

General-purpose acousto-optic connectionist processor

Thomas Naughton, MEMBER SPIE
Zohreh Javadpour, MEMBER SPIE
John Keating, MEMBER SPIE
Department of Computer Science
National University of Ireland
Maynooth, Ireland
E-mail: tomn@cs.may.ie

Miloš Klíma
Jirí Rott
Department of Radioelectronics
Faculty of Electrical Engineering
Czech Technical University
Prague, Czech Republic

Abstract. We present a reconfigurable optical implementation of an acousto-optic algebra processor, based on a calomel (single crystal mercurous chloride) matrix-vector multiplier. Two successful applications are demonstrated: neural network training and curve detection. In the first, a perceptron learns two-input Boolean functions. In the second, a connectionist model of the Hough transform is generalized to handle arbitrary curves. Experiments for line detection and circle detection are performed, and the Hough transform's throughput is solely limited by the update latency of the acousto-optic unit. © 1999 Society of Photo-Optical Instrumentation Engineers. [S0091-3286(99)01007-7]

Subject terms: acousto-optics; optical processor; optical neural network; Hough transform; calomel.

Paper PPB-01 received Nov. 8, 1998; revised manuscript received Jan. 6, 1999; accepted for publication Jan. 28, 1999.

1 Introduction

The application of artificial neural networks (ANNs), or experimental analyses of neuron dynamics, require hardware solutions suitable for large-scale implementations. In particular, reconfigurable implementations that facilitate increasing numbers of network interconnections are highly desirable. Electronic and optical implementations have been intensively investigated with many proposed architectures presented in the literature.¹⁻⁵ It has been shown that optical implementations provide a more suitable solution to the problem of increased interconnections. The optical implementations are not without problems, however. Research has shown that imaging aberrations and light detection issues severely restrict the number of neurons available in the proposed architectures.⁶⁻⁹

Recent experiments by Gao et al.¹⁰ use lenslet arrays in a coaxial arrangement to successfully implement an optical neural network with 32×32 neurons. Mori et al.¹¹ have an optical matrix-vector multiplier that uses 2-D (two-dimensional) structures [liquid crystal display (LCD) panels] to implement a two layer network. Our system uses a matrix-vector multiplier based on a combination of LCD and acousto-optic (AO) devices and can be used for both types of connectionist processing (Hopfield and multilayer). In common with other implementations, the matrix of weights is displayed on a 2-D LCD panel, but in contrast, the input vector is encoded in the acoustic wave of an AO unit with a calomel substrate. Calomel (mercurous chloride) is the best developed member of a promising class of photoelastic materials, the mercurous halides, whose unusual optical and acoustic properties may significantly extend the performance of AO devices.¹²⁻¹⁶

In this paper, we describe the experimental model and present the successful results of implementing two sets of experiments using the architecture. The first involves the training of a single-layer perceptron ANN using several two-input logical functions. In the second set, we imple-

ment a Hough transform (HT). We employ the connectionist model of Seth and Datta¹⁷ for our HT and extend their work by utilizing reconfigurable SLMs to enable weight updates and increasing the dimensions of the input space. In addition, we extend Seth and Datta's straight line implementation to one that detects circles and show that our fully generalized model can be used to detect any arbitrary shape.

1.1 Calomel

A class of photoelastic materials, the mercurous halides, exhibit unusual optical and acoustic properties that look promising for many signal processing and spectrum analyzing tasks. Single crystal mercurous chloride (Hg_2Cl_2), or calomel,¹²⁻¹⁶ is the best developed member of this class and boasts a high spectral transmission range, low acoustic attenuation and a very low shear mode acoustic velocity. In addition, calomel combines the most advantageous property of each of tellurium dioxide and fused quartz (two more widespread AO materials) with its high AO figure of merit M_2 (Ref. 18) and capability to withstand large optical power densities,¹⁹ respectively.

Calomel AO cell fabrication details and experimental data,¹²⁻¹⁵ comparisons between the properties of calomel and various other AO materials,¹⁶ and details of our particular calomel AO cell²⁰ can be found in the literature. Although a reliable preparation technique¹² has been developed for mercurous chloride single crystals, the material's low hardness and high fragility cause significant difficulties when fabricating an AO unit.

