Communicated by Nabil H. Farhat

An Optoelectronic Architecture for Multilayer Learning in a Single Photorefractive Crystal

Carsten Peterson*
Stephen Redfield
James D. Keeler
Eric Hartman
Microelectronics and Computer Technology Corporation,
3500 West Balcones Center Drive, Austin, TX 78759-6509 USA

We propose a simple architecture for implementing supervised neural network models optically with photorefractive technology. The architecture is very versatile: a wide range of supervised learning algorithms can be implemented including mean-field-theory, backpropagation, and Kanerva-style networks. Our architecture is based on a single crystal with spatial multiplexing rather than the more commonly used angular multiplexing. It handles hidden units and places no restrictions on connectivity. Associated with spatial multiplexing are certain physical phenomena, rescattering and beam depletion, which tend to degrade the matrix multiplications. Detailed simulations including beam absorption and grating decay show that the supervised learning algorithms (slightly modified) compensate for these degradations.

Most network models are based on "neurons," V_i, interconnected with synaptic strengths T_{ij}, and a local updating rule:

$$V_i = g(\sum_j T_{ij} V_j) \tag{1.1}$$

where $g(x)$ is a nonlinear gain function.

The majority of neural network investigations are performed with software simulations, either with serial or parallel architectures. Many applications would benefit tremendously from custom-made hardware that would facilitate real-time operation. Neural network algorithms require a large number of connections (the matrix elements T_{ij} of equation 1.1). Optics offers the advantage of making these connections with light beams.

Several architectural proposals for optical implementations now exist (Psaltis and Farhat 1985; Soffer *et al.* 1986; Psaltis *et al.* 1982). Most

*Present Address: Department of Theoretical Physics, University of Lund, Solvegatan 14A, S-22362 Lund, Sweden.

deal with Hebbian learning (no hidden units) using either spatial light modulators (SLMs) or photorefractive crystals. The latter technology, in which the T_{ij}-elements are represented by holographic gratings, is particularly well suited for neural network implementations. The gratings decay naturally, and this can be exploited as a beneficial adaptive process. In Psaltis *et al.* (1982), a multilayer architecture of several photorefractive crystals was designed to implement the backpropagation algorithm with hidden units.

In this letter, we present a *single crystal architecture* that is versatile enough to host a class of supervised learning algorithms, all of which handle hidden units. In contrast to other approaches we use *spatial multiplexing* (Anderson 1987; Peterson and Redfield 1988) rather than angular multiplexing of the interconnects. This spatial multiplexing implies rescattering and beam depletion effects for large grating strengths and large numbers of neurons. We demonstrate, on the simulation level, how the supervised learning models we consider implicitly take these effects into account by making appropriate adjustments.

Photorefractive materials can be used as dynamic storage media for holographic recordings. A recording takes place as follows: With the object (1) and reference (2) beam amplitudes defined as

$$\vec{E}_1 = \vec{A}_1 e^{-i\vec{k}_1\cdot\vec{r}} \tag{1.2}$$

$$\vec{E}_2 = \vec{A}_2 e^{-i\vec{k}_2\cdot\vec{r}}, \tag{1.3}$$

the intensity pattern of the two-wave mixing takes the form[2]

$$I = |\vec{E}_1 + \vec{E}_2|^2 = I_1 + I_2 + 2\sqrt{I_1 I_2}\cos[(\vec{k}_2 - \vec{k}_1)\cdot\vec{r}] \tag{1.4}$$

where intensities $I = |\vec{E}|^2$ have been introduced. The refractive index is locally changed in the photorefractive material proportional to this periodic intensity pattern.

