Pergamon

Neural Networks, Vol. 8, No. 3, pp. 481–488, 1995

0893-6080/95 $9.50 + .00

0893-6080(94)00077-8

CONTRIBUTED ARTICLE

Studies on Object Recognition From Degraded Images Using Neural Networks

A. RAVICHANDRAN AND B. YEGNANARAYANA

Indian Institute of Technology

(*Received* 23 *July* 1993; *revised and accepted* 28 *June* 1994)

Abstract—*The objective of this paper is to study the performance of artificial neural network models for recognition of objects from poorly resolved, noisy, and transformed (scaled, rotated, translated) images, such as images reconstructed from sparse and noisy data in a sensor array imaging context. Noise and sparsity of data in the imaging context result in degradation of quality of the reconstructed image as a whole, instead of affecting it in the form of local corruption of the image pixel information as in many image processing situations. Hence, (i) neighbourhood processing methods for noise cleaning may not be suitable, (ii) feature extraction cannot be reliably performed, and (iii) model-based methods for classification cannot easily be applied. In this paper, we show that neural network models can be used to overcome some of the difficulties in dealing with degraded images as obtained in an imaging context.*

Keywords—Object recognition, Degraded images, Transformation invariance, Sensor array imaging, Sparsity, Image reconstruction, Preprocessing, Neural networks.

1. INTRODUCTION

In this paper we address the problem of recognition of objects from degraded images obtained through reconstruction from sparse and noisy data, as in the case of sensor array imaging. The aim of a sensor array imaging (Yegnanarayana, Mariadassou, & Saini, 1990) system, such as an underwater acoustic imaging system, is to obtain an image of an object by transmitting acoustic waves and sensing the wave reflected from the object using an array of sensors as shown in Figure 1. We consider a simulated sensor array imaging context to generate array data for imaging the objects.

Due to sparsity of the received data and noise, the reconstructed images are poorly resolved and noisy. It is difficult to recognise an object from these reconstructed images due to lack of visual clues required for recognition. The task of automatic recognition becomes much more complex when the object moves relative to the receiver array. In this case recognition is required independent of changes in scale, position, and orientation of the object in the image. As the number of expected targets becomes larger, it becomes difficult for a human observer to recognise the object from such poor quality images.

Requests for reprints should be sent to Dr. B. Yegnanarayana, Department of Computer Science & Engineering, Indian Institute of Technology, Madras-600 036, India. E-mail: yegna @ iitm.ernet.in.

The difficulty in building a recognition system to deal with degraded images arises from the fact that noise and sparsity of data in the imaging context result in degradation of quality of the reconstructed image as a whole instead of affecting it in the form of local corruption of the image pixel information. Hence, (i) neighbourhood processing methods for noise cleaning may not be applicable, (ii) feature extraction cannot be reliably performed, and (iii) model-based methods for classification cannot easily be applied.

To deal with the degraded images, it is desirable that the processing mechanism adapts itself to the task domain and to the data being processed. Representation and description suitable to the task domain must be acquired by *learning* from examples and not through explicit specification. Data must be processed by activating the knowledge in a *context-sensitive* manner and not in a preprogrammed manner. In this paper, we describe how simple artificial neural networks, which exhibit capabilities for adaptation, even if only to a limited extent, can be put to use in a complex pattern recognition task.

The task addressed in this paper is recognition of a fixed set of planar objects. Studies reported here are made by simulating noise and sparsity as obtained in a simplified model of a sensor array imaging situation. We limit the study to simple metric transformations (rigid object transformations in two dimensions), namely, scaling, translation, and in-plane rotation. The

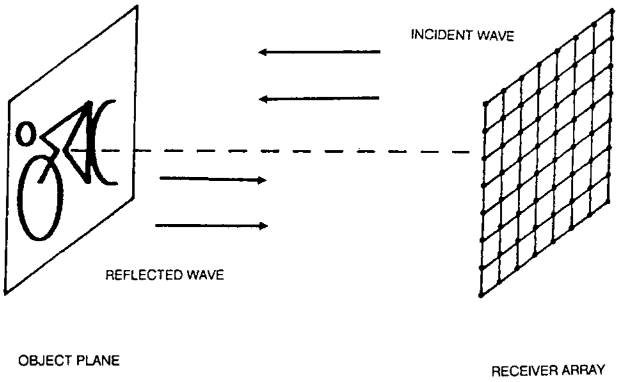

FIGURE 1. A simplified sensor array imaging setup.

type and details of the objects in the images are important factors that may decide the complexity of the recognition task. In this study we consider two sets of objects of varying complexity: (a) complex set—Olympic game symbols, and (b) simpler set—characters of the English alphabet.