2 Architecture

The operation of our processor can be visualized from the schematic in Fig. 1. Using optical components each row of the matrix is directed through, and is thus multiplied by, the vector encoded in the second SLM. Summing the resulting values in each row produces a vector representing the inner




Fig. 1 Space-integrating AO matrix-vector multiplier.

matrix-vector product. The implementation of our space-integrating²¹ optical processor can be seen schematically in Fig. 2 and in the photograph in Fig. 3. The values of the weight matrix and the image vector are controlled by a computer program executing on a personal computer.

As shown in Fig. 2, the collimated beam from the HeNe laser is first expanded. The beam expander consists of lenses L1 and L2 and has an expansion ratio of 1:50. In the vertical direction (side view), cylindrical lens L3 operates as a Fourier lens in the standard 1-*f* setup, such that the matrix columns are Fourier transformed at L3's focal plane. The AOU's acoustic beam was centered along this plane such that it coincided with the zero-order of the LCD signal's 1-D Fourier transform. In the horizontal direction (top view in Fig. 2), lens L3 does not influence the laser beam. The input vector is encoded in the acoustic wave of the AOU, which operates in the Bragg regime. As a result, the first AO diffracted order (given sufficient time and space to spatially expand) is a 2-D signal containing the multiplication of the matrix coefficients from each row of the LCD and this input vector. The row summation and imaging on the detector is performed by lenses L4 and L5 and the sampled vector interfaced to a PC.

The system consists of a 20 mW HeNe laser, an amplitude modulating Sony LCX012AL TFT (thin film transistor) LCD panel (640×480 pixels), and an 8-bit intensity resolution Panasonic WVCD50 CCD (charge-coupled device) camera with the automatic gain functionality disabled. By extracting a predefined pattern of pixels from our 2-D CCD image we could simulate a 1-D detector of arbitrary width. The AOU consists of a calomel substrate orientated for a shear mode acoustic beam, with a 75 MHz center




Fig. 2 Schematic of the optical setup: LCD, matrix spatial light modulator; AOU, AO unit encoding a vector; and CCD, detector.




Fig. 3 Photograph of the optical setup.

frequency. The crystal was grown by BBT Materials Processing, Czech Republic, and the cell fabricated by Isle Optics, United Kingdom. Through tests with step-function gradient intensity images of varying dynamic ranges we have determined empirically that our LCD/AOU/CCD combination enables 6 bits of intensity resolution on the LCD panel.

2.1 Fourier Optical Model

Coherent optical systems are straightforward to analyze since they can be approximated using the equations of Fourier optics. If we ignore scaling factors due to lens configurations and let (x,y) describe each space domain and (α,β) describe each frequency domain along our optical system, the detected output signal s_{DET} can be written as

$$s_{DET}(\alpha, -y) = \left| C(\alpha) \int \int_{-\infty}^{\infty} \left[A(\beta) g(x) \int_{-\infty}^{\infty} f(x,y) \exp(i2\pi\beta y) dy \right] \times \exp[i2\pi(\alpha x + \beta y)] dx d\beta \right|^2, \quad (1)$$

where f is the LCD panel signal and g is the signal encoded in the acoustic beam. The 1-D signal A is a binary-valued spatial frequency bandpass filter, infinite in the x -direction, which corresponds to the height of the acoustic beam along its length. We presume that the acoustic beam undergoes negligible absorption and divergence. If such divergence is to be modeled, then a real-valued 2-D filter should be used for A . Finally, C is a binary-valued 1-D spatial frequency bandpass filter, infinite in the y -direction, which corresponds to the width of the linear CCD sensor.

2.2 Simulated Detector Plane Signals

Using the preceding idealized 3-*f* model, several simulations of the output plane signal were performed with MATLAB from The Mathworks Inc. We see from these experiments that the limitations of a finite acoustic beam height and finite linear detector width are extremely important considerations when implementing a coherent AO matrix-vector multiplier.