The so-called grating efficiency η is to a good approximation proportional to the incoming beam intensities

$$\eta^{1/2} \propto \sqrt{I_1 I_2} = A_1 A_2. \tag{1.5}$$

Consider a crystal where grating strengths η_{ij} have been created with the interference of equation 1.4. Let a set of light beams $\vec{E}_i = \vec{A}_i e^{-i\vec{k}_i\cdot\vec{r}}$ impinge on these gratings. The outgoing electric fields can be written as

$$\vec{E}_j = \vec{A}_j e^{-i\vec{k}_j\cdot\vec{r}} = \sum_i \eta_{ij}^{1/2} e^{i\vec{K}_{ij}\cdot\vec{r}}\vec{A}_i e^{-i\vec{k}_i\cdot\vec{r}} = \sum_i \eta_{ij}^{1/2}\vec{A}_i e^{-i\vec{k}_j\cdot\vec{r}} \tag{1.6}$$

[2]Neglecting a constant phase shift due to the relative phases of the beams.

where $\vec{K}_{ij} = \vec{k}_i - \vec{k}_j$. That is, a matrix multiplication of amplitudes

$$A_i = \sum_j \eta_{ij}^{1/2} A_j \tag{1.7}$$

is performed by the photorefractive medium. Thus, identifying the amplitudes A_i with the neuronic values V_i, and $\eta_{ij}^{1/2}$ with the connection strengths T_{ij}, the matrix multiplication of equation 1.1 can be implemented. Correspondingly, equation 1.5 implements a Hebb-like learning rule.

The reconstruction, or readout, process is partially destructive. The efficiency decays exponentially with the readout duration, for a given read energy density. In the past this grating decay has been a problem because the use of the neural network would gradually fade the recordings. A technique has recently been discovered for controlling this destruction rate by choosing appropriate applied fields and polarizations of the object and reference beams[3] (Redfield and Hesselink 1988).

Equation 1.7 can be implemented either with angular (Psaltis *et al.* 1982) or spatial multiplexing (Anderson 1987; Peterson and Redfield 1988) techniques. In Redfield and Hesselink (1988) it was observed that at most 10–20 gratings can be stored at the same localized region with reasonable recall properties using angular multiplexing. For this reason we have chosen the spatial multiplexing approach, which corresponds to direct imaging (see Fig. 4).

With direct imaging, the intensities from the incoming plane of pixels become depleted and rescattered when passing through the crystal, causing the effective connection strengths to differ from the actual matrix elements T_{ij} (see Fig. 1). For relatively small systems and grating ($\eta^{1/2}$) sizes these effects are negligible; in this case the identification of T_{ij} with $\eta_{ij}^{1/2}$ is approximately valid. However, these effects are likely to pose a problem for realistic applications with large numbers of neurons.

We estimate the rescattering and beam depletion effects on $I_i^{out} = A_i^2$ (equation 1.7) by explicit simulation of the reflection [coefficient η_{ij}] and transmission [coefficient $(1-\eta_{ij})$] of the intensity arriving at each grating.[4] In Figure 1 we show the emergent light given random $\eta_{ij}^{1/2}$ values on the interval [0,0.1]. The data represents an average of different random input patterns of intensity strengths in the interval [0,1]. The major effects are in the "end" of the crystal. Clearly, if the $\eta_{ij}^{1/2}$ were set according to a pure (unsupervised) Hebbian rule, a large network would produce incorrect answers due to the depletion and rescattering effects. We demonstrate below how to overcome this problem with *supervised learning* algorithms together with a *temperature gradient* procedure for the output amplifiers.

In our architecture, the input and output neurons are planes of n^2 pixels and the connection matrix is a volume that can map at most n^3

[3]This technique yields a recording half-life of O(10hrs) for continuous readout.

[4]It is sufficient for our purposes to investigate this effect on a slice of the volume.

