

Design and Demonstration of an Opto-Electronic Neural Network using Fixed Planar Holographic Interconnects

Paul E. Keller and Arthur F. Gmitro
Optical Sciences Center
University of Arizona
Tucson, Arizona 85721

keller@martha.radiology.arizona.edu and gmitro@zen.radiology.arizona.edu

Abstract

Implementation of an opto-electronic Hopfield style associative memory neural network is discussed with emphasis on the construction of an experimental system employing binary amplitude holograms.

1. Introduction

A key element of most neural network systems is the massive number of weighted interconnections used to tie relatively simple processing nodes together in a useful architecture. The inherent parallelism and interconnection capability of optics make it a likely candidate for the implementation of the neural network interconnection process. While there are several optical technologies worth exploring, we are looking at the capabilities and limitations of using fixed planar holographic interconnects in a neural network system and have implemented an initial test system using planar holograms and opto-electronic nodes.

2. System

All neural network systems consist of nodes (simple non-linear elements crudely imitating biological neurons) and weighted interconnections (synapses) between nodes. The basic system we have looked at employs optical interconnects and electronic nodes in a feedback architecture. A prototype is shown in Figure 1.




Figure 1: Prototype opto-electronic neural network system.

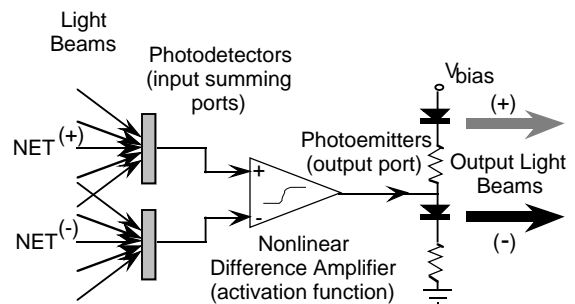


Figure 2: Single node in the system.

Each node is composed of an input summing port, non-linear transfer device, and an output port. In an opto-electronic system, a differential pair of detectors is operated as an input to the node; signals with positive (excitatory) weights arrive at one detector, and signals with negative (inhibitory) weights arrive at another detector. These detectors sum up the intensity of each optical signal arriving at the node. A threshold operation is electronically applied to the detected signal to produce an output signal. The output signal of the node modulates an optical source. Figure 2 illustrates an idealized node.

An individual node drives an optical beam that illuminates a single subhologram. Each subhologram stores the connection weights between that node and all other nodes. A subhologram is designed as a Fourier transform hologram and used in a coherent optical system so that the diffracted connection pattern is independent of subhologram position.

3. Design

The Hopfield¹ auto-associative memory model was chosen as a means to test the interconnect capability of planar holographic optical interconnects in the experimental opto-electronic neural network. This neural network

tries to associate each pattern presented to it with a pattern that it was trained on during an initial batch training process. Using the Hopfield outer product formulation, a training set of patterns was used to construct the fixed interconnection weights. The Hopfield model is a globally interconnected neural network; all nodes are connected to all other nodes with a deterministic strength or weight. These weights were then encoded in an array of binary amplitude subholograms.

Much effort went into the construction of the holograms used in the interconnect process. After examining a variety of computer generated hologram (CGH) techniques for accuracy of reconstructed interconnect weights, computation time, required space bandwidth product, and diffraction efficiency, two techniques, error diffusion and random search, were found to satisfy many of these criteria.

Both techniques were used to produce binary amplitude holograms. The general design techniques are as follows. The weights connecting a single node output to all node inputs are represented as an intensity pattern in the detector plane. The optical amplitude at the detector plane is given by the square root of the intensity. To reduce the dynamic range required to encode a hologram, a random phase function is added. Since amplitude holograms must produce a hermitian diffraction pattern, the mapped weights are shifted off the optic axis and a hermitian conjugate is added. To compensate for the sinc function roll-off in the connection pattern due to finite sized hologram pixels, the connection weights are multiplied by an inverse sinc function weighting. This predetermined diffraction pattern is then inverse Fourier transformed, and the transformed data values are normalized using the extreme amplitude values for the entire set of subholograms. Lexicographically scanning these sampled values, each sample is binarized. Since the data is continuous valued, an error is produced by the binarization. This error is propagated to the adjacent unbinarized pixels of the hologram; this is the error diffusion process.² The net effect of the error diffusion process is to reduce the total quantization error across the entire hologram. What remains is a high frequency binarization error that manifests itself as diffracted light far off the optic axis in the detector plane. The location of this diffracted light is controlled by the method used to distribute the binarization error in the hologram. To improve interconnection weight accuracy and to confine the diffracted spots of light to the center of each detector cell, each hologram is replicated 4-times vertically and horizontally (16 replicas). The error diffusion algorithm has shown the best performance for a non-iterative CGH design process.