The first set of experiments sought to determine the effect of a finite AO interaction height. Figure 4 contains a sample optical signal immediately behind the LCD panel, which combines the LCD sampling frequency with an arbitrary modulation function. The acoustic beam acts as a bandpass filter on the Fourier spectrum of each matrix column. Figure 5 shows the effect of a finite AO interaction



Fig. 4 Input signal for the interaction height experiments that simulates a portion of the LCD panel with arbitrary pixels turned on or off.

height on spatial resolution in the output plane. As less vertical spatial frequencies make it through the system, the separate horizontal spectra of the output signal (which correspond to individual summed rows of the matrix-vector point multiplication) become less distinguishable. The implications for numerical computations are obvious. An acoustic beam whose height is not sufficient to pass enough vertical spatial frequency information will cause the individual output vector data values to overlap in the detector plane. If the AO interaction height cannot be changed, then




Fig. 5 Results from the interaction height experiments when Fig. 4 is the input plane signal: (a) the unfiltered signal appearing in the AO plane, (b) corresponding output plane signal showing clearly separated rows of information, (c) the spatially filtered signal due to a finite acoustic beam height, and (d) corresponding output plane signal.




Fig. 6 Results from the detector width experiments: (a) the input signal immediately behind the LCD panel, (b) the corresponding output plane signal assuming a detector of arbitrary width, and (c) the input signal as it appears to a detector of finite width which only measures a portion [that shown in (d)] of the output plane signal.

either a new Fourier lens configuration to rescale the frequency spectrum or a spatial separation of the inputs in the LCD plane is required.

The model also tells us that the diffracted orders due to the vertical pixel structure do not affect our computation. Figure 5(a) contains three horizontal orders, -1 , 0 , and $+1$. By filtering in the vertical direction [Fig. 5(c)] we ensure that the -1 and $+1$ orders, which contain the vertical sampling frequency information, do not propagate further through the system.

A second set of simulation experiments were undertaken to determine the effects of employing a linear detector of finite width. As with the previous experiments, the issue here is that a device of finite dimensions is attempting to measure a (theoretically) infinite Fourier spectrum and in doing so performs a bandpass filtering operation (see Fig. 6). For these experiments a constant acoustic vector of 1° s was used. A finite detector width will mean that we, in effect, measure the matrix-vector product of a blurred version of our inputs. Once again, the model tells us that if we filter out the $+1$ and -1 orders due to the horizontal LCD sampling frequency [those orders in Fig. 6(b) not preserved in Fig. 6(d)], the horizontal pixel structure does not affect our computation.

2.3 Processor Development

The time-invariant nonuniformities in our system were due to the laser (Gaussian beam profile), AOU (attenuation, anisotropy, inhomogeneous regions), and optics (dust, aberrations). All time-independent nonlinearities were successfully compensated for in our system since temporally they act as constants. Time-dependent nonuniformities (thermal

variations in the equipment, vibrations, airborne dust, electrical noise), optical interference, and oversampling (LCD, camera and frame grabber) conspire to reduce the linearity of the system. However, physical experiments with the system (presented later in the paper) show that these time-dependent nonuniformities do not appreciably affect the processor and, in the case of ANN training, are actually a desirable influence.

For certain experiments, the AO unit was not directly used in the computation. It lay in the optical path and contained an acoustic beam of constant amplitude (and constant frequency) representing a vector of 1's. This served to evaluate the other equipment and to verify the effective linearity of the system while still taking into account the physical properties of our calomel AO cell. To effect a multiplication, both area and intensity modulation was employed at the LCD plane. This was achieved by using the LCD panel to represent a lower resolution matrix of data values, where each data value was encoded in a square of pixels. With grid squares of side 10 pixels, area modulation effectively enabled 100 resolvable levels. The two images were combined in a computer memory and the resultant image was sent to the LCD panel. Under illumination, the signal immediately behind the LCD panel represented the product of the area and intensity modulating matrices.

