3-1-1992

Detection, location, and quantification of structural damage by neural-netprocessed moire profilometry

Barry G. Grossman

Frank S. Gonzalez

Joel H. Blatt

Jeffery A. Hooker

Detection, location, and quantification of structural damage by neural-net-processed moiré profilometry

Barry G. Grossman
Frank S. Gonzalez
Joel H. Blatt
Jeffery A. Hooker

Detection, location, and quantification of structural damage by neural net processed moiré prolifometry

Barry G. Grossman, MEMBER SPIE
Frank S. González, MEMBER SPIE
Florida Institute of Technology, Electrical and Computer Engineering Department
150 West University Boulevard, Melbourne, FL 32901-6988

Joel H. Blatt, MEMBER SPIE
Jeffery A. Hooker, MEMBER SPIE
Florida Institute of Technology, Physics and Space Sciences Department
150 West University Boulevard, Melbourne, FL 32901-6988

ABSTRACT

The development of efficient high speed techniques to recognize, locate, and quantify damage is vitally important for successful automated inspection systems such as ones used for the inspection of undersea pipelines. Two critical problems must be solved to achieve these goals: the reduction of the non-useful information present in the video image and the automatic recognition and quantification of extent and location of damage. Artificial neural network processed moiré profilometry appears to be a promising technique to accomplish this. Real time video moiré techniques have been developed which clearly distinguish damaged and undamaged areas on structures, thus reducing the amount of extraneous information input into an inspection system. Artificial neural networks have demonstrated advantages for image processing, since they can "learn" the desired response to a given input and are inherently fast when implemented in hardware due to their parallel computing architecture. Video moiré images of pipes with dents of different depths were used to train a neural network, with the desired output being the location and severity of the damage. The system was then successfully tested with a second series of moiré images. The techniques employed and the results will be discussed.

1. INTRODUCTION

A new real-time automated inspection technique to recognize, locate and quantify damage has been developed using a unique combination of video moiré techniques and artificial neural networks. Contour maps generated by non-contact real time video moiré techniques provide easily interpreted surface shape images. The automated inspection technique has been demonstrated using both a backpropagation network and a forward-only counterpropagation network. A software implementation using neural networks, is particularly useful for applications which have changing inspection requirements. The neural network architecture and training processes allow this inspection technique to be used in a wide variety of applications.

2. VIDEO MOIRÉ SYSTEM

Moiré patterns or fringes are seen when a regular grating projected on a surface is viewed through a similar grating. These patterns were first described by Lord Rayleigh[1], who speculated that they might be useful for surface inspection. In traditional projection moiré, a ruling or grating is projected upon a surface at an angle. The distorted ruling on the surface is then viewed at a different angle through an optical system containing an identical transmission ruling located in an intermediate image plane. The moiré system used in this paper is an outgrowth of our work with acousto-optic cells[2], and is unique in that it uses video technology to generate moiré patterns from a single variable spacing projection grating and does not utilize a second physical viewing grating[3]. Zoom projection technology is also employed to give a system with real time variable field of view and variable depth resolution.

A Michelson interferometer operated in an off axis tilted mirror format produces a series of equidistant straight lines separated by a distance of $\lambda/2\alpha$ and described by the equation

$$x = \frac{n\lambda}{2\alpha} - \frac{\varepsilon}{\alpha},$$

 0-8194-0751-8/92/$4.00

where λ is the wavelength, α is the angle between the mirrors, and ε is the on-axis mirror separation. If the interferometer is illuminated by a laser beam, the expanded straight line fringes can be projected on a target as variable resolution structured illumination. This interferometric grating projector produces the number of resolution elements needed for practical moiré work. In the classical projection or shadow system the distorted target grating is viewed through a physical reference grating. This reference grating has the effect of a binary optical filter in that where the grating is opaque no information is transmitted and where the grating is clear information is passed. It is the selectively transmitted distorted grating information along with the reference grating information that causes the classical moiré interference pattern. By using a special video processor this situation may be duplicated with a totally projection system. In this operation part of the video signal of one camera is made transparent to the video signal from a second source. This is accomplished by thresholding on a given signal level within the video signal. The result is to combine two video signals into one, not by superimposing the two but actually replacing part of one scene with another. The complete video moiré system is shown below in Fig. 1.