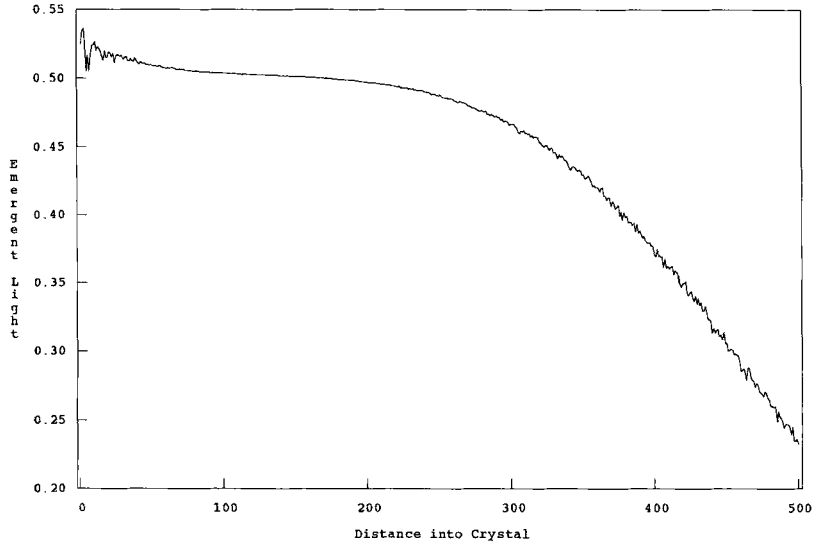

Figure 1: Emergent light (I_i^{out}) as function of distance (number of grating sites) into the crystal (i) for random $\eta_{ij}^{1/2}$ values and input intensities (see text).

connections. If we want full connectivity we are thus short one dimension (Psaltis *et al.* 1982). The volume can serve only $n^{3/2}$ neurons with full connectivity, so we need an $n^{3/2} \rightarrow n^{3/2}$ mapping between the SLMs. There are many ways to accomplish this. We have chosen to use multiple pixels per neuron. The input plane is organized as follows: each of the n rows contains $\sqrt{n}$ neurons replicated $\sqrt{n}$ times. In the output plane, each row contains $\sqrt{n}$ replicas of the sequence $i, i+1, \ldots, \sqrt{n}$, where $i = 1$ for the first row, $i = 1 + \sqrt{n}$ for the second row, $i = 1 + 2\sqrt{n}$ for the third row, etc. By deliberately omitting selected elements, architectures with limited (e.g., layered) connectivity can be obtained.

We begin by describing how to implement the mean-field-theory (MFT) learning algorithm (Peterson and Hartman 1988). We then deal with feedforward algorithms. The MFT algorithm proceeds in two phases followed by a weight update:

1. *Clamped phase.* The visible units are clamped to the patterns to be learned. The remaining units then settle (after repeated updating) according to

$$V_i = g(\sum_j T_{ij} V_j / T) \tag{1.8}$$

where the "temperature" T sets the strength of the gain of the gain function $g(x) = \frac{1}{2}[1 + \tanh(x)]$.

2. *Free phase.* The input units alone are clamped and the remaining units settle according to

$$V_i' = g(\sum_j T_{ij} V_j'/T). \tag{1.9}$$

In each of these phases, the units settle with a decreasing sequence of temperatures, $T_o > T_1 > \ldots > T_{final}$. This process is called simulated annealing (Rumelhart and McClelland 1986). Equations 1.8 and 1.9 are normally updated asynchronously. We have checked that convergence is also reached with synchronous updating, which is more natural for an optical implementation.

After the two phases, updating (or learning) takes place with the rule

$$\Delta T_{ij} = \beta(V_i V_j - V_i' V_j') \tag{1.10}$$

where β is the learning rate. As it stands, equation 1.10 requires storing the solutions to equations 1.8,1.9 and subtracting, neither of which is natural in an optical environment. There are no fundamental obstacles, however, to doing intermediate updating. That is, we update with

$$\Delta T_{ij} = \beta V_i V_j \tag{1.11}$$

after the clamped phase and

$$\Delta T_{ij} = -\beta V_i' V_j' \tag{1.12}$$

after the free phase. We have checked performance with this modification and again find very little degradation.

The grating strengths $\eta_{ij}^{1/2}$ must necessarily be positive numbers less than 1. However, most neural network algorithms require that T_{ij} can take both positive and negative values, constituting a problem for both optical and electronic implementations. The most straightforward solution of several possible solutions (Peterson *et al.* 1989) to this problem is to have two sets of positive gratings, one for positive enforcement (T_{ij}^+) and one for negative (T_{ij}^-). The negative sign is then enforced electronically with a subtraction and equation 1.1 reads

$$V_i = g[\sum_j (T_{ij}^+ V_j - T_{ij}^- V_j)]. \tag{1.13}$$

In the modified MFT learning algorithm, the adjustment of equation 1.11 is always positive while the adjustment of equation 1.12 is always negative. So the clamped phase need only affect positive weights and the free phase need only affect negative weights. In Figure 2, generic read and write cycles for MFT are shown (only a slice of the volume is