Since negative numbers cannot be readily represented with a light intensity, an algorithm to allow real-valued perceptron weights in a -1 to $+1$ range to be represented in a 0 to $+1$ range was developed in advance. Our matrix-vector product processor must compute

$$y_i = \sum_j a_{ij}x_j, \tag{2}$$

where each input $x_j \in \mathbf{R}^+$ is a positive real-valued number, and each weight $a_{ij} \in \mathbf{R}$ and each output $y_i \in \mathbf{R}$ are (possibly negative) reals. Introducing a constant $K \in \mathbf{R}^+$

$$y_i = \sum_j (a_{ij} + K - K)x_j, \tag{3}$$

and expanding

$$y_i = \sum_j (a_{ij} + K)x_j - K \sum_j x_j, \tag{4}$$

enables us to define the following encoding scheme

$$y_i = \sum_j c_{ij}x_j - R_{\max} \sum_j x_j, \tag{5}$$

where R_{\max} is the largest positive value that can be represented in our system and $c_{ij} = (a_{ij} + R_{\max})$ is the encoded weight. If we ensure that $R_{\max} + R_{\min} = 0$, then since $c_{ij} \in \mathbf{R}^+$ our encoding scheme will transform the set of real-valued weights a_{ij} to the positive real values c_{ij} required by an intensity modulating SLM. By subtracting the constant $R_{\max} \sum_j x_j$ at the end of the calculation we return our matrix-vector product to the correct value in the -1 to $+1$ range.

3 Computational Theory and Representation

ANNs are an obvious application for exploiting the high connectivity possible with optical processors. The self-correcting nature of ANNs also means that they are somewhat tolerant to individual low precision or incorrect calculations throughout their learning cycle. In general, calculating the weight updates for one layer of an ANN involves a matrix-vector inner product, such as that shown in Eq. (2), where the signal to neuron y_i is the sum of the vector of inputs x_j multiplying y_i 's set of input synaptic weights a_{ij} . In this paper, we implement a single-layer perceptron²² with binary inputs, bipolar outputs, real-valued weights, delta learning rule and threshold transfer function. The learning tasks were of a linearly separable class of problems: several two-input logic functions. To represent these Boolean operations with a perceptron the classification output vector was transformed from the more common unipolar set $\{0,1\}$ to the bipolar set $\{-1,1\}$. We have already shown how a negative weight space may be represented by an intensity modulating SLM.

Detecting curves in an image is an important elementary task in machine vision. The HT has long been recognized²³⁻²⁵ as a robust technique for the detection of analytically defined shapes in a scene. The position of the highest peak in Hough space indicates the position (and possibly orientation, scale, etc.) of the most likely curve in image space matching our description. The HT can also be generalized²⁶ to reveal the presence of a shape that cannot be parametrically formulated (a fingerprint for example).

Due to the HT's wide use and applicability, an optical ANN implementation¹⁷ of the transform is highly desirable. Such an implementation would combine the massively parallel connectivity of neural network paradigms with the parallelism of optical systems. This should significantly lower the currently restrictive computational complexity of the transform on digital electronic architectures for large images.

3.1 Connectionist Model of the HT

Seth and Datta¹⁷ have proposed a neural network model for the HT that transforms image space into Hough space via a parameter space. Input images are represented as vectors and each neuron in image space u_i is connected to all neurons in parameter space through interconnection weights. Each weight w_{ij} and each neuron in the parameter space can have the value 0 or 1, and the value of each element in Hough space A_j is given by

$$A_j = \sum_{i=1}^m v_j^i = \sum_{i=1}^m w_{ij}^i u_i, \tag{6}$$

where v_j^i is the j th neuron of the i th parameter space.