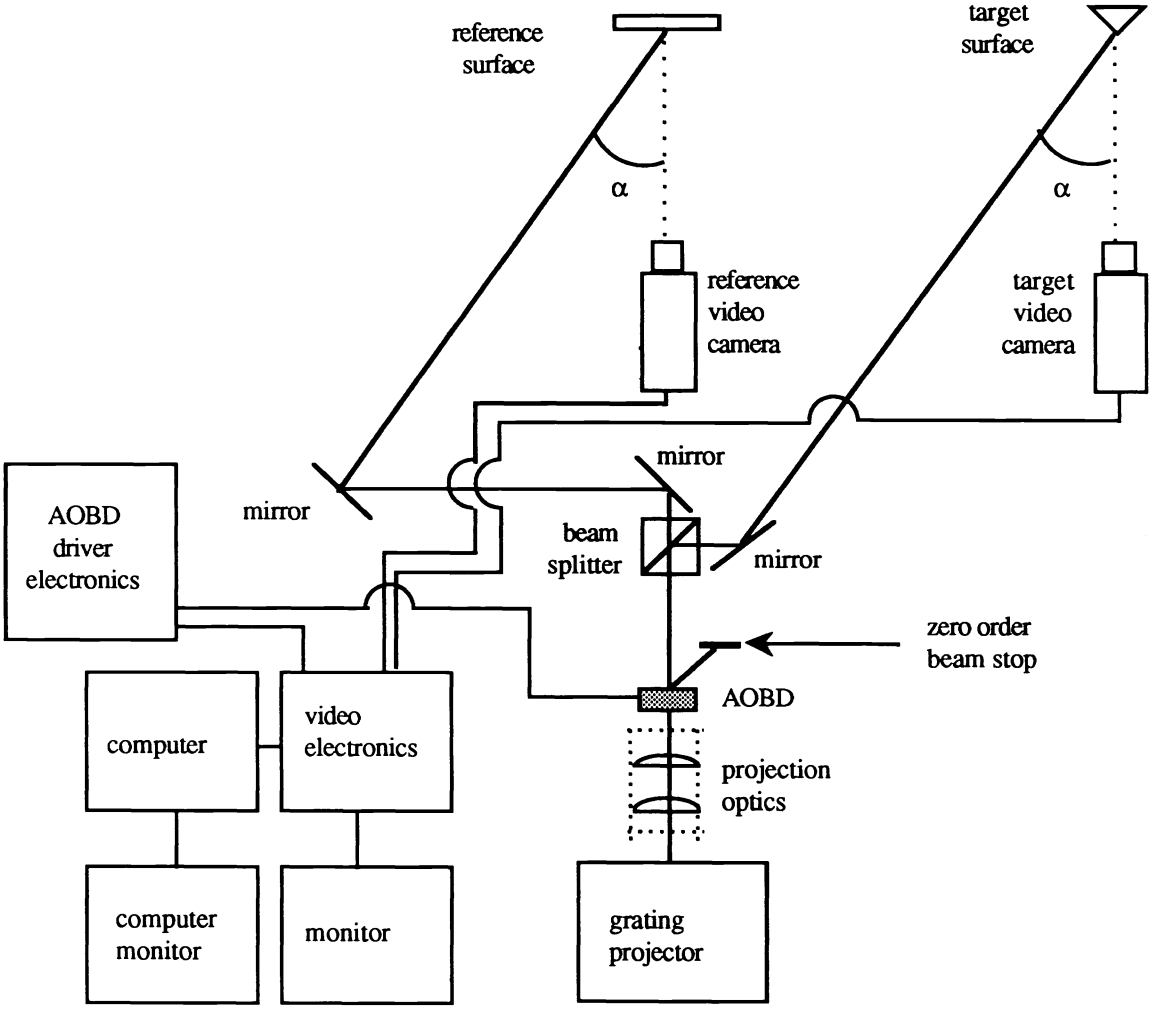

Fig. 1. Video projection moiré system.

The system consists of an Argon-ion illuminated Michelson interferometer grating projector with a beam splitter cube which separates the projected beam into the target and reference gratings. Viewing the target and reference surfaces from symmetric angles are two low light level, high resolution (650 TV lines) video cameras. The video signals from these cameras are fed to the video processor. The modified and unmodified video output from the video processor are delivered to two video monitors on which the original and altered video may be viewed simultaneously and to a computer which is used to collect and store data in the form of images via a video frame grabber board.[4]

3. ARTIFICIAL NEURAL NETWORKS

Research on artificial neural networks has boomed in recent years as a result of the discovery of various learning algorithms and the success in applying these networks to solve difficult problems. In general, most artificial neural networks consist of layers or "arrangements" of simple processing units or "neurons" interconnected by weighted links. The architectural organization and the respective training procedures among these networks varies considerably. Current research efforts at the Florida Institute of Technology and many other international institutions are yielding an ever increasing amount of variations and new approaches.

The structure of our backpropagation trained feed-forward network is shown in Fig. 2. The backpropagation learning algorithm is a gradient descent method used to search for the optimal parameters in multilayer feed-forward neural networks.[5] Artificial neural network-based applications have proven to be of comparable or superior performance to conventional methods in diverse areas such as pattern recognition[6], image compression[7], vehicle navigation[8], speech reproduction[9], and sensor and fiber optic processing.[10-12]

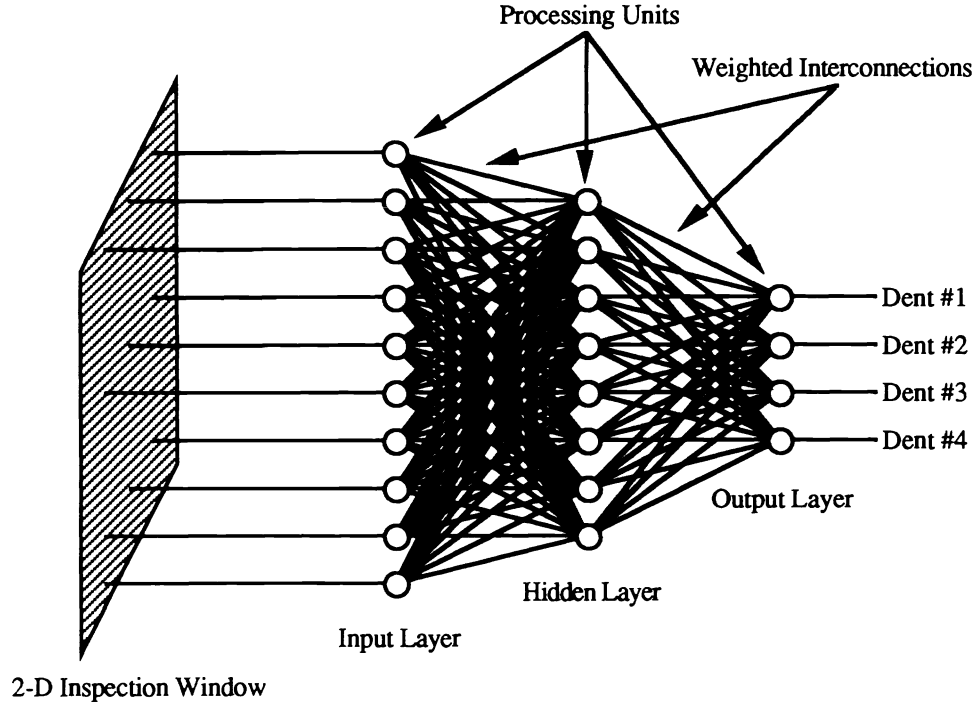

Fig. 2. Feed Forward Network.

Our feed forward network consists of three fully interconnected layers of processing units. The input layer serves as an interface between the 2-D moiré images and the network. Each processing unit in the first layer distributes its input value towards all the second (hidden) layer units through weighted links. The input of each of the hidden layer units is the sum of the product of the input signals and the connection weights through which they pass, given by:

$$zin_j = \sum_i (x_i v_{ij})$$

where x_i is the value of the i^{th} input layer unit and v_{ij} is the weight located between the i^{th} input layer unit and the j^{th} hidden layer unit. Each hidden layer unit applies a nonlinear activation function to its net input. The bipolar sigmoid was used as the activation function, given by:

$$Zout_j = \left[\frac{2}{1 + e^{-\sigma_j zin_j}} \right] - 1$$

where $zout_j$ is the output signal of the j^{th} hidden layer unit, σ_j is an adjustable slope parameter associated with the j^{th} hidden layer unit, and zin_j is the input signal of the j^{th} hidden layer unit. These activation signals are sent thru weighted connections towards the output units . As with the hidden layer units, output layer units receive the sum of the product of hidden layer activations and their associated weights:

$$yin_k = \sum_j (zout_j w_{jk})$$

An activation function is applied to the net input of the output units, producing the outputs. The bipolar sigmoid function, as described above, was used with the output layer.