We consider a simulated sensor array imaging setup to generate the sensor array data. By varying the parameters of the source, receiver, and medium characteristics, it is possible to generate data that will produce images over a wide range of degradations in a controlled manner.

To demonstrate the effectiveness of neural network models for different cases, the following studies are made. Object recognition studies from images with degradations due to noise and sparsity alone are discussed in Section 2. Studies for the case of degradations due to transformations alone are discussed in Section 3. Finally, studies for the case of degradations due to noise, sparsity, and transformations are discussed in Section 4.

2. DEGRADATION DUE TO NOISE AND SPARSITY ALONE

In this section we consider the situation where the images of objects are poorly resolved and corrupted by noise. In these situations feature extraction is not likely to be successful. Because the goal is just classification and not generation of a description of the image, a pixelwise template-matching approach may be preferred over feature-based methods. We have used a Hamming network to learn the templates of the images and to match them by correlation (Lippmann, 1987).

The objective is to test the performance of the Hamming network under different levels of degradation and to identify the limits up to which the network can perform reliably. Sparsity of samples and noise are used to generate images of different levels of degradation. The knowledge base in the present studies consists of 20 Olympic games symbols, five of which are shown in Figure 2a. These symbols were created by first digitising the printed symbols from a newspaper using a scanner and editing the symbols. Each image consists of 128 × 128 pixels. Three different recognition experiments were performed using this data.

Experiment 1: Effect of Sparsity

Images were reconstructed from sparse data collected using receiver arrays of different sizes. Data at the sensor array were simulated assuming array sizes of 64 × 64, 32 × 32, 16 × 16, and 8 × 8 sensors and using data at two frequencies in each case. Images were reconstructed using an iterative reconstruction algorithm based on the Projection onto Convex Sets (POCS) with the constraints of finite support on the object plane and the known data on the receiver plane (Yegnanarayana et al., 1990). Figure 2b shows the images reconstructed from an 8 × 8 array data. Each of these images was converted into a binary image and presented to the neural network that was already trained to the original 20 images. The templates of the images are stored using the training algorithm for the Hamming network described in Lippmann (1987). The recognised symbol for each reconstructed image is given in the figure.

The results are summarised in Figure 3a. All the 20 patterns are correctly classified when the array size is

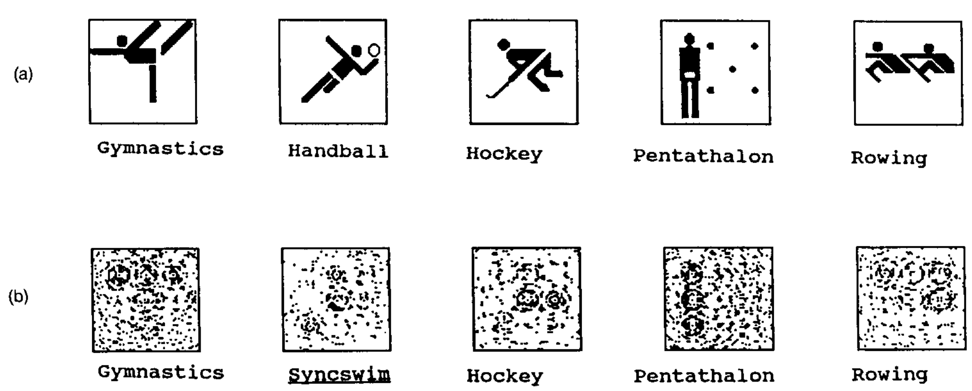

FIGURE 2. Some Olympic game symbols and the corresponding images reconstructed from sparse data. (a) A subset of Olympic game symbols in the knowledge base (each image has 128 × 128 pixel points). (b) Images reconstructed from data collected at two frequencies from an 8 × 8 array. The decision of the network is given. The underlined one is misclassified.

greater than or equal to 16 × 16. This performance is very impressive because it is difficult for us to identify visually discriminating features in many of these images.

Experiment 2: Performance With Data at Multiple Frequencies

Images were reconstructed from data collected at multiple frequencies of the incident plane waves in the simulated sensor array imaging setup (see Figure 1). Sensor array data were simulated assuming an 8 × 8 array with data collected at different frequencies. The images reconstructed from one, two, four, and eight different frequencies data were presented to the network for recognition. Iterative algorithms based on POCS were used to reconstruct the images as before (Yegnanarayana et al., 1990). The recognition results are summarised in Figure 3b. Classification performance improves even for smaller array sizes like 8 × 8 array, when the data are collected at more frequencies.