The feed-forward architecture and binary representation of neurons are two significant features of this model which make it suitable for optical implementation. Seth and Datta have already presented an algorithm that can be used to obtain the values of the weight matrix for line detection. The following pseudo-code algorithm could be used for circle detection,




Fig. 7 Weight space matrices: (a) matrix to detect lines in a 4×4 pixel image with slopes from -1 to 1 (interval of 0.25) and intercepts from 0 to 3 (interval of 1) and (b) matrix to detect circles with various centers and radii in a 6×6 pixel image.

$$r_{\max} = (x_{\max}^2 + y_{\max}^2)^{1/2}$$

$$t = (r_{\max} - 1) / r_{\max}$$

For all image pixels (x_i, y_i)

For all possible (a, b, r)

$$e_{ij} = r_j^2 - (x_i - a_j)^2 - (y_i - b_j)^2$$

if $(-t < e_{ij} < t)$ then

$$w_{ij} = 1$$

otherwise $w_{ij} = 0$

where interconnection error e_{ij} is compared against threshold t , and where a and b define the center and r the radius of the circle. After updating the weights, the weight matrix contains the information required to map every possible point in image space to parameter space. Therefore, any shape in the image space with parameter values within the appropriate range can be detected using the same weight matrix. In effect, each row of the weight matrix represents one possible configuration of the shape in the image, as shown in Fig. 7.

Although these algorithms map closed form analytically described curves to a weight space, the weight space itself is fully generalized. Any arbitrary shape can be represented by simply building up a weight matrix representing all allowable configurations (varying position, orientation, scale, pose, etc.) of that shape.

4 Implementation and Experimental Results

An ANN with simple individual functions and a straightforward weight space update algorithm was chosen: a single-layer perceptron. It had been thought that any inaccuracies in calculating the weight space updates during the perceptron's training phase would be detrimental to learning. It has been shown recently, however, that synaptic noise (which we equate with the random analog inaccuracies introduced by time-dependent nonuniformities in our optical processor) in a multilayer perceptron during its training phase will enhance fault-tolerance, generalization ability and learning speeds.²⁷

For the following perceptron experiments, a constant vector was encoded in the acoustic beam. The complete set of binary input vectors was optically encoded on the LCD panel as a matrix using intensity modulation while the weight space was encoded using area modulation. By presenting the input vectors simultaneously as a matrix, each inner product could be computed in parallel and a complete learning epoch calculated instantaneously. The perceptron's nonlinear transfer function was effected through a software thresholding operation on the captured CCD signal. New weights were calculated using the average of the four weight adjustment vectors (one corresponding to each input vector). The new weight space was then scaled to the maximum LCD resolution before the SLM was updated by the PC. The delay between epochs is introduced by the update latency of the LCD (refreshing at 32 kHz vertically and 60 kHz horizontally), the matrix-vector computation itself (3 ns), the frame rate of the CCD sensor (effectively 50 Hz), and a small software routine. The software delay increases linearly (as opposed to quadratically) with an increase in the number of neurons.

Currently, the acoustic signal propagation is not synchronized with the detector. A sequence of frames (with adjustable phase) is captured while the signal propagates through the AO cell and the correct frame extracted manually. This procedure is satisfactory for a single pass arrangement, however, to automate the process we require strict synchronization between the AO cell, the detector, and possibly the laser.

4.1 AO Perceptron

A perceptron to learn selected two-input Boolean operations was successfully implemented on the architecture. The four input vectors were

$$\mathbf{X}_1 = \langle 0, 0 \rangle, \quad \mathbf{X}_2 = \langle 0, 1 \rangle, \quad \mathbf{X}_3 = \langle 1, 0 \rangle, \quad \mathbf{X}_4 = \langle 1, 1 \rangle.$$

The perceptron also had a bias input of 1 , and an associated weight. At each epoch, the four input vectors were sent through the system, generating four outputs. These outputs are plotted against epoch number in Fig. 8 for the AND and NOR operations. To aid visualization, lines are drawn between the discrete values in time. The outputs are shown here before application of the transfer function and therefore give a quantitative indication of the robustness of the system. In particular, the distance of each output from the threshold can be seen, as well as how symmetrically the set of outputs is distributed on either side of the threshold.