The backpropagation learning algorithm is used to train the network to perform the required mapping approximations. This algorithm is a supervised process that requires multiple presentations of the training set. During training, the training set images are presented serially at random order and the computed output is compared to the expected target output. The difference between the expected and the computed outputs, e_k, is propagated back to the network and serves as the basis for the weight updates. The update for the weights between the j^{th} hidden layer and k^{th} output layer is given by:

$$\Delta_{jk} = \alpha\, e_k\, f'(yin_k)\, z_j$$

where α is a learning rate factor ($0 < \alpha < 1.0$), e_k is the error of the k^{th} output layer unit , $f'(yin_k)$ is the derivative of the activation function of the k^{th} output layer unit, and z_j is the output of the j^{th} hidden layer unit. The update for the weights between the i^{th} input and the j^{th} hidden layers is determined in terms of the error signals of the units and weights to which they connect. This update is given by:

$$\Delta_{ij} = \alpha \sum_k [e_k\, f'(yin_j)\, f'(yin_k)\, w_{jk}]\, x_i$$

where $f'(yin_j)$ is the activation derivative of the j^{th} hidden layer unit, w_{jk} is the weight between the j^{th} hidden layer unit and the k^{th} output layer unit, and x_i is the output of the i^{th} input layer unit. All the weights are updated after the end of each backward pass. Once the network has successfully converged, the weights remain constant for inspection or further training.

The forward-only counterpropagation network, the second network we have implemented, is illustrated in Fig. 3. This network is the combination of a self organized topological feature map[13-14] (Kohonen network) and an additional layer that uses the Grossberg outstar learning rule.[15] The Kohonen's self organizing neural network is able to cluster input vectors according to the similarities in their components. Vector patterns that are most similar would cluster into the same vector code (Kohonen's cluster unit). A competitive learning rule is used to choose a winning cluster unit. The winning cluster unit is selected based on the closeness between the input and weight vectors. This closeness can be determined by the euclidean distance or dot product. The cluster unit with the best input to weight correlation and some of its neighbor units are allowed to update their weights. The weight update for the winning clustering unit is given by:

$$v_{ij}^{new} = v_{ij}^{old} + a(t)[x_i - v_{ij}^{old}]$$

where v_{ij} is a weight connecting the i^{th} input and the j^{th} clustering (Kohonen) units, a(t) is a time-dependent decreasing learning rate ($0 < a(t) < 1.0$), and x_i is the i^{th} input unit. The topological structure of the neighborhood may be one or two dimensional. The initial neighborhood should start large enough to bring the weight vectors closer to the region of the input vectors. This neighborhood is gradually reduced to allow update only for the weights associated with the winning units. The winning unit is allowed full weight update throughout the training process. The amount by which the neighboring units update weights is usually lower than for the winning unit. In our implementation, we allowed the two neighboring units to update their weights half the amount of the winning unit. Once trained, the interconnection weights between the input and clustering layer will contain the information coding needed to perform the clustering. Kohonen's self-organizing feature maps have been successfully applied to pattern classification, speech recognition, and smart structures sensing.[16-18] The additional layer connected to the clustering units of the Kohonen layer serves to associate the winning cluster[s] into specific output units. The weight update for this fully interconnected layer is given by:

$$w_{jk}^{new} = w_{jk}^{old} + \beta(t)[y_j - w_{jk}^{old}]K_j$$

where w_{jk} is the weight between the j^{th} clustering unit and the k^{th} output unit, y_j, is the expected state for the k^{th} output unit, K_j is the output of the j^{th} clustering unit, and ß(t) is a time-dependent decreasing learning rate ($0 < ß(t) < 1.0$). This equation is known as the Grossberg's outstar learning rule and is a special case of the delta rule for when only one competition unit wins. Training of the output layer starts at the end of the Kohonen training, once stable clustering formations have been developed at the Kohonen layer. Only a few training passes are required for the weights in this output layer to reach an optimum state.

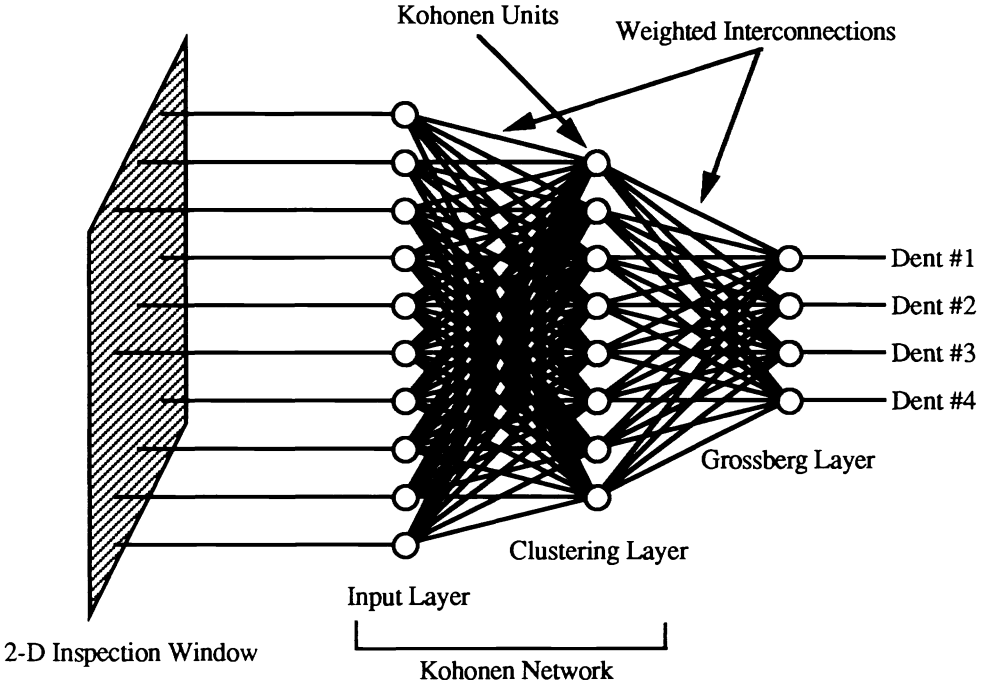

Fig. 3. Forward-Only Counterpropagation Network.

4. EXPERIMENTAL SET-UP

The targets, consisting of 4 cm diameter pipes 12.5 cm long, were mounted an a translation-rotation stage which allowed each pipe to be translated horizontally and rotated about the translation axis. The pipes were heated and dented to several depths (given in Table I) by pressing a 6.35 mm (0.250 in) ball bearing into the pipe with an arbor press.

Table I. Depths of the dented pipes.

Dent	Depth (mm)
(a) Pipe #0	no dent
(b) Pipe #1	2.2479
(c) Pipe #2	2.3622
(d) Pipe #3	3.6322
(e) Pipe #4	4.5297

The pipes were then sprayed with flat grey paint to cover the marks made by heating the pipes and to give them a uniform appearance. As can be seen in Fig. 4, the dents are nearly invisible without the structured illumination.