Experiment 3: Effect of Noise

Images were reconstructed from data with various levels of noise. Random noise with Gaussian distribution was added to the simulated sensor array data collected at two frequencies with a 16 × 16 array. The recognition results obtained are summarised in Figure 3c. It appears that degradation effects due to noise are more severe than the degradation caused due to sparsity of data. However, for the case of 16 × 16 data, the network classifies satisfactorily up to a noise level of 0 dB.

In the above study it was assumed that transformation and spatial distortion of objects do not occur, and the network used direct pixelwise matching. In practice, these conditions are seldom met because objects being imaged move relative to the imaging system. Therefore, it is difficult to recognise objects in such images using the Hamming network.

3. DEGRADATION DUE TO TRANSFORMATIONS ALONE

If images are clean and noise free, there exist methods for overcoming the effects of metric transformations. In this section we explore this possibility of transformation invariant recognition of objects from noise-free images. We have used a transformation-invariant feature space based on the theory of geometric moments, and a neural network classifier that uses this feature space for object recognition.

3.1. Issues in Transformation-Invariant Recognition

Neural approaches for transformation-invariant recognition were actively pursued in the 1980s (Fukushima, 1983; Giles, Griffin, & Maxwell, 1988; Widrow & Winter, 1988). In these approaches, attempts were made to obtain invariance either by suitably designing the structure or by training the networks. Neither of these approaches is realistic from an engineering point of view (Barnard & Casasent, 1991). Since 1990, *hybrid* systems that use invariant feature spaces to handle transformations, and neural networks for classification have been studied (Barnard & Casasent, 1991; Khotanzad & Lu, 1990). There are two major advantages in this approach:

(i) The requirements on the classifier are reduced because the number of features required is much less than the number of pixels.
(ii) Invariance for all input objects is ensured.

The following are the disadvantages of using invariant feature spaces:

(i) The input image is not directly input to the classifier. Preprocessing is required to compute the features, which may be computationally expensive.
(ii) Not all feature spaces are equally suitable for a given problem. Each feature space has its own shortcomings. Moment-invariant feature spaces have difficulties when noise is present. Feature spaces, such as wedge ring samples of the magnitude of the Fourier transform and the magnitude of the Fourier transform in log-polar coordinates, are not invariant to all possible transformations (Barnard & Casasent, 1991).

The invariant feature approach is a global feature-matching approach useful only in simplified situations where we can assume that the object is segmented (object is separated from its uniform background) and that there is no occlusion. This enables us to define invariant features computed directly from the image. If the object and the background are not distinct, local feature extraction and normalisation have to be done one after the other to achieve transformation invariance.

3.2. A Two-Stage Recognition Approach for Transformation-Invariant Recognition

We describe a two-stage approach for recognition. In the first stage, we use invariant feature spaces to handle transformation and in the second stage we use a neural network for classification.

Methods based on the theory of geometric moments have been used for normalisation and invariant feature extraction (Hu, 1962). If the object is compact and has only a few details, invariant measures stable over a wide range of spatial transformations can be designed. If the transformations are metric, then it is possible to design such an invariant feature space. Based on theories of invariant algebra that deal with properties of algebraic expressions that remain invariant under general linear transformations, Hu (1962) derived com-

