Optoelectronic Implementations of Neural Networks

Demetri Psaltis
Alan A. Yamamura
Ken Hsu

Steven Lin
Xiang-guang Gu
David Brady

T HOUGH WE ARE NOW SURROUNDED BY COM-puters performing a seemingly endless array of tasks, there is an apparent dichotomy between problems best solved by machines and those best solved by humans. Machines far exceed human capabilities at problems that can be solved using simple logic operations on a large scale and at great speed, but biological systems remain far superior at problems involving a high degree of unpredictability, such as speech recognition, machine vision, and motor control. By incorporating some of the properties of biological systems in artificial neural networks, we hope to create machines that might address this second class of problems. In the simplest case, a neural network consists of a large number of relatively simple processing elements, or neurons, and a set of connections between these neurons. The connections define the functionality of the network, and their strengths are determined through adaptation or "learning." While the basic concept of an artificial neural network is fairly old, a resurgence in neural network research has been spurred by an increased understanding of biological systems, the development of a variety of neural architectures and training methods, and technological advancements in the implementation of neural computers.

The primary advantage of optics lies in the optical systems ability to provide the massive interconnections between processors required in most neural network models.

Optics has emerged as a promising technology for neural network implementation. Most of the proposed optical designs [1] can be cast into a canonical architecture with two-dimensional planes of processors connected optically via the third dimension (Figure 1). The primary advantage of optics lies in an optical system's ability to provide the massive interconnections between processors required in most neural network models. Massive connectivity is comparatively difficult in electronic systems because electrical signals must travel on physical wires that consume area and require careful design to minimize interference and crosstalk. Optical signals, on the other hand, can propagate through free space and pass through one another without interacting. Imaging systems, like those in the eye, easily provide millions of one-to-one connections between points on opposite sides of a lens. The most promising method of optical interconnections involves the use of holography. Though holograms are best known for their ability to generate three-dimensional images, they are more generally capable of redirecting light in a programmable fashion. Holograms can be used to connect each neuron with others in the same, or adjacent, processing planes. In contrast with implementations, in which the specifications of the connection patterns must be stored separately from the connections themselves, holographic media can simultaneously provide both the massive physical interconnectivity and the large memory required to specify the connections. This duality is particularly useful in adaptive networks.

The ability of neural networks to learn via adaptation of the connection pattern is one of the greatest strengths of the neural approach. Fixed holograms can be used in applications where learning is accomplished prior to actual use. Often, however, the network must continuously adapt itself to changing conditions for better overall performance. Dynamic holographic media, such as photorefractive crystals, allow continuous modification of interconnection patterns. In photorefractives, the interference between desired input and output light patterns move charges in the crystal, creating an electric field distribu-


Fig. 1. Optical neural computer architecture.

Volume Fourier Interconnections

Fig. 2. Optical heteroassociative memory.

tion which modulates the index of refraction. This varying index acts as a holographic grating that reproduces a desired output in the presence of its associated input. The portion of light emanating from one neuron, which is redirected as light converging to another, is determined by parameters in the writing process. By controlling these parameters, the efficiency of the coupling between each pair of neurons may be specified.

The ability of neural networks to learn via adaptation of the connection pattern is one of the greatest strengths of the neural approach.

The simplest form of holographic associative memory is realized by recording the interference between associated image pairs [2]. This is equivalent to recording the hologram of an image using its associated pattern as the reference beam. A well-known property of holography is that when a hologram is illuminated with the reference beam used during recording, the reconstructed wave is a reproduction of the object that was stored in the hologram. In this manner, we can say that the reference and object are associated with each other. An object is reconstructed accurately by its reference beam only if the reference is a simple, uniform wave. The patterns are formed on appropriately selected sampling grids of fractal dimension 3/2 to obtain a faithful reconstruction of the object beam when the reference is a complex image [3]. A schematic diagram of this simple form of holographic associative memory is shown in Figure 2. The interference patterns between two images placed at the input and training planes are holographically stored in the photorefractive crystal. Presentation of either of the two images at the input plane reconstructs its associated image as the output.