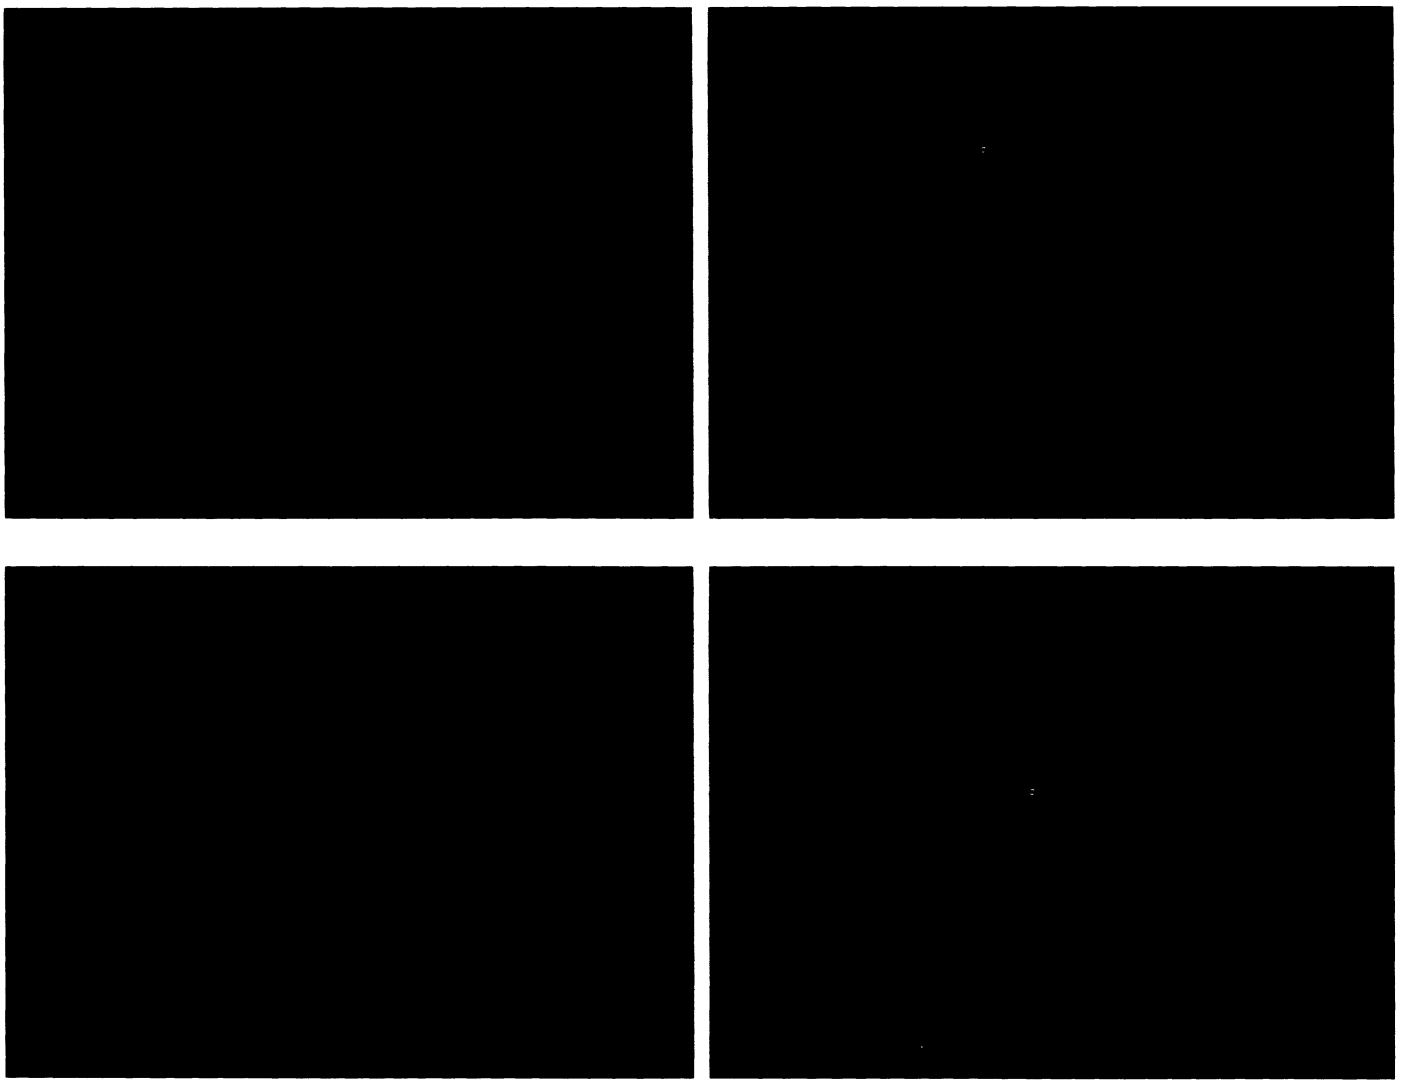

Fig. 4. Dented pipes.

Fig. 5 shows the moiré contours of the dented pipes. A Macintosh II microcomputer equipped with a Quick Capture frame grabber digitized the images which are then stored in 8-bit greyscale TIFF files with the aid of Image, a digital image processing package.[19] The intensity of the digitized greyscale images range from 0 (white) to 255 (black). The images were then preprocessed by averaging, sampling, and scaling. These capabilities were added to the Pascal source code of *Image* in order to reduce memory and computational time requirements involved during training or inspection. We also added code to convert the preprocessed images into text files that can be ported to non-Macintosh computer environments. Both neural networks and accompanying weight generation codes have been successfully ported from our original Pascal implementation on a Macintosh to various computer systems including a VAX 8200, Harris HCX-9, and CRAY Y-MP. A backpropagation neural network, already added to *Image*, would serve as the basis for the prototype autonomous inspection system.

5. SIMULATION RESULTS

5.1 Initial Results

During the first stage of our research, we studied the performance of the backpropagation neural network to locate the center of a fixed size, continuously valued spot positioned along the horizontal axis of images composed of 50x640 pixels. The area

surrounding the spot was set to white (0 intensity). The spot's diameter was approximately 48 pixels and its intensity varied from 39 to 213. The original images were preprocessed from 50x640 pixels down to 13x160 values. We scaled the intensity distribution to continuously-valued real numbers in the interval [-1.0, 1.0].

