc., the network must learn to distinguish male from female. This is a simpler problem than general face recognition, and there is some evidence that it is carried out by a separate population of cells in the inferior temporal cortex (Damasio et. al., 1990).

Several investigators have previously applied PDP networks to the problem of gender classification (Golomb et al., 1989; Cottrell and Metcalfe, 1989). However, the hidden unit representations developed in those models were not analyzed in detail. Moreover, we wanted to directly compare the representations developed by different types of networks (PDP, RBF) when confronted with the exact same training and test sets.

## Methods

Using pictures from several college yearbooks, we created a database of 1400 faces, all cropped to the same size (50x70 pixels) (see Figure 1). A subset of these faces were picked to form the training sets and the remaining faces were then used as the test set. The training sets typically included between 20 and 60 faces and the test set contained 1200 novel faces.

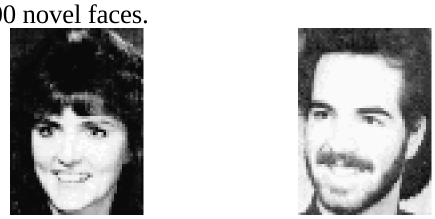

Figure 1. Examples of training faces

The networks used in our simulations were standard 3-layered networks (50x70 input, 10 hidden, 1 output). The networks were trained on a desired output of 1.0 for male faces and 0.0 for female faces. All simulations were conducted using the NEXUS neural network simulation environment (Sajda and Finkel, 1992).

## RESULTS

### Generalization

Networks trained using either PDP or RBF performed remarkably well. For most of our simulations, we consistently achieved close to 90% correct classification on the test set. This performance was in spite of the fact that our database consisted of faces posed at different orientations and photographed under varied lighting conditions.

One observation we made was that the networks had trouble learning "unusual" faces. Part of the learning curve for a PDP network trained on 21 faces is shown below. The network had trouble with a female face with very short hair and a male face with long hair (see Figure 2), a common problem even for humans.

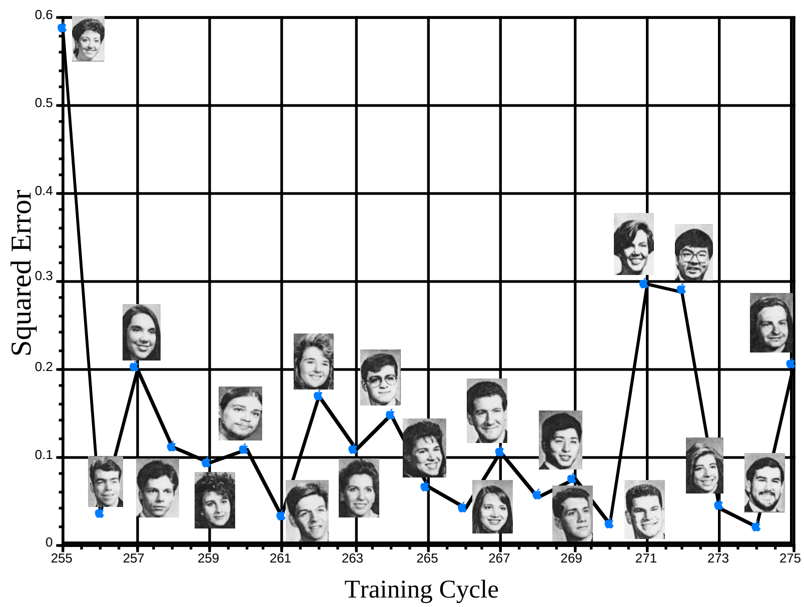


Figure 2. Learning Curve for PDP network trained on 21 faces

## Noise

We also investigated the robustness of the network's performance in the presence of noise in the input images. A subset (200) of the test set was modified by adding varying degrees of uniform white noise. Both the PDP and the RBF networks showed remarkable tolerance to noise, even for signal-to-noise ratios of down to 1.0 (see Figure 3). Since the network maintains its performance in spite of significant changes in pixel values, this suggests that the network is using distributed or configural representations that are not easily degraded in the presence of noise.

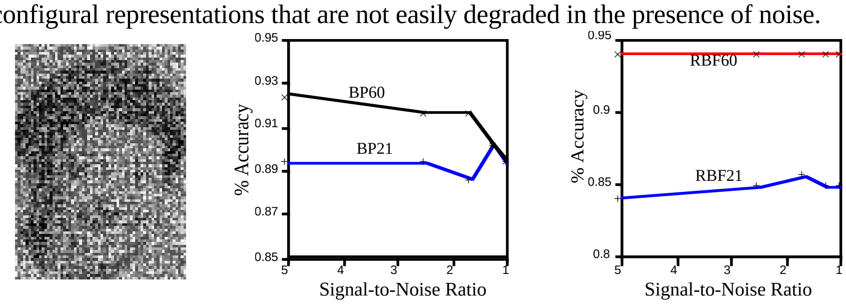


Figure 3. Image with signal-to-noise ratio of 1.0 (left); performance of PDP network (middle) and RBF network (right) on noisy images. BP21, BP60 are the PDP networks trained with 21 and 60 images. RBF 21 and RBF60 are the RBF counterparts.

### Image Orientation

Humans have a remarkable ability to generalize between novel views of previously seen objects. This ability however, does not extend to upside-down faces. The networks also exhibit these properties. 200 faces were again picked from the test set and flipped horizontally or vertically. These images were then presented to the networks. The graphs below (see Figure 4) show that both the PDP and RBF networks, perform just as well for horizontally flipped faces but are significantly impaired at classifying upside-down faces.