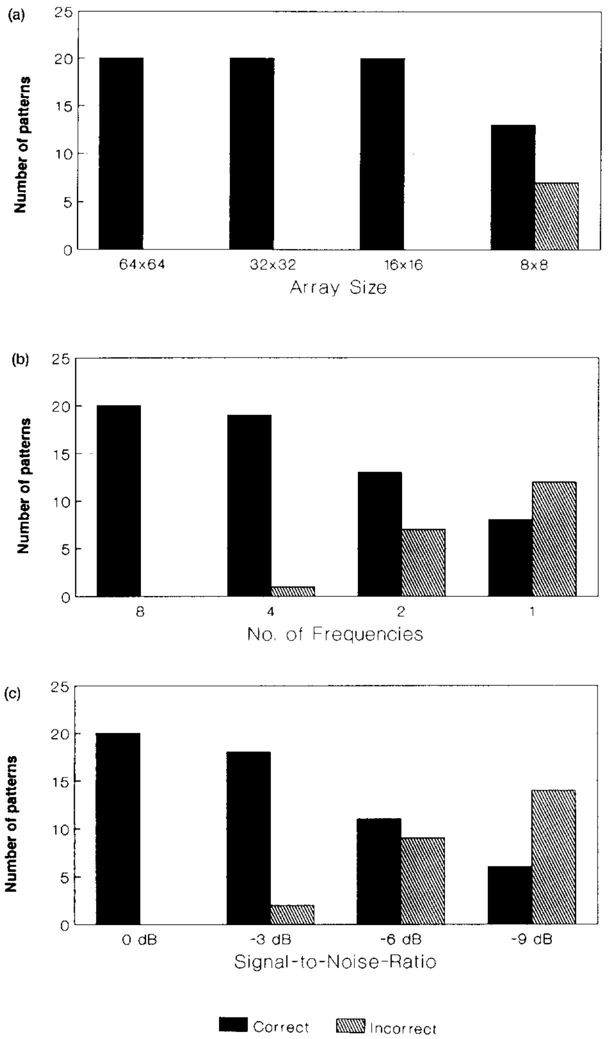

FIGURE 3. Number of patterns recognised for different cases: (a) with different sparse arrays (64 × 64, 32 × 32, 16 × 16, and 8 × 8 sensors); (b) when data at different frequencies are used with an 8 × 8 sensor array; (c) for different levels of noise in the received data with a 16 × 16 array.

binations of moment values that are invariant with respect to scale, position, and orientation. The effectiveness of this space for classification purposes has been studied over the years (Prokop & Reeves, 1992). The present implementation closely follows Khotanzad and Lu (1990).

The utility of the moment invariants is illustrated through the following experiment. Figure 4 shows several Olympic game symbols represented in a two-dimensional feature space formed by the first two moment invariants. It must be noted that some of the symbols that are very different in image shape are close to each other in the feature space. Hence, moment-invariant features may not correspond to visually discriminating features employed by the human visual system.

To perform classification based on these features, we have used a multilayered feedforward neural network, trained using the error back-propagation algorithm (Rumelhart, Hinton, & Williams, 1986). In our study, we have used a network with six nodes in the input layer corresponding to the six moment-invariant features and 20 nodes in the output layer corresponding to the 20 objects. The network has one hidden layer. For classification, the network was trained using noise-free normalised patterns.

3.3. Experimental Studies

Eight different images from each of the 20 Olympic games symbols are generated, consisting of varying scales, orientations, and translations of each image. Some of these images (six for each symbol) are shown in Figure 5.

Two images per symbol were used for training and the remaining six for testing. Classification accuracy of 100% was obtained up to a scale reduction that causes a 3:1 reduction in the length of the image (i.e., for an original image size of 128 × 128 points, 3:1 reduction results in an image of approximately 40 × 40 points). Beyond this, further reduction causes misclassification in many cases due to loss of object detail in the reduced image. Hence, for the set of Olympic games symbols this appears to be the level up to which the approach can be reliably used.

To study the effect of details in the image on the recognition performance, we have tested the approach for classification of 10 characters of the English alphabet. These character images contain far fewer details than the olympic symbols. Therefore, they could be scaled down to a much lower value without severe distortion. Eight different images from each of these characters were generated by scaling, rotation, and translation. Two of these were used for training and the remaining six for testing. Classification accuracy of 100% was obtained for all the test data up to a scale reduction that causes a 12:1 reduction in the length of the image (i.e., a reduced image of size 10 × 10 points). Thus, the moment feature approach works better in the case of objects with simple shapes.

The results point out that if normalisation is done separately, then the tasks of feature extraction and classification are simplified. However, the approach has the following limitations:

(i) This technique is not easily generalised to provide invariance against other nonlinear transformations of the pattern. For example, this approach may not be suitable for hand-printed character recognition, where distortions are not necessarily linear transformations.

(ii) In situations where the scene consists of multiple objects, the system must be capable of paying attention to the individual objects in a scene, for each of which the invariance of perception must separately be valid.