An example of a set of associations produced by this system is shown in Figure 3. In this experiment, the Arabic numerals 1, 2, and 3 were associated with their Chinese counterparts. Up to 111 such associations have been recorded in our laboratory onto a single crystal. The system of Figure 2 can also be used to implement error-driven learning algorithms. The simplest example for training a single layer is the perceptron algorithm [4] [5]. This algorithm uses a single-output neuron to dichotomize a set of input patterns. A pattern is presented at the input, and the light diffracted by the hologram is compared to the desired output. If the output for a given pattern is too low (high), the connections taking that image to the output are increased (decreased). A learning curve for an experimental implementation of this algorithm is shown in Figure 4. The diagram shows the output of the system (vertical axis) as a function of the number of learning cycles. Ten patterns of 100 pixels each are separated into two random classes. The output for incorrectly classified


Fig. 3. Experimental demonstration of the holographic associative memory.

patterns is shown below the centerline. After 188 presentations, all patterns are correctly classified, as is shown by the response above the centerline for presentations 188–198. Since no corrections are needed, learning stops at this point. We have correctly classified up to 15 images using the same process.

Processing elements acting as neurons comprise the final component of an artificial neural network. Since the holographic connections are essentially linear, we require nonlinear processing elements to perform complex computations. Though optical nonlinearities are more difficult to find than electronic ones, they do exist. For example, optical prop-


Presentation Number

Fig. 4. Experimental demonstration of the optical implementation of the perceptron learning algorithm.

(a)   (b)

Fig. 5. 10 × 10 array of optoelectronic neurons monolithically fabricated in GaAs.

erties like the index of refraction or absorption can be modified by the intensity of light in some nonlinear optical media, allowing the design of nonlinear processing elements, such as optical switches.

A more promising approach in the near term, is the use of hybrid optoelectronic processing elements. Gallium Arsenide (GaAs) is an excellent material for this purpose, since it can be used to fabricate both fast electronic circuits and optical sources and detectors. Figure 5 shows a GaAs device with a 10 × 10 array of neurons [6]. Each neuron has a phototransistor (to detect the optical input) and a second bipolar transistor in a Darlington pair that provides sufficient current gain to drive a Light Emitting Diode (LED) for communication with other neurons. The neuron provides a nonlinear thresholding characteristic through saturation of the transistor pair. The threshold level of the neuron is controlled by an external current source connected to the base of the phototransistor. We can implement more complex nonlinear operations at each neuron using additional detectors and electronic circuitry. The advantage of this combined optoelectronic approach is that it exploits the strengths of both optics and electronics with optics providing the connections and electronics providing the nonlinear processing.

It is, however, possible to implement a complete hybrid neural computer using available technology developed for conventional computing [7]. Figure 6 shows a neural computer architecture in which an analog integrated circuit implements a single layer of a multilayer feed-forward neural network. On the chip, electronic neurons from one layer connect to the next through an electronic crossbar circuit. Connections between neurons contain synapses that regulate the amount of current that flows between neuron pairs. Every neuron then receives a current from each of its inputs equal to the activity of the input times a weight, and then applies a nonlinear threshold function to the sum of its total synaptic excitation current before storing the result. Photodiodes are placed in each synapse location to simultaneously load all the connection weights optically from the top of the chip. The network on the chip is bidirectional, so that after data passes through one layer of the network, the disk spins to load the next set of weights before data passes through the crossbar.

In this manner, a feed-forward network can be implemented with as many layers as can be stored along a track on the disk. Optical storage of the weights provides a number of advantages. First, optical disks provide high-density, high-capacity storage, which is necessary for the potentially large number of weights needed for different neural network "programs." This is particularly true in image recognition applications. A typical 5-in. optical disk can store almost 10^{10} bits at micron spacing. Second, electronic implementations face a bottleneck in loading the weights because as the number of neurons, N, in each layer increases, the number of weights per layer increases as N^2 (for a feed-forward network that is fully connected between layers); at best, the number of available pads on a chip increases as $O(N)$. The optical storage format allows the weights to be quickly read out in parallel with a simple imaging system [7]. For example, with 10^4 synapses on a chip and cycle time limited by disk access time to 10 μs, the optoelectronic implementation performs 10^9 operations/s.