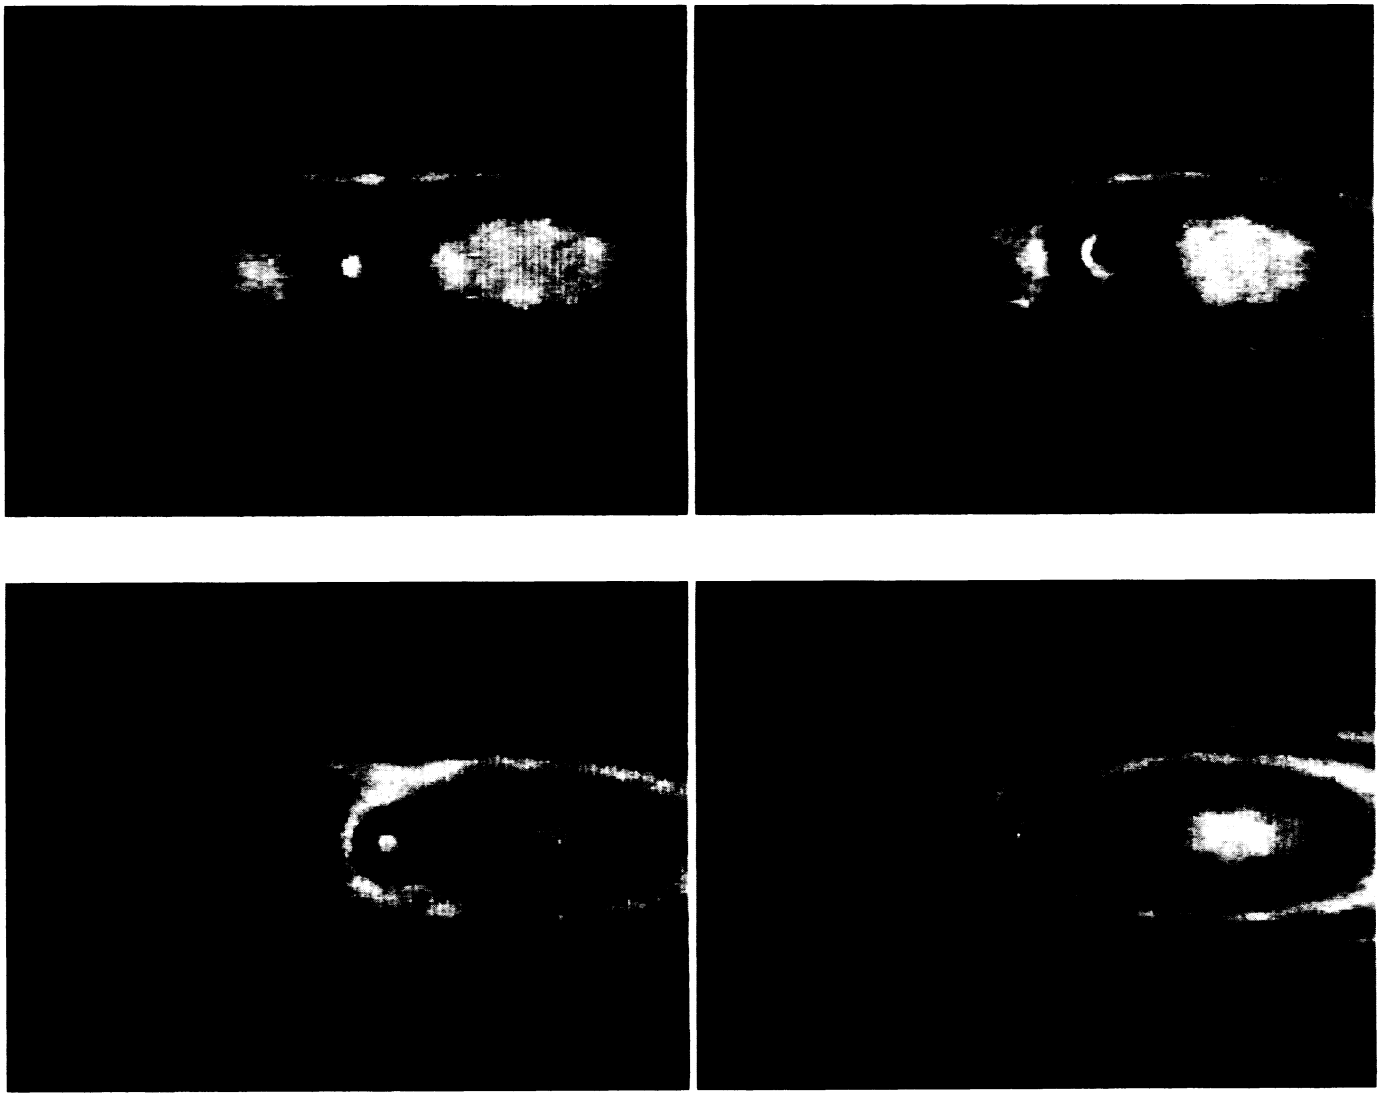

Fig. 5. Moiré contours of dented pipes.

With our first approach, we tested the entire image at once by feeding a pixel to each unit in the input layer. The network consisted of 2080 input layer units, a variety of hidden layer units, and one unit in the output layer. The first training sessions involved 10 preprocessed images. The training errors were typically in the vicinity of 1% for 1000 iterations. The network was able to accurately locate the spot location for the training set. However, test results for images with spots at locations between the training spots did not correlate well with their actual location. Further tests were performed with more training images and different combinations of the initial parameters. Although the selection of initial parameters proved to affect convergence, test results remained unacceptable. Increasing the number of training images from 10 to 18 did improve test results, while the convergence speed remained unchanged. A training set of 49 images slightly improved test performance, but increased the convergence time. In our last attempt with this technique we increased the training set to 67 images, but the test results remained comparable to the results with 49 images, while the convergence time increased further. We concluded that the inability of the network to correctly locate the spot at locations other than those of the training set was due to the network's stringent requirement to process the whole image at once.

We decided to approach the problem differently, by training the network with two smaller rectangular image segments: one with the spot and the other without the spot. These training images were preprocessed from 50x50 to 13x13 pixels by averaging and sampling. The intensity distribution was scaled to continuously-valued real numbers in the interval [-1.0, 1.0]. The network was trained to produce a high output if the spot was on the center or a low output if there was no spot. Fast training convergence was achieved since now the training set was much smaller; only two preprocessed images were involved each consisting of only 169 values. For damage detection, an additional algorithm translated the testing pipe images through this input layer window while the network analyzed the image. With this method the network was able to locate the spots to within 1 pixel using backpropagation. Memory requirements with this approach were low enough and permitted us to implement a forward-only counterpropagation network. Test results for this network came out as accurate as with backpropagation.