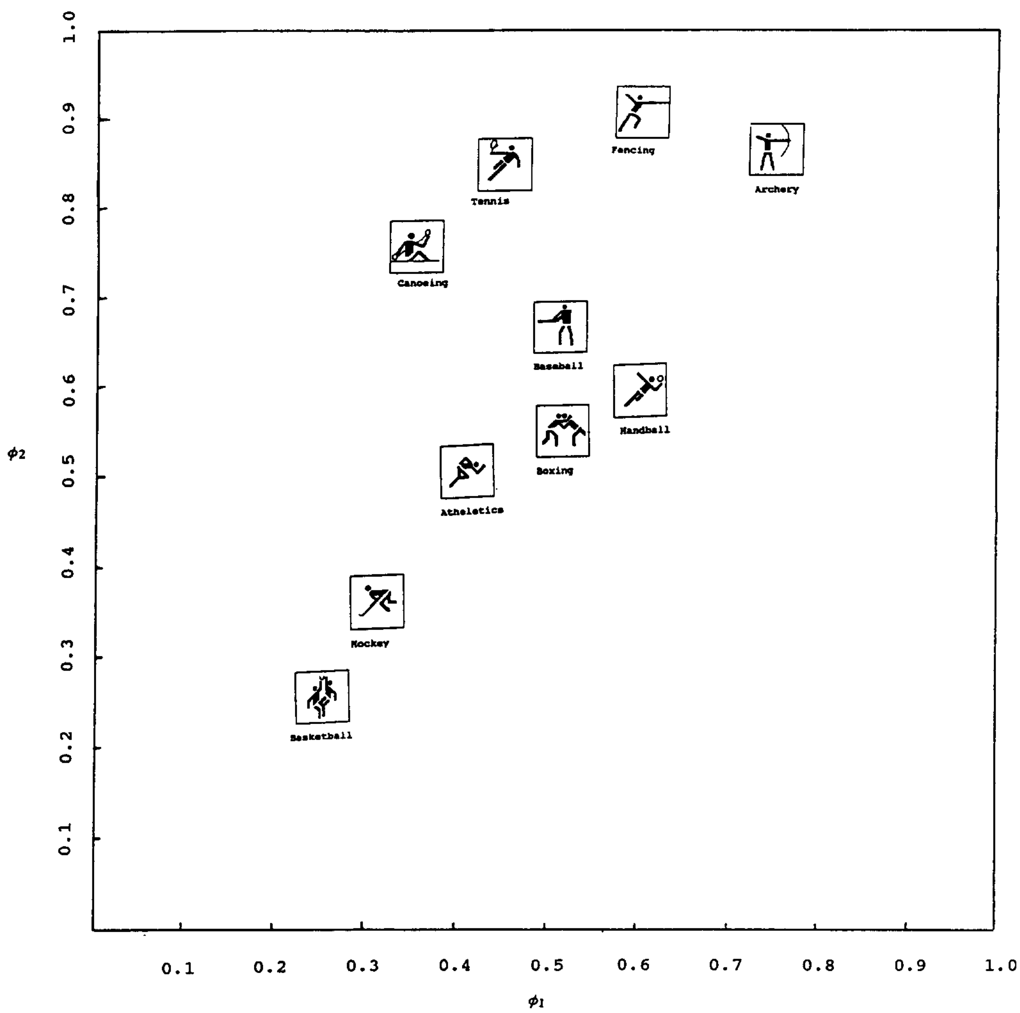

FIGURE 4. Some Olympic game symbols represented in a two-dimensional feature space formed by the first two moment-invariant features, $\phi 1$ and $\phi 2$. Values represent logarithm of the absolute values of the features, normalised to unity.

Though the approach has these limitations, it seems effective in the present context, where we have assumed no more than one object in the scene at any time. The object assumed was planar and rigid, and the problem of occlusion was not addressed.

4. DEGRADATION DUE TO NOISE, SPARSITY, AND TRANSFORMATIONS

4.1. Need for Preprocessing and Noise Suppression

Noisy and transformed images are difficult to recognise using the two-stage approach. This is because during the computation of moments, we do not distinguish between pixels of the object and noise. If the moment features are extracted directly from the noisy image, the estimates are not accurate. There is thus a need for a preprocessing stage before features can be extracted.

Suppressing noise and segmenting an image into object in the foreground and noise in the background is a nontrivial task. In general, this requires physical and semantic knowledge about generic class of objects and even the specific object (Grossberg, Mingolla, & Todorovic, 1989). However, if the image can be modelled as a single two-dimensional object placed on a uniform background, general purpose models may be useful. These include models for general classes of local features such as blobs, edges, etc., as well as models that describe how such features can be grouped into aggregates. Traditional segmentation algorithms are two-stage sequential processes: local features are detected in the first stage, and they are grouped in the second stage. In the presence of noise and data sparsity, such a strictly sequential process, where labels are first determined and processed later, may not work. This is because when data is noisy, labels are ambiguous. For example, each pixel in a binary image may be interpreted either as an image pixel or as a noise pixel. Hence, there is need for an interactive process where locally ambiguous interpretations compete to achieve a globally less ambiguous interpretation. A neural architecture is ideally suited for such a task.

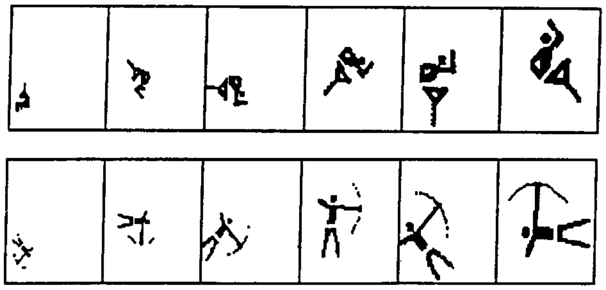

FIGURE 5. Examples of rotated, scaled, and translated images of Olympic game symbols used to study transformation-invariant recognition of objects.