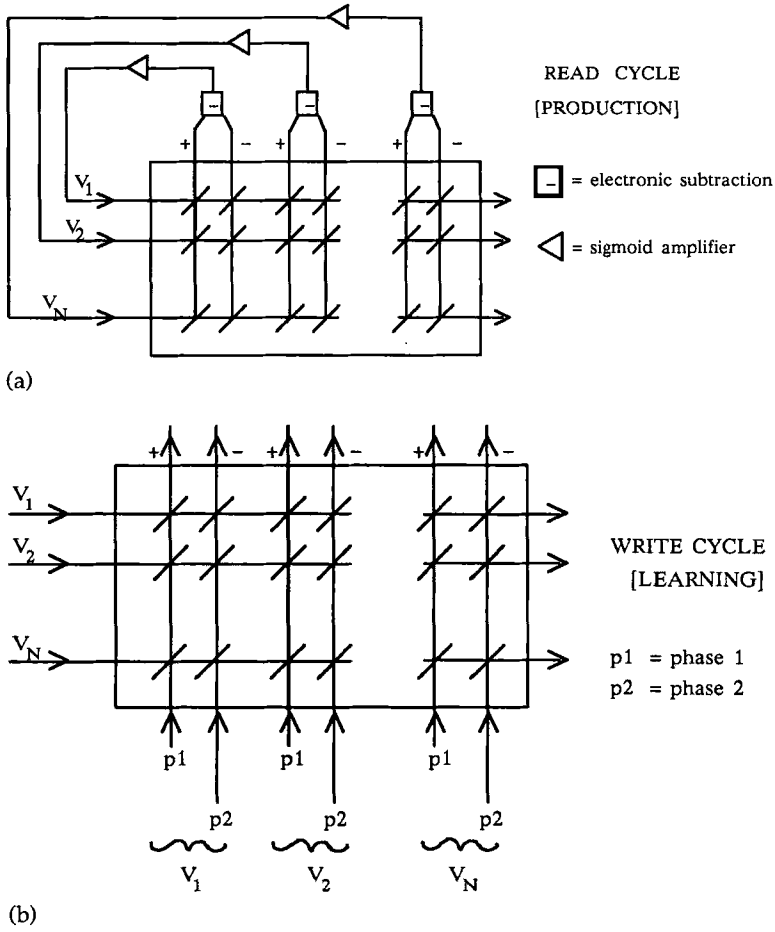

Figure 2: (a) An MFT read (production) cycle. (b) An MFT write (learning) cycle.

shown). As can be seen, each connection strength is represented in the crystal by two gratings, T_{ij}^{+} and T_{ij}^{-}. Thus n neurons require $2n^2$ connections (for full connectivity). In the read (production) cycle, the reference beam is iteratively processed through the crystal until it has settled according to equation 1.8 (or equation 1.9).

The write (learning) cycle (Fig. 2), is slightly more complicated and differs between the clamped and free phases. In the clamped phase, the beam again settles according to equation 1.8. It is then replicated as two

beams, "reference" and "object". The two beams impinge simultaneously on the crystal with the object beam hitting only the T_{ij}^{+} columns. The interconnect weights are then adjusted where the beams cross. The free phase proceeds in the same way except that the object beam hits only the T_{ij}^{-} columns.

We now briefly discuss the implementation of feedforward models, again using a single crystal with spatial multiplexing. Optical implementations of the backpropagation (BP) algorithm have been investigated elsewhere (Psaltis *et al.* 1982), but these investigations have focused mainly on angular multiplexing with a multi-crystal architecture.

We restrict the discussion to three layer networks (one hidden layer) with input-to-hidden and hidden-to-output connections. For such networks, symmetric feedback output-to-hidden connections are required for BP (as they are in MFT). The neurons are arranged on the SLMs in layers. We denote the input-to-hidden weights as T_{jk} , hidden-to-output as W_{ij} , and output-to-hidden as $W_{ji} = W_{ij}$.