Initially the weight space was set to zero, which corresponds physically to random values determined by the time-variant nonlinearities in the system. It was found that normalization of the weight space (rescaling to the full intensity resolution of the LCD) between epochs resulted in improved convergence. In addition, the inherent noise in the optical system helped weights avoid settling into local minima. For example, while learning the NOR operation [Fig. 8(b)] the four outputs have settled at constant values representing a temporary solution. However, sufficient noise in the system at epoch 11 enables the outputs to climb out of this local minimum, producing a more desirable distribution either side of the zero threshold.




Fig. 8 Graphs of the four unthresholded outputs corresponding to the input vectors X_1 to X_4 (defined in the body of the paper) plotted against epoch number for each of two Boolean functions. The threshold level is shown as a dotted line in each graph. (a) When the outputs have stabilized (by epoch 17) the bipolar thresholded outputs for X_1 to X_3 have the value -1 and X_4 has the value $+1$, which is the correct response for a perceptron that has learned the AND operation. (b) A similar graph where input X_1 is the only vector that generates a $+1$ output showing that the perceptron has been successfully trained with the NOR operation.

The learning coefficient in each of these examples was $\eta=0.05$. Several different values for η were tested; however, the convergence time remained typically between 25 and 35 epochs for each of six logical functions tested. More importantly, the development of the weight space over time was similar for each η chosen (see Fig. 9). This insensitivity to η is thought to be due to the weight space normalization between epochs.

As mentioned earlier, although the AOU was not initially used in a processing capacity, the success of these experiments showed the effective linearity of our AO processor. This enabled us to confidently implement an HT on the system and utilize the increased vector update rate of our AO SLM. No processing capability is lost in going from a 2-D (LCD) to a 1-D (AO) input representation (unless, for example, one employs a lenslet array for interconnection¹⁰) since a 2-D representation requires the input vector to be replicated on each row.

4.2 AOHT

In the first HT experiment we sought to detect straight lines in a 4×4 pixel image. This repeats the original experiment of Seth and Datta.¹⁷ In our reconfigurable implementation, the weight matrix of Fig. 7(a) was used. The input image is shown in Figs. 10(a) and 10(b). The highest peak in row 28 of Hough space [Fig. 10(c)] indicates that a line has been detected in the input image with a slope of -1 and an intercept of 3. The second experiment verifies this result. The weight matrix and the image vector are simultaneously encoded on the LCD panel using area and intensity modulation respectively. The detected CCD image [Fig. 10(d)] once again reveals a maximum value in the 28th row. Hough spaces are characteristically noisy in contrast to the highly resolved peaks of correlation planes, and this is reflected in our noisy optical measurements.




Fig. 9 Adjustments to weight w_2 during NAND learning experiments with four different values for the learning coefficient η . Although the learning times for each η are not identical, we can see from the significant overlap in the plots that the development of the weights over time is somewhat invariant to the learning coefficient chosen.




Fig. 10 Line detection with the AO HT: (a) the 4×4 input image, (b) the input image represented as an electrical signal for modulation of the AOU, (c) the Hough space from the first experiment showing maximum optical energy in row 28, and (d) maximum optical energy in row 28 of the Hough space from the second experiment.

The third experiment involved detecting a circle in a square image of side 6 containing multiple objects (see Fig. 11). The circle detection matrix of Fig. 7(b) was used as the weight space, and once again the input matrix was encoded using intensity modulation on the LCD panel. The result of




Fig. 11 Circle detection with the AO HT: (a) a schematic of the 6×6 input image, (b) a 1-D representation of this input for AO modulation, and (c) result of the matrix-vector product detected by CCD camera showing a peak in the 129th row of Hough space.

the transformation to Hough space revealed a maximum value in the 129th row, which corresponds to the correct detection of a circle with radius 2 and center (3,3).

5 Conclusions and Future Work

A reconfigurable optical computing platform, based on a calomel AO matrix-vector multiplier, was designed and built. We outlined the major considerations that arise when building such a system and presented the successful results of training a bipolar-output perceptron with several (linearly separable) logical operations to analyze the processor's soft learning capabilities. It was possible to compensate for all of the system's time-invariant nonlinearities and we found that those remaining inherent nonlinearities introduced a level of noise and crosstalk that actually aided learning.