4.2. Design of the Noise-Suppression Network

We propose a multiscale processing approach. In the case of a compact rigid object with homogenous surface, measured surface properties remain the same over a range of scales and neighbourhoods, whereas noise is specific. In other words, surface features that are stable when the scale is varied can be considered as *features* of the object and those that disappear abruptly can be labelled as noise. This generalisation holds good for compact objects.

We propose an analog neural network in which measurements and hypotheses from differently sized neighbourhoods are integrated through cooperative and competitive interactions to perform noise suppression and object extraction. The proposed network consists of three stages. In the first stage, surface patches are detected at three different scales by measuring local image contrast. A surface patch represents a region of the object. Due to the presence of noise and the fact that the detector windows overlap, the detector outputs are ambiguous. Hence, in the second stage, by competitive interaction between adjacent detectors of the same scale, the ambiguity is reduced. In the third stage, the surface patches at different scales interact to obtain a combined output.

The present network design is inspired by the CORT-X filter described in Carpenter, Grossberg, and Mehanian (1989) for extracting sharp boundaries from noisy binary images. The major differences between the CORT-X filter and the present preprocessor arise from the fact that the function of the CORT-X filter is to extract the boundary and complete it, giving rise to a coherent meaningful boundary, whereas the function of the proposed network is to suppress noise and segment the object. Hence, although the CORT-X filter uses oriented contrast detectors to detect object boundaries, the proposed network uses unoriented contrast detectors to detect surface patches on the object.

Stage 1: Unoriented Contrast Detection

Unoriented contrast detectors exist corresponding to every point in the image. Each contrast detector has a receptive field with a smaller (inner) square at the center and a larger (outer) area surrounding it. It is sensitive to the amount and spatial scale of image contrast at a given image location. Thus, each detector hypothesises a surface patch of its scale present in the image.

Let $I(x,y)$ denote the intensity of the input image at position (x,y) in the lattice. The total excitatory input to the detector $E_s(x,y)$ is obtained by integrating the total activation in the inner square of its receptive field:

$$E_s(x,y) = \frac{\iint_{\text{inner}} I(x,y)\,dx\,dy}{\iint_{\text{inner}} dx\,dy}.$$

where s is the index of the size of the receptive field.

Similarly, the total inhibitory input $F_s(x,y)$ is obtained by integrating the total activation in the outer surrounding area of its receptive field:

$$F_s(x,y) = \frac{\iint_{\text{outer}} I(x,y)\,dx\,dy}{\iint_{\text{outer}} dx\,dy}.$$

The output of the contrast detector is defined as

$$C_s(x,y) = \max[E_s(x,y) - \alpha_s F_s(x,y) - \beta_s, 0],$$

where α_s (>1) is a contrast parameter, β_s ($0 < \beta_s < 1$) is a threshold parameter, and the max operator ensures that the output signal is nonnegative.

Large-scale filters suppress noisy pixel distributions effectively, but localise the image patches poorly due to their broader spatial sampling. Small-scale filters, on the other hand, are sensitive to noise but are more reliable in localisation. Hence each filter in itself is insufficient to suppress noise and to extract localised image patches. We need multiple scale interactions. These are described in Stage 3.

Stage 2: Spatial Competition Within Each Scale

The aim of this competitive stage, realised by an on-center off-surround network, is to reduce ambiguity in hypothesis among spatial neighbours and select more probable hypotheses. Each detector output excites the cell activity at the next layer, which represents the same position and scale, while inhibiting cell activities at the neighbouring locations. As a result, cells that have high activity suppress activities of nearby cells that have lower activity due to noise. Under equilibrium condition, the cells that initially had higher activity saturate in a winner-take-all fashion, and those with low initial activity are cut off.

The output $D_s(x,y)$ of the cell at position (x,y) and scale s is given by

$$D_s(x,y) = \frac{C_s(x,y)}{1 + \gamma_s \Sigma_x \Sigma_y C_s(x,y) G_s(x,y)}$$

where $C_s(x,y)$ is the activity of the detector at the input, $G_s(x,y)$ is the competition kernel realised in the

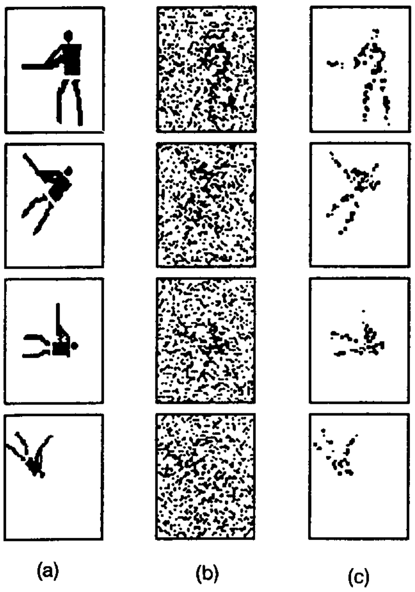

FIGURE 6. (a) Transformed images of the Olympic game symbol for baseball. (b) Corresponding images obtained by reconstruction from data collected by a sparse 16 × 16 sensor array at two different frequencies. (c) Images in (b) after noise suppression.