Random search is an iterative process used to improve the connection accuracy of the error diffusion holograms. Starting with an error diffusion hologram, this method determines whether a perturbation (flipping a pixel in the hologram from opaque to transparent or vice versa) improves the accuracy. A perturbation is kept only if the accuracy is improved. This process is repeated until convergence. The main disadvantage of the random search process is the massive computation required. This process is related to the simulated annealing process except that no annealing takes place. The simulated annealing process allows accuracy degrading perturbations of the hologram to be kept with a probability modeled by the Maxwell-Boltzmann distribution.³ Simulated annealing, in theory, is able to find the globally optimum solution; in practice, limited computation requires compromises that may or may not produce good results. We have found that the random search algorithm produces holograms with almost the same performance as simulated annealing but requiring far less computation time.

For a large scale problem, electron beam fabrication would be required to produce the array of subholograms. For the small scale problem that was implemented, a photolithography process was used; the hologram mask was printed onto a sheet of film using a laser film writer and photographically reduced onto a holographic plate.

4. Experiment

The experimental opto-electronic neural network is illustrated in Figure 3. An initial pattern of 8 by 8 pixels is fed into the system by a computer; this pattern represents the initial state of the neural network. The pattern is written onto a Hughes Liquid Crystal Light Valve SLM using a high-intensity projection television. This binary pattern is polarization encoded onto the coherent optical laser beam by the SLM. The polarization beam splitting cube reflects only the vertical component of this polarized signal so that a binary amplitude pattern illuminates the hologram array. Each pixel of the pattern illuminates an individual subhologram. There are 64 nodes with 4096 bipolar interconnections in the experimental system.

The Fourier transform (Fraunhofer diffraction pattern) of the hologram array is produced at the back focal plane of the lens. To reduce scatter, the low frequency information of the diffraction pattern is filtered out. A relay lens is used to image the filtered Fourier plane onto a video camera. The light beams (diffraction from the hologram plane) arriving at the detector plane constitute the input to the node plane. In a practical opto-electronic neural network, each electronic node would take the difference between the signal on its positive-weight detector and its negative-weight detector, threshold the result, and drive an optical source such a laser diode to be either on or off. For our experimental test system, a video camera is used to detect the optical input signals. The video signal is fed into the computer where it is digitized by a video frame buffer. The computer splits up the video frame into a grid and sums up the intensity in each cell to simulate a detector array. The difference and thresholding operations are performed digitally and the output stored in a video frame buffer, where the video output represents

the next iteration of the network. This forms the new network state, which illuminates the hologram plane, and the process continues until the network converges to a stable state.

As this experimental system was described, a node can take on two values; a value of 0 is represented by a dark pixel, and a value of 1 is represented by a light pixel. The performance of a Hopfield style neural network is significantly improved by using bipolar node values instead of unipolar node values. As an experimental test of bipolar nodes, a two step process was used. During the first step, a pattern was projected onto the SLM, and the detected pattern on the video camera was stored. During the second step, the inverse of the pattern was projected onto the SLM, and its detected pattern on the video camera was subtracted from the first detected pattern. Polarization encoding is another method for constructing bipolar nodes. A node value of +1 is encoded as horizontally polarized light, and a node value of -1 is encoded as vertically polarized light. With bipolar weights and bipolar state values, four detectors and two polarizers are used in the input summing port of the node.

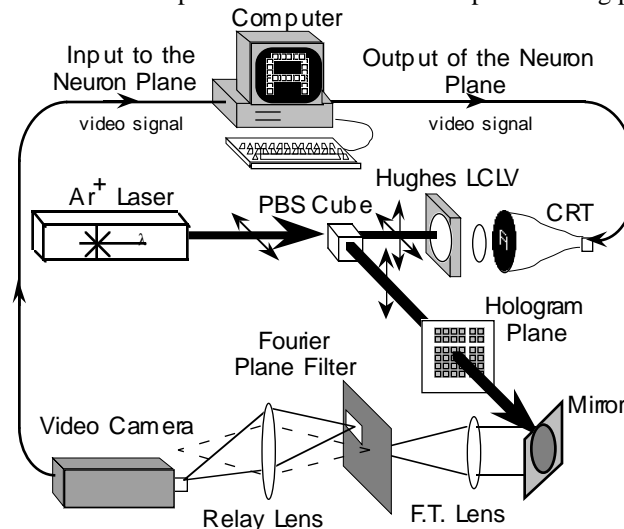


Figure 3: Experimental opto-electronic neural network used to test and evaluate the performance of planar holographic interconnects.