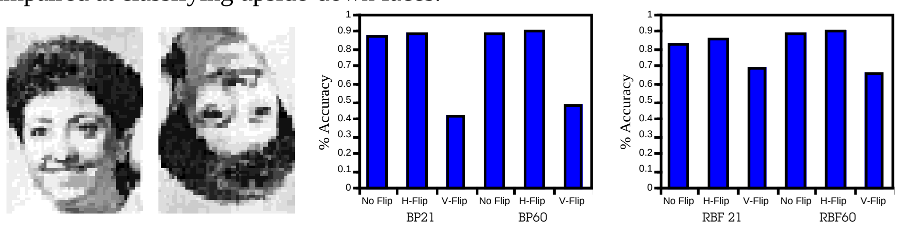


Figure 4. Horizontally flipped face, vertically flipped face, the PDP and the RBF network performances

## ANALYZING NETWORK BEHAVIOR

### Network Connectivity

In both the PDP and the RBF networks, we observed a small number of hidden units whose connection maps look like averaged faces, as well as a number of units that seem to have connection maps that are relatively accurate replicas of particular faces in the training set (see Figure 5). When we looked at the variance of the hidden unit's activities across test images, we found that the variance of the first type of unit is much larger than that of the second type. This suggests that the network may be using a small number of hidden units to broadly classify the faces and the rest of the hidden units to code for the outliers.

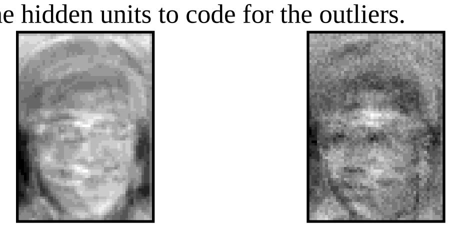

Figure 5. Connection map of cells representing average faces (left) and outliers (right).

**Network Inversion**

To further investigate the representation developed by the networks, we used network inversion to obtain the optimal stimulus for each desired output. The networks seemed to have identified the region directly below the ears and the region of the forehead as the best guides to the classification problem. This is a surprisingly reliable cue as most males have short hair and have much more exposed forehead than females. The network seems to have learned that facial hair occurs predominantly in males. This is intriguing since psychological studies have shown the importance of “external features” (hair style, face outline) in human face recognition (Shepherd et al, 1981).

**Status of the Template Model**

In order to gain insight into how difficult a problem gender classification is, and to determine whether performance levels in the 90th percentile should be judged significant, we developed a simple alternative network approach. A “male” template was made by linearly averaging all the male faces in the database (680). A “female” template was similarly constructed from 443 female faces. These templates are shown in figure 7.

Figure 7. Averaged male template (left) and averaged female template (right)

A simple network architecture then measured the difference between each input pattern and the two templates -- whichever match was closer was the winner. The network was tested on the whole database and its performance turned out to be very similar to the trained networks. The network also performed comparably to the PDP and RBF trained networks in terms of noise tolerance as well as horizontally and vertically flipped faces. This suggests that crude template matching, where the templates are constructed as linear combinations of inputs, may account for most of the performance of the PDP and RBF networks. Additional accuracy can then be achieved by adding specialized detectors for the outlying faces in each gender.

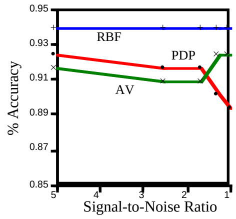


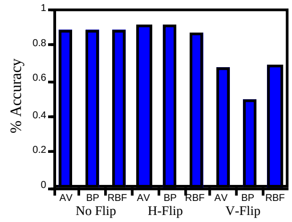


Figure 8. Comparison of various networks on noisy images (left) and image orientation (right).

## CONCLUSION

Our results suggest that PDP and RBF networks both adopt a classification strategy that combines a statistical decision scheme (hidden units representing combinations of faces) with an exemplar or prototype scheme (units committed to representing individual outliers). Human conceptual categorization has similarly been shown to depend on a combination of statistical and prototype schemes (Smith and Medin, 1981). This may suggest that both training algorithms have captured the essential structure of the gender-recognition problem.

## REFERENCES


Brunelli R., Poggio T. Face recognition--features versus templates. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 15:1042-1052, 1993.

Buhmann, J., Lades, M., von der Malsburg, C. Distortion invariant object recognition by matching hierachically labeled graphs. *IJCNN International Joint Conference on Neural Networks 1989.* 1:155-159, 1989.

Damasio, A. R., Tranel, D., Damasio, H. Face agnosia and the neural substrates of memory. *Annual Review of Neuroscience.* 13:89-109, 1990.

Cotrell, G. W., Metcalfe, J. EMPATH: Face, emotion and gender reocgnition using holons. In Touretsky & Lippmann, editors, *Advances in Neural Information Processing Systems 3*, 572-577. Morgan Kauffmann, 1991.

Golomb, B. A., Lawrence, D. T., Sejnowski, T. J. SEXnet: A neural network indentifies sex from human faces. In Touretsky & Lippmann, editors, *Advances in Neural Information Processing Systems 3*, chapter 572-577. Morgan Kauffmann, 1991.

Sajda, P. and Finkel, L. H. NEXUS: A simulation environment for large-scale neural systems. *Simulation*, 59:358-364, 1992.

Shepherd, J. W., Davies, G. M., Ellis, H. D. Studies of cue saliency. In Davis, Ellis, Shepherd (Eds), *Perceiving and remembering faces*, London:Academic Press, 1981.

Smith, E. E., Medin, D. L. *Categories and Concepts*. Cambridge: Harvard University Press, 1981.

Turk, M. and Pentland, A. Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3:71-86, 1991.