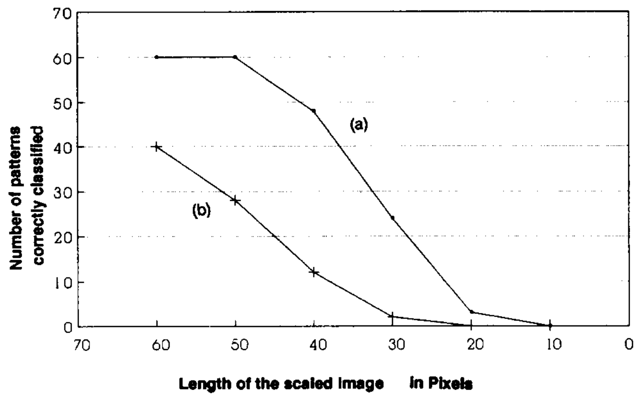

FIGURE 7. Transformation-invariant recognition of Olympic game symbols from images obtained by reconstruction from data collected by (a) 32 × 32 array and (b) 16 × 16 array at two frequencies. Graph shows the number (out of a set of 60 test patterns) of objects (maximum size 128 × 128) correctly classified as the size of the image is reduced.

present implementation by a Laplacian kernel, and γ_s is a parameter for controlling the effect of competition.

Stage 3: Multiple Scale Interaction

The responses of various detectors are combined to retain evidence that is available at multiple scales and remove those unsupported across scales. This is done through cooperative interactions:

$$M(x,y) = D_1(x,y)[\sum_{xy} D_2(x,y)U(x,y)][\sum_{xy} D_3(x,y)U(x,y)]$$

where $D_1(x,y)$, $D_2(x,y)$, and $D_3(x,y)$ are the responses of detectors at three different scales in the present implementation, $U(x,y)$ is an unoriented excitatory kernel (Carpenter, Grossberg, & Mehanian, 1989), and $M(x,y)$ is the combined output. The purpose of the kernel is to make the effect of the larger-scale detectors more diffuse spatially owing to their broader receptive fields. The noise-suppressed output $M(x,y)$ is analog, ranging from 0 to 1, and is thresholded to yield a binary image.

4.3. Experimental Studies in Preprocessing and Classification

A subset of Olympic games symbols consisting of 10 symbols was chosen for this study. Noisy, incomplete, and transformed images of these were obtained by reconstruction from simulated sparse data. We have considered 32 × 32 array and 16 × 16 array data for preprocessing and classification studies. Besides Olympic games symbols, images of alphabet characters were also used to test the classification performance for the case of objects with simpler shapes.

Figure 6 illustrates the effect of the preprocessing stage. Figure 6a shows the transformed images of the Olympic games symbol for baseball, used as objects in imaging simulation. Corresponding images reconstructed from data collected by a 16 × 16 array are shown in Figure 6b and the preprocesser outputs are shown in Figure 6c. Even to the human observer, the preprocessed images in Figure 6c are clearer than those in Figure 6b. This is because many unnecessary details have been removed by preprocessing, reducing the strain on the observer who can now concentrate his attention on the discriminating features. It is still difficult for us to recognise the objects from the preprocessed images due to poor resolution and missing fea-

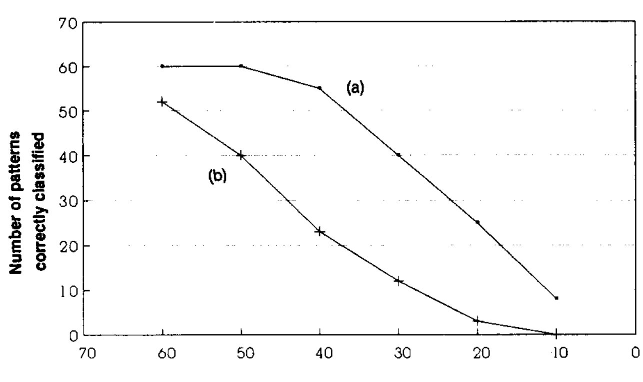

FIGURE 8. Transformation-invariant recognition of alphabet characters from images obtained by reconstruction from data collected by (a) 32 × 32 array and (b) 16 × 16 array at two frequencies. Graph shows the number (out of a set of 60 test patterns) of objects (maximum size 128 × 128) correctly classified as the size of the image is reduced.

tures. However, given the preprocessed image along with the list of the original symbols, the task of the human observer is simplified. Even though it may not be possible to rely on the system for stand-alone recognition performance, it should still be possible to use the output as an aid to human decision making.