Subsequently, we chose to implement a connectionist model of the HT and verified its operation by detecting lines and circles. Our Hough space is fully generalizable and can be used to represent any arbitrary shape with varying position, orientation, scale, or pose. The advantages of our AO implementation over other techniques include a computational space-complexity independent of image size (up to the maximum allowable resolution of our LCD panel), a transformation throughput limited only by the aperture time of the AO cell, and a reconfigurable weight matrix easily controllable by a computer program on a PC.

Calomel's exceptional optical properties make it a promising material for AO signal processing. We have found, however, that these properties come at the cost of reduced quality in the final fabricated unit through a combination of crystal growth complexity and strict handling requirements.

Our short term goals are focused on expanding the ANN and HT implementations to utilize the full space-bandwidth product of our LCD panel. It will then be possible to accommodate the weight space of an ANN layer with up to 640 inputs and up to 480 neurons, while only linearly increasing the computation time due to the thresholding software routine. [The theoretical computing power of our optical setup (based on a 640×480 pixel LCD panel and an $85 \mu s$ AO unit update time) is calculated as 3.6×10^9 real-valued multiplications per second.] In addition, acquiring a custom built linear detector with an appropriate pixel aspect ratio and ability to synchronize with the AO unit will remove our current necessity to use an area-scan CCD camera.

Hard learning (solving problems whose solution spaces are not linearly separable) is not possible using a single layer perceptron but it is possible to implement more sophisticated networks using a multiple pass optical arrangement. These experiments have addressed both the constructional and operational difficulties of such an arrangement, currently under development, which is a tellurium dioxide-based optical neural network capable of hard learning.

Acknowledgments

This work was sponsored in part by research grants from Enterprise Ireland (HEIC-HE/95/471) and ESPRIT (ROPE-LTR21920). The authors would like to sincerely thank the reviewers of this paper for their helpful observations and suggestions.