5.2 Moiré Inspection Results

Having validated our neural network techniques, we then began working with the moiré images. We trained the networks to quantify the size of the defects into five categories. Fig. 6 shows the moiré images of the five "window" segments, one undented and four which have a dent of differing depth, that served as the initial training set.

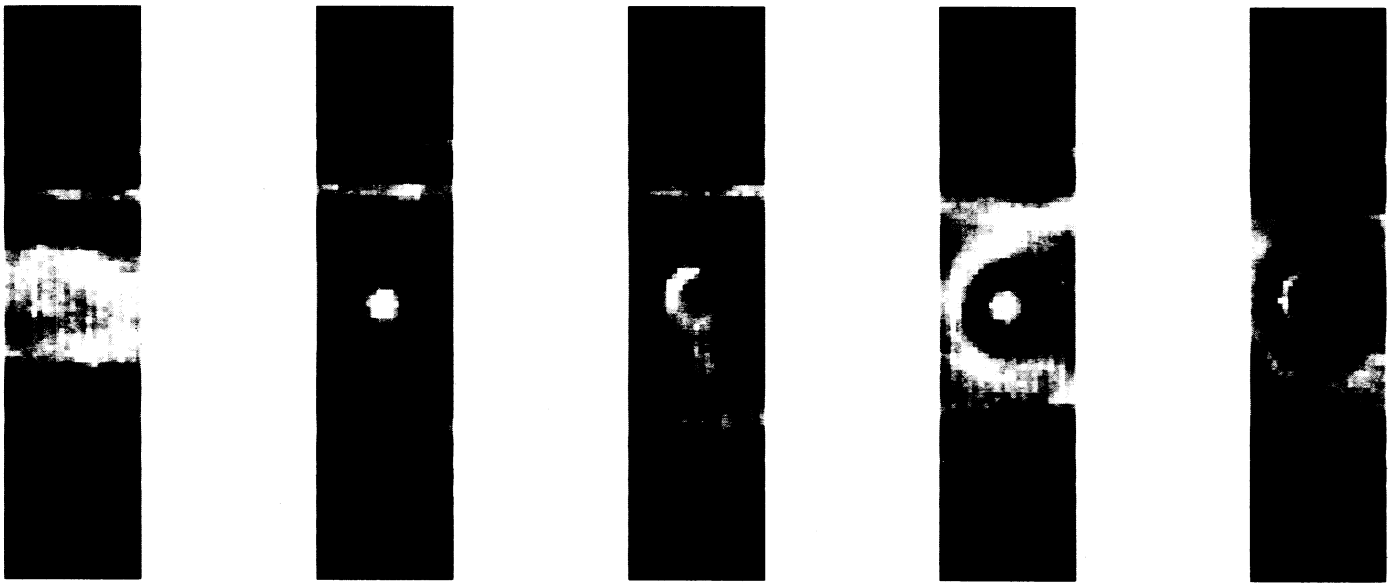

(a) (b) (c) (d) (e)

Fig. 6. Sample training images.

These rectangular windows were preprocessed from 410x100 pixels to 160x25 pixels by averaging and sampling. The intensity of distribution was scaled to [-1.0, 1.0]. Fig. 7 illustrates the pipeline inspection concept using the moving window approach. As the window moves along the pipe, the network produces four output signals that indicate the presence of a dent and its depth. The depth of the dent is given by one of the four output units. The dent location is found at the position with the largest output. Under most inspection conditions these outputs could be biased to prevent the system from false alarms due to low or ambiguous responses.

Our initial results demonstrated poor inspection capability, specially for the pipes with the largest dents. Although the backpropagation network had converged to under 1.0% error for all the training images, inspection results yielded erroneous results all along the pipes. The forward-only counterpropagation results were less confusing but still unacceptable. For both networks, results were only satisfactory at the location where the moving inspection window exactly matched the location of the dents. Careful observation of these results made apparent the major reasons for the observed poor performance. The intensity variation along the horizontal axis, which typically increased linearly from greyscale values of 170 at the left to 90 at the right, made obvious the requirement of training images from multiple locations along the pipes. Inspection results using these

additional training sections demonstrated dramatic improvement, especially along the undented regions of the pipes. The erroneous signals that persisted at the locations where the inspection window approached the dents were eliminated by including training images containing dents shifted 25 and 50 pixels to both sides.

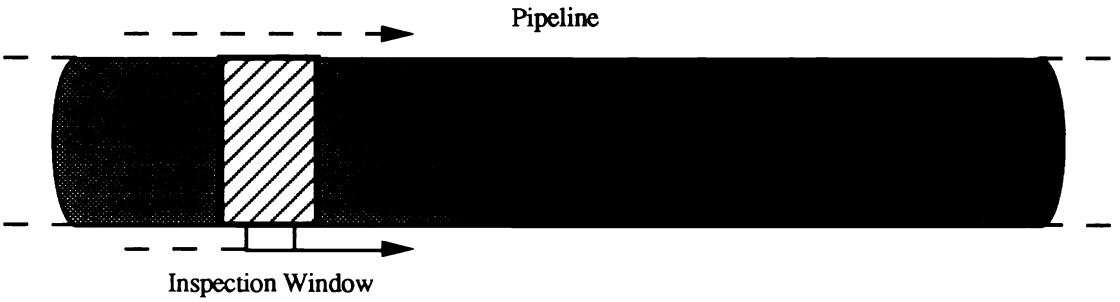

Fig. 7. Pipeline inspection using the moving window approach.

We have also found that the initial selection of various parameters particularly learning rate, number of hidden layer units, and training set were important factors determining the convergence and accuracy of the backpropagation network. A learning rate in the vicinity of 0.1 generally produced best results. Learning rates higher than 0.2 caused saturation making the network unable to converge. We have tested the backpropagation network with 5, 10, 15, 20, 25 and 50 units in the hidden layer. Ten or more hidden layer units consistently produced the best results. Inspection results with 5 units in the hidden layer contained undesirably high responses. The selection of the initial random weight range was a definitive determining factor that affected the inspection results. The network was able to converge to less than 1.5% error for several initial weight ranges. However, an initial range in the interval [-0.1, 0.1] consistently demonstrated superior results. Typical inspection results using the backpropagation network are shown on Fig. 8.

As the window moves along the pipe, the network produces four output signals that indicate the presence of a dent and its depth. At the location of a dent, the output unit with the highest correlation produces a high signal, while the other output units remained low. Each of the plots in Fig. 8 shows the inspection results for a network with four output units. The calculated location (in pixels) of the dents, as indicated by the peak output signals, were 330 for dent #1, 370 for dent #2, 362 for dent #3, and 330 for dent #4. Fig. 8 (a) shows the results for an image with the smallest dent. The calculated dent location came out only three pixels away from the actual location. The signals from the other three output units were very low and rarely became noticeable, as is the case at location 80. The results with the other dented pipes (plots b, c, d) demonstrated similar characteristics. The small errors are due to the data reduction involved during preprocessing. The inspection plot for the undented pipe is not shown since all the output signals were lower than 0.02.