Transformation-invariant features were extracted from the preprocessed images and are given as input to the neural network classifier for recognition. Figure 7 shows the results for 32×32 and 16×16 array data cases. We observe that 100% recognition is obtained for 32×32 array data case (Figure 7a) up to a scale reduction that causes 2:1 reduction in the length of the image (i.e., when the reduced image is of size 50×50 pixels). This is to be compared with the recognition performance in the case of noise-free images where 100% accuracy was obtained up to a scale reduction that causes 3:1 reduction in the length of the image. Hence, with increased image degradation, the amount of distortion that can be tolerated is reduced.

For 16×16 array data case (Figure 7b), we observe that for a scale reduction that causes a 2:1 reduction in the length of the image, the classification accuracy is only about 65%. Although this classification performance is quite impressive because the corresponding images are extremely poor in quality, this percentage is too low to be reliable in a practical system. Moreover, this falls very rapidly to nearly zero when the scale reduction in length is as low as 4:1.

To test the recognition performance with a set of objects with simpler shapes, experiments were conducted on the alphabet data. Results are shown in Figure 8. For 32×32 array case (Figure 8a), 100% classification accuracy was obtained up to a scale reduction that causes a 3:1 reduction in length of the image (when size of the reduced image is 40×40 points), which is better than for the Olympic games symbols.

For 16×16 array case (Figure 8b) near 100% classification accuracy was obtained only up to a scale reduction that causes a 2:1 reduction in length of the image, which is better than for Olympic games symbols case using 16×16 array data, but worse than the alphabet case using 32×32 array data.

5. CONCLUSION

In this paper, we have studied the performance of artificial neural network models that can be trained to identify objects from poorly resolved, noisy, and transformed images, such as images reconstructed from sparse and noisy data.

We have made a systematic study of the object recognition problem by identifying aspects of complexity, namely, types and levels of image degradation, transformational variability, nature of symbol set, etc., and performing the study in steps of increasing complexity. This enables us to identify the limits up to which the systems and approaches can be useful.

The systematic study has revealed a desirable characteristic of neural network models, namely, graceful degradation of performance with increasing complexity. In situations where complexity is less, these neural network-based methods can be used as stand-alone object recognition schemes. As the complexity of the recognition task increases due to increased detail in the image, and/or due to degradations in imaging, the output of the system can still be used as an aid to human decision making.

REFERENCES

Barnard, E., & Casasent, D. (1991). Invariance and neural nets. *IEEE Transactions on Neural Networks, 2,* 498–508.

Fukushima, K. (1983). Neocognitron: A neural network for a mechanism of visual pattern recognition. *IEEE Transactions on Systems, Man and Cybernetics,* **13,** 826–834.

Carpenter, G. A., Grossberg, S., & Mehanian, C. (1989). Invariant recognition of cluttered scenes by a self-organising ART architecture: CORT-X Boundary Segmentation. *Neural Networks,* **2,** 169–181.

Giles, C. L., Griffin, R. D., & Maxwell, T. (1988). Encoding geometric invariances in higher order neural networks. In D. Z. Anderson (Ed.), *Neural information processing systems* (pp. 301–309). Denver, CO: AIP.

Grossberg, S., Mingolla, E., & Todorovic, D. (1989). A neural network architecture for preattentive vision. *IEEE Transactions on Biomedical Engineering,* **36,** 65–84.

Hu, M. (1962). Visual pattern recognition by moment invariants. *IRE Transactions on Information Theory,* **IT-8,** 179–187.

Khotanzad, A., & Lu, J. (1990). Classification of invariant image representations using a neural network. *IEEE Transactions on Acoustics, Speech and Signal Processing,* **38,** 1028–1038.

Lippmann, R. P. (1987). An introduction to computing with neural nets. *IEEE ASSP Magazine,* **2,** 4–22.

Prokop, R. J., & Reeves, A. P. (1992). A survey of moment-based techniques for unoccluded object representation and recognition. *CVGIP: Graphical Models and Image Processing,* **54,** 438–460.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In *Parallel distributed processing.* Cambridge, MA: MIT Press.

Widrow, B., & Winter, R. (1988). Neural nets for adaptive filtering and adaptive pattern recognition. *IEEE Computer,* **21,** 25–39.

Yegnanarayana, B., Mariadassou, C. P., & Saini, P. (1990). Signal reconstruction from partial data for sensor array imaging applications. *Signal Processing,* **19,** 139–149.