5. Results

The best performance of the associative memory neural network would come from a network storing randomly generated patterns, but since patterns of distinct structure (vertical lines, horizontal lines, diagonal lines) are generally encountered in vision and pattern recognition tasks, it was decided to use a set of ordinary typewriter characters (letters, numbers, symbols) to construct the test network. Using the Hopfield outer product formulation, a training set of three patterns, ABX , was used to determine the interconnection weights.

A prime feature of auto-associative memory neural networks is the convergence of the network to the ideal stored pattern when the input pattern is corrupted. By randomly flipping the pixels of the training set, a test set of corrupted patterns was generated. These patterns were presented to the experimental opto-electronic neural network, a computer simulation of the opto-electronic neural network, and a computer simulation of the ideal neural network. From the simulation, it was found that the auto-associative neural network constructed with random search holograms performed almost identically to the same neural network with ideal interconnect weights. The experiment, while not performing quite as well as the simulation, did come close for both unipolar and bipolar state values. The results with error diffusion holograms were not as good as with the random search holograms but show that error diffusion based holographic interconnects are a good trade-off between system performance and CGH computation time for bipolar state values. Figure 4 illustrates the performance of the experimental opto-electronic neural network with a test set composed of corrupted versions of the letter B . This figure is a graph of the probability that the network converges to the original letter B as a function of the number of corrupted pixels in the input pattern. The total number of pixels in the input is 64. From this graph, it is apparent that when the number of incorrect pixels in the input pattern becomes too large, the network does not converge to the ideal pattern. Similar responses were found with the other stored patterns.

The small differences between the experimental results and the simulation results were caused by aberrations in the Fourier transform lens and relay lens, non-uniformity of the video camera, high frequency roll off in the holograms due to loss of resolution during the hologram fabrication process, and RF interference in the electronics produced by the argon ion laser's plasma discharge tube.

The experimental opto-electronic neural network system along with its computer simulation shows that a planar hologram can be used to implement the interconnect weights of a neural network. The results we have found with the experiment agree well our analytic calculations of neural network performance.⁴

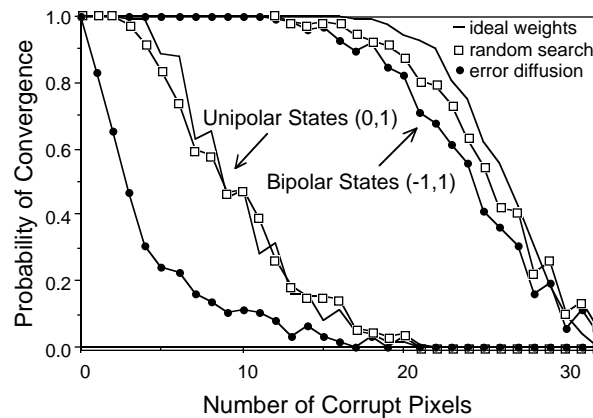


Figure 4: Performance of the ideal associative memory and the opto-electronic implementations. The graph plots the probability of convergence of the network to the correct state versus the number of corrupted pixels.

6. Conclusions

We have demonstrated that a system employing planar holographic optical interconnects can be used to implement a neural network architecture and that the performance of an optically implemented Hopfield style network comes close to that of an arbitrary system employing ideal interconnect weights.

This research is supported by the Optical Circuitry Cooperative at the University of Arizona.

7. References

1. J.J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," Proc. Natl. Acad. Sci. USA **79**, 2554-2558 (1982).
2. S. Weissbach, F. Wyrowski, and O. Bryngdahl, "Quantization Noise in Pulse Density Modulated Holograms," Opt. Comm. **67**, 167-171 (1988).
3. M. Kim, M. Feldman, and C. Guest, "Optimum encoding of binary phase-only filters with a simulated annealing algorithm," Optics Letters **14**, 545-547 (1989).
4. A. Gmitro, P. Keller, and G. Gindi, "Statistical Performance of Outer-Product Associative Memory Models," Appl. Opt. **28**, 1940-1948 (1989).