For all our tests using the backpropagation network the correct output unit always produced the highest signal. The performance of the networks was tested by studying the effects of noise and up/down shifts. Noise was added by randomly "spraying" the images with various pixel intensities. Inspection results with noisy images slightly reduced performance. The signals from the correct output units were typically reduced from 0.99 to 0.95. The results with the up-shifted images produced correct output signals that ranged from 0.2 to 0.8. The inspection results with the down-shifted images produced correct output signals ranging from 0.5 to 0.6 for the smaller two dents, and roughly 2 to 4 times the noise level for the larger dents. Although the expected output signal for some of the down-shifted dented images were lower than 0.1, it always much higher than the other outputs. Biasing the output units could be used as a selection parameter to help determine the existence of a dent.

Inspection results with the forward-only counterpropagation network were as encouraging as with backpropagation. The results with the noisy images successfully identified dent depth and location. The up-shifted test produced correct output signals that ranged from 0.5 to 0.7 with no trace of erroneous signals. The results with the down-shifted images were slightly better than with the backpropagation network, since the correct output signals ranged from 0.5 to 0.7. However, the up-shifted test results frequently indicated the presence of a different dent next to the correct dent location.

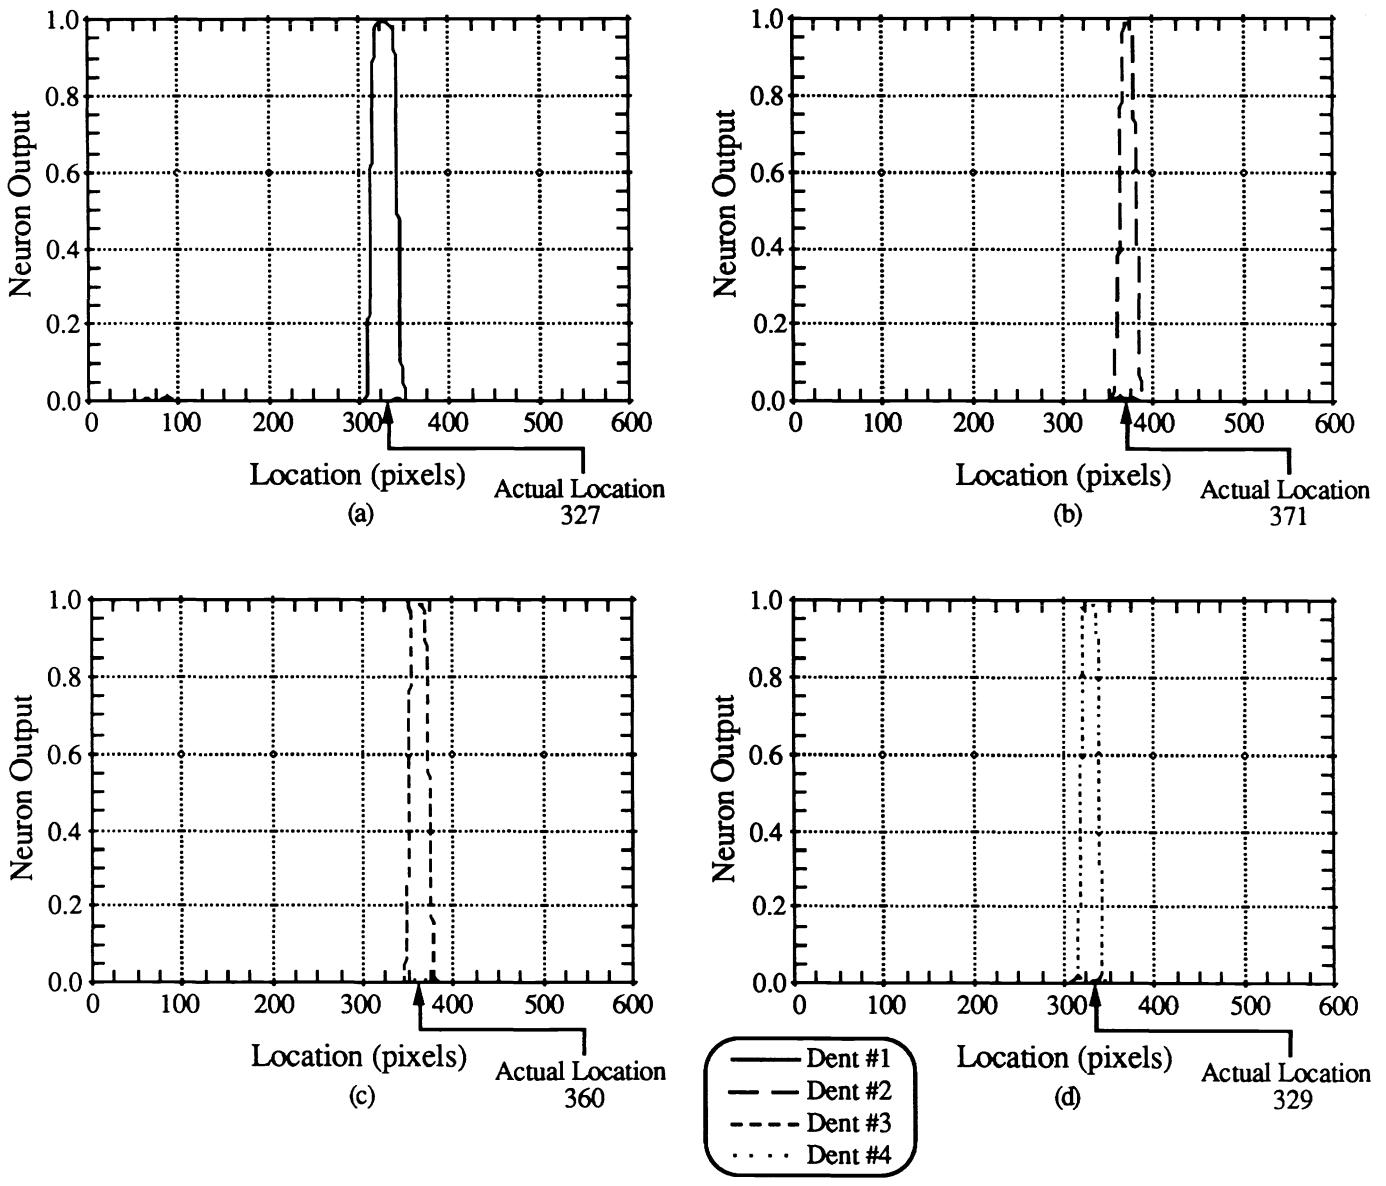

Fig. 8. Measured inspection results using the backpropagation network.