References

1. C. Mead, *Analog VLSI and Neural Systems*, Addison-Wesley, Reading, MA (1989).
2. M. A. Arbib, Ed., *The Handbook of Brain Theory and Neural Networks*, MIT Press, Cambridge (1995).
3. H. J. Caulfield, J. Kinser, and S. K. Rogers, "Optical neural networks," *Proc. IEEE* **77**, 1573–1582 (1989).
4. Special issue on neural networks, *Appl. Opt.* **26**(35), (Dec. 1987).
5. Special issue on neural networks, *Appl. Opt.* **32**(8), (Mar. 1993).
6. T. Lu, S. Wu, X. Xu, and F. T. S. Yu, "Two-dimensional programmable optical neural network," *Appl. Opt.* **28**, 4908–4913 (1989).
7. S. Wu, T. Lu, X. Xu, and F. T. S. Yu, "An adaptive neural network using a high resolution video monitor," *Microwave Opt. Technol. Lett.* **2**, 252–257 (1989).
8. F. T. S. Yu, S. Wu, X. Yang, S. Jutamulia, and G. M. Storti, "Optical neural network using novel electron trapping materials," *Opt. Mem. Neural Netw.* **1**, 49–54 (1992).
9. R. C. Frye, E. A. Rietmann, and C. C. Wong, "Back-propagation learning and nonidealities in analog neural network hardware," *IEEE Trans. Neural Netw.* **2**, 110–117 (1991).
10. S. Gao, J. Yang, Z. Feng, and Y. Zhang, "Implementation of a large-scale optical neural network by use of a coaxial lenslet array for interconnection," *Appl. Opt.* **36**, 4779–4783 (1997).
11. M. Mori, Y. Yagai, T. Yatagai, and M. Watanabe, "Optical learning neural network with a pockels readout optical modulator," *Appl. Opt.* **37**, 2852–2857 (1998).
12. C. Barta, "Preparation of mercurous chloride monocrystals," *Krist. Tech.* **5**, 541 (1970).
13. M. Gottlieb, A. P. Goutzoulis, and N. B. Singh, "Fabrication and characterization of mercurous chloride acoustooptic devices," *Appl. Opt.* **26**, 4681–4687 (1987).
14. A. P. Goutzoulis and M. Gottlieb, "Characteristics and design of mercurous halide Bragg cells for optical signal processing," *Opt. Eng.* **27**(2), 157–163 (1988).
15. Space Technology Ireland Ltd., "Single crystal calomel—physical properties," Data sheet CAL-DAT-01, Maynooth, Ireland (1993).
16. S. M. P. McKenna-Lawlor, D. J. Gleeson, C. J. O'Driscoll, M. Klíma and C. Barta, "Fast optical signal processing aboard spacecraft based on calomel," in *Proc. 44th Congress of the Int. Astronautical Federation*, pp. 1–6, Graz, Austria (Oct. 1993).
17. M. Seth and A. K. Datta, "Optical implementation of a connectionist model of Hough transform," *Opt. Eng.* **35**(6), 1779–1784 (1996).
18. R. D. Guenther, *Modern Optics*, Wiley, New York (1990).
19. A. P. Goutzoulis and D. R. Pape, Eds., *Design and Fabrication of Acousto-Optic Devices*, Marcel Dekker, New York (1994).
20. M. Klíma, J. Rott, P. Dvořák, D. Gleeson, S. McKenna-Lawlor and J. Keating, "Model of 2D optical correlator for fingerprint identification," *IEEE Aerosp. Electron. Syst.* **12**(7), 3–9 (1997).
21. A. VanderLugt, *Optical Signal Processing*, Wiley, New York (1992).
22. S. Haykin, *Neural Networks: A Comprehensive Foundation*, Prentice Hall, Englewood Cliffs, NJ (1994).
23. P. V. C. Hough, "Methods and measures for recognising complex patterns," U.S. Patent No. 3069654 (Dec. 1962).
24. V. S. Nalwa, *A Guided Tour of Computer Vision*, Addison Wesley, Reading, MA (1993).
25. E. Trucco and A. Verri, *Introductory Techniques for 3-D Computer Vision*, Prentice Hall, Englewood Cliffs, NJ (1998).
26. D. H. Ballard, "Generalizing the Hough transform to detect arbitrary shapes," *Pattern Recog.* **13**(2), 111–122 (1981).
27. A. F. Murray and P. J. Edwards, "Analogue synaptic noise—a hardware nuisance, or an aid to learning?" in *MicroNeuro'93: Proc. 3rd Int. Conf. on Microelectronics for Neural Networks*, pp. 121–129, UnivEd Technologies, Ltd., Edinburgh, Scotland (Apr. 1993).



Thomas Naughton received his BSc in computer science and experimental physics in 1995 from the National University of Ireland, Maynooth (previously St. Patricks College), where he is a PhD student and teaches computer science. His research interests include optical information processing and computational complexity theory, with a current emphases on acoustic-optic systems and nonlinear joint transform correlation. He was previously with Space Technology Ireland Ltd. and at the Czech Technical University in Prague. Naughton holds one patent and is a member of SPIE, OSA, IEEE and ACM.



Zohreh Javadpour received her BSc degree in physics from Tarbiat Moalem University in Tehran, Iran, in 1987, and her honors degree, master of computer science, in 1996 from the National University of Ireland (NUI), Maynooth, where she is presently a PhD student. Her research areas are optical neural networks, pattern recognition techniques and applications.



John Keating is a lecturer in the National University of Ireland, Maynooth. He received his BSc in mathematics and experimental physics and his PhD in experimental physics in 1986 and 1990, respectively. His current research interests include artificial neural networks and optical computing. He is currently a member of the Institute of Physics.



Miloš Klíma received his MSc in 1974 and his PhD in 1978 from the Czech Technical University in Prague. Currently he is an associate professor and a group head with the Department of Radioelectronics, Faculty of Electrical Engineering, the Czech Technical University, Prague. He has been involved in R&D activity in the field of laser beam control for imaging purposes, imaging TV systems under the space program Interkosmos, smart image sensors and more recently optical methods for information processing. His educational activities are related to photonics, television and image processing.

Jirí Rott: Biography and photograph not available.