Though purely electronic versions could match this speed, the circuitry to sustain required MHz data transfer rates would
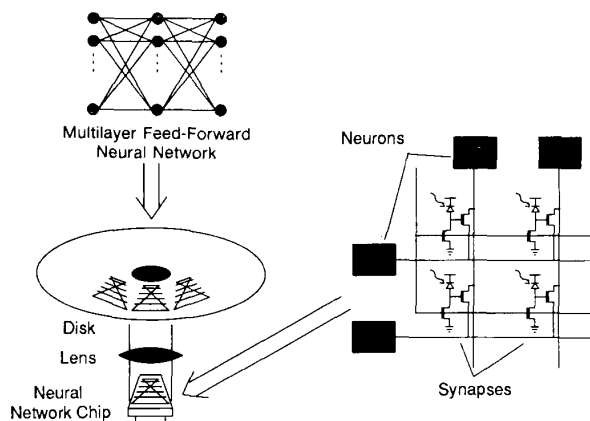

Multilayer Feed-Forward
Neural Network

Neurons

Disk

Lens

Neural
Network Chip

Synapses

Fig. 6. Hybrid optical and electronic multilayer feed-forward neural network.

Fig. 7. Optical autoassociative memory.

Partial Input

Rotation

Translation

Fig. 8. Input and output images from optical autoassociative memory.

limit the number of neurons on the chip and provide direct access to less than a hundredth of the memory directly accessible with the optical disk. The weights can be determined in a previous training phase using learning algorithms such as error back-propagation [8]. Therefore, mature optical disk and analog integrated circuit technologies developed for conventional computers can be used to implement a hybrid optoelectronic neural computer.

Next we describe an experimentally demonstrated network in which optics plays a larger role (see Figure 7) [9] [10]. The heart of the system is the Hughes Liquid Crystal Light Valve (LCLV) with 400 × 400 resolvable elements serving as a plane of 160,000 neurons. The reflectivity of a liquid crystal at each point of the output side of the LCLV is determined by the intensity of light at the same point on the input side. The complete network acts as an autoassociative memory—recalling an image stored in memory from a partial or corrupted input. The input is imaged onto the photoconductive input side of the LCLV. Light from an argon laser strikes the reflective output side of the LCLV (the returning beam contains a thresholded version of the input). Part of the image is sent to a CCD camera and viewed on a monitor as the output. The rest of the image enters an optical correlator. The correlator uses a lens to take the Fourier transform of the image and multiply it with the transforms of the library images stored in a hologram, while a second lens retransforms the products. The hologram is recorded such that the correlations of the input with each library image are spatially separated. A pinhole array samples the correlation peaks, and the light from each peak reconstructs its associated image through a second hologram. Therefore, the output of the second hologram is the linear combination of each stored image weighted by its correlation with the input image. This combined image is fed back to the input side of the LCLV, where each stored image positively reinforces its component in the input. As the data traverses the loop many times, the image which best matches the input will dominate. Figure 8 demonstrates the properties of the loop where noise-free library images can be recalled with partial inputs, even with small differences in rotation, translation, or scale. The sequence of photos show the system's operation. In each of the three step sequences, the left photo shows the external input; the center photo shows the system's response to the input and the external input superimposed; and the photo on the right shows the stored image locked onto by the loop after the input was removed.

Neural networks have made a comeback because of their potential at tasks that have proven difficult for conventional computational methods and machines. Optics has emerged as a leading technology for the implementation of neural computers because its capabilities match well with the requirements of most neural network models. The weakness of optics in providing complex nonlinear processing is mitigated by the fact that many neural network models require only simple processing elements. The strength of neural networks lies in distributed processing and massive interconnections which optics easily provides. It is possible today, to implement effective hybrid optoelectronic neural computers using optical disk and analog-integrated technology developed for more conventional computers. We have experimentally demonstrated neural network implementations using almost all optical components, but perhaps the most effective solution will involve using two-dimensional, optoelectronic processing planes with optical sources and detectors for communications and electronics for complex processing.

Neural networks have made a comeback because of their potential at tasks that have proven difficult for conventional computational methods and machines.

Acknowledgments

References

[1] *Applied Optics*, Special Issue on Neural Networks, vol. 26, no. 23, Dec. 1, 1987.