Typical cpu time consumption for both networks are shown in Table II. The backpropagation network consisted of 2575 input units, 10 hidden layer units, and four output units. The forward-only counterpropagation network consisted of 2575 input units, 29 clustering units, and four output units. For both networks, the training set consisted of 29 images. The backpropagation network required 15000 iterations to reach errors lower than 1.5%, while the forward-only counterpropagation networks only required 1500 to achieve stable clusterings.

The forward-only counterpropagation network required less cpu time because a lower number of iterations. Although the backpropagation network required more training time, inspection time per image was faster due the lower number of units in the hidden layer. Training time requirements can be reduced by modifications of the networks and by proper selection of the training set. Inspection time can be reduced by scanning the pipe at larger step sizes, i.e. skipping several pixels until high output signals are produced. With this modification, inspection time per image using a Mac IIci can be reduced from 1:48 min to 40 sec. Further inspection time reduction can also be achieved by using a different size for the inspection window and by modifying the network.

Table II. CPU time consumption during training and inspection.

	Backpropagation		Forward-only Counterprop	
Computer	Training 29 images 15000 iters	Ispection per image	Training 29 images 1500 iters	Ispection per image
Mac IIci	4:43 hr	1:48 min	50:01 min	4:04 min
Harris HCX-9	1:50 hr	25 sec	16:20 min	76 sec
CRAY YMP	43.5 sec	0.42 sec	1.77 sec	0.64 sec

6. CONCLUSIONS

This automated inspection technique, although still in its infancy, has demonstrated encouraging performance. Current work concentrates on improving the generalization and in reducing training and inspection time. We are also developing software to locate damage in two dimensions (both two axis displacement and one axis displacement plus rotation). We are confident that as our research continues, we will be able to reduce the processing required so that this technique can be implemented on a microcomputer, and will provide the high performance real time inspection capability required for many applications.

7. ACKNOWLEDGMENTS

This work has been partly supported by the National Science Foundation under the award number ASC910034P. The authors would like to thank Jutta Baerman, Fred Bristol, and Charles Inman for their valuable help with the preparation of the pipes and translation-rotation stage, Leris Padrón and James Wells for their technical assistance with the Harris HCX-9 workstation, and Dr. Laurenne V. Fausett for her lectures and advice on neural networks.

8. REFERENCES

1. Lord Rayleigh, "On the Manufacture and Theory of Diffraction-gratings," The Phi. Mag. of Journal of Science, vol. 47, pp. 81-93, Feb. 1874.
2. Blatt, J.H., Ho, H-C.C, Young, E.H: "Generation of Moiré Contours with Acoustic-Optic Cells," Optical Engineering, 28:9:996-998, September 1989.
3. Blatt, J.H., Hooker, J.A, Young, E.H, Belfatto R.V., "Video Applications to Moiré Metrology," Journal of Laser Applications, vol. 2, no. 3 and 4, 35, Summer/Fall 1990.
4. Blatt, J.H. , Hooker, J.A., Ho, H.-C. C., and Young, E. H., "Signal-to-Noise Improvements In Moiré Prolifometry by Acousto-Optic Scanned Gratings and Video Processing", Paper 1614-26 to be presented at the SPIE conference on "Optics, Illumination, and Image Sensing for Machine Vision VI" at ROBOTICS '91, Advances in Intelligent Robotic Systems, 10-15 November, 1991, Boston Massachusetts.
5. Rumelhart, D.E., Hinton G.E., Williams R.J., "Learning Internal Representations by Error Propagation," in PDP: vol. 1, Chapter 8, MIT Press, Cambridge, MA, 1986.
6. LeCun, Y., Boser, B., Denker, J.S., Henderson, D., Howard, R.E., Hubbard, W., Jackel, L.D., "Backpropagation Applied to HandWritten Zip Code Recognition," Neural Computation 1, pp. 541-551, 1990.
7. Cotrell, G.W., Munro, P., Zipser, D., "Learning Internal Representations from Gray-Scale Images: An Example of Extensional Programming," 9th Conference of the Cognitive Science Society, Seattle, 1987, pp. 462-473, Hillsdale: Erlbaum.
8. Pormeleau, D.A., "ALVINN: An Autonomous Land Vehicle in a Neural Network," Advances in Neural Information Processing Systems I, ed. D.S. Touretzky, pp. 305-313, San Mateo: Morgan Kaufmann, Denver, 1988.
9. Sejnowsky, T.J., Rosenberg, C.R., "Parallel Networks that Learn to Pronounce English Text," Complex Systems 1, pp. 145-168, 1986.

10. Grossman, B., Hou, H., Nassar, R., Ren, A., Thursby, M., "Neural Network Processing of Fiberoptic Sensors and Sensor Arrays," Proceedings of SPIE OE/Fibers '90, September 1990.
11. Grossman, B., Nassar, R., Hou, H., "Electrooptic Neural Processors for Smart Sensors and Structures," Proceedings of IEEE Southcon '90, September 1991.
12. Grossman, B., Hou, H., Nassar, R., "Optical Processors for Smart Structures," Proceedings of SPIE OE/Aerospace Sensing '90, April 1990.
13. Kohonen, T., "Self-Organized Formation of Topologically Correct Feature Maps," Biological Cybernetics, vol. 43, pp. 59-69, 1982.
14. Kohonen, T., Representation of Sensory Information in Self-Organizing Feature Maps and Relation of these Maps to Distributed Memory Networks," Optical & Hybrid Computing, SPIE vol. 634, Bellingham, WA: SPIE, 1987.
15. Hecht-Nielsen, R., *Neurocomputing*, Addison-Wesley, 1990.
16. Lu, S-Y., "Pattern Classification using Self-Organizing Feature Maps," 1990 IEEE INNS International Joint Conference on Neural Networks. vol, I, pp. 471-480.
17. Kohonen, T., Mäkisara, K., Saramäki, T., "Phonotopic Maps -Insightful Representation of Phonological Features for Speech Recognition," Proceedings of the 7th International Conference on Pattern Recognition, Montreal 1984, pp. 182-185, NY:IEEE.
18. Grossman, B., Gao, X., Thursby, M., "Composite Damage Assessment Employing an Optical Neural Network Processor and an Embedded Fiberoptic Sensor Array," Proceedings of SPIE OE/Fibers '90, September 1991.
19. Image by Wayne Rasband: Internet, BitNet: wayne@helix.nih.gov (CompuServe: 76067,3454)