We have developed a procedure for implementing BP in the crystal such that the updates to the W_{ij} weights are exact, and the updates to the T_{jk} weights are correct to first-order in β. Hence, for learning rates small enough compared with the magnitude of W_{ij} and $g'(h_i)$, this procedure is faithful to the exact BP algorithm. While space does not allow a detailed description (see Peterson *et al.* 1989), the procedure entails a two-stage backward pass analogous to the MFT modification of equations 1.11,1.12: for each pattern, the set of "positive" ("negative") gratings are written with the weight change component due to the positive (negative) term in the error expression.

Other learning algorithms for feedforward networks, such as the Kanerva associative memory (Kanerva 1986) and the Marr and Albus models of the cerebellum (Keeler 1988), can be implemented in a very similar (and even simpler) manner.

We have conducted simulations of the modified MFT and BP learning algorithms on the spatially multiplexed architecture. In addition to rescattering, beam depletion, and double columns of weights, the simulations contained the following ingredients:

Temperature gradient. Beam depletion (see Fig. 1 and absorption, below) is the one effect in the crystal for which the MFT and BP algorithms were not able to automatically compensate. We found that we could counterbalance the asymmetry of the emergent light by varying the gain (or inverse temperature) across the crystal. The gain increases with depth into the crystal. Without this technique, none of our simulations were successful.

Absorption. The crystal absorbs light, both in read and write phases, exponentially with depth into the crystal.

Grating decay. During continuous exposure to illumination, the crystal gratings decay exponentially in time.

Simulations were performed for three different problems: random mappings of random (binary) vectors, the exclusive-or (XOR) or parity problem, and the 6×6 mirror symmetry problem (Peterson and Hartman 1988).

Both MFT and BP successfully learned all three problems in simulations. In Figure 3 we show the results for the mirror symmetry problem using 12 hidden units and 36 training patterns. As can be seen from this figure, the supervised learning algorithms have the expected property of adjusting the weights such that the various physical effects are taken care of.

The system configuration has two principal optical paths, a reference path (1) and an object path (2) (see Fig. 4). Each path has a spatial filter, a beam splitter, a SLM, and an imaging lens system. The object path ends with a CCD array. The photorefractive crystal is SBN and an argon ion laser is used as a coherent light source. Thresholding $(1/2(1 + \tanh(x)))$

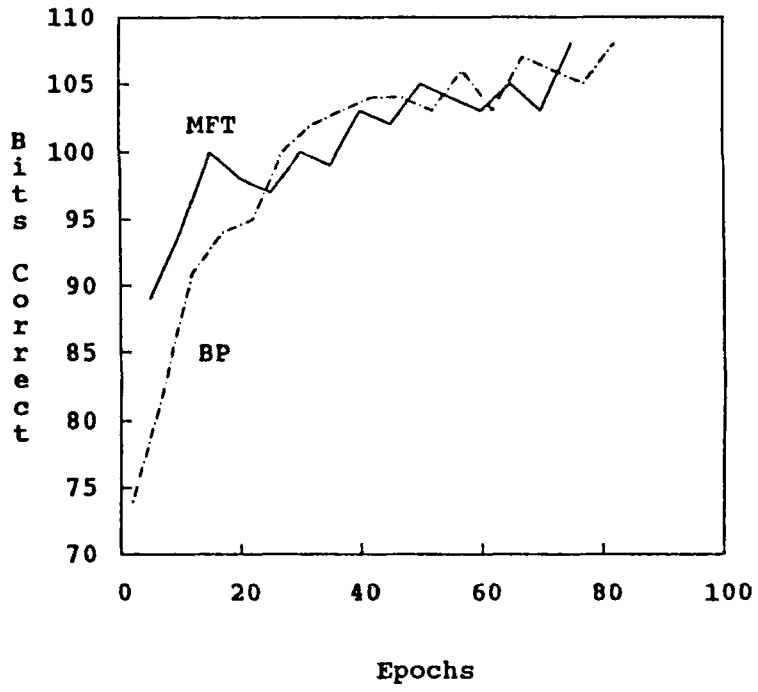

Figure 3: Learning performance of MFT and BP on the 6×6 mirror symmetry problem (horizontal axis is number of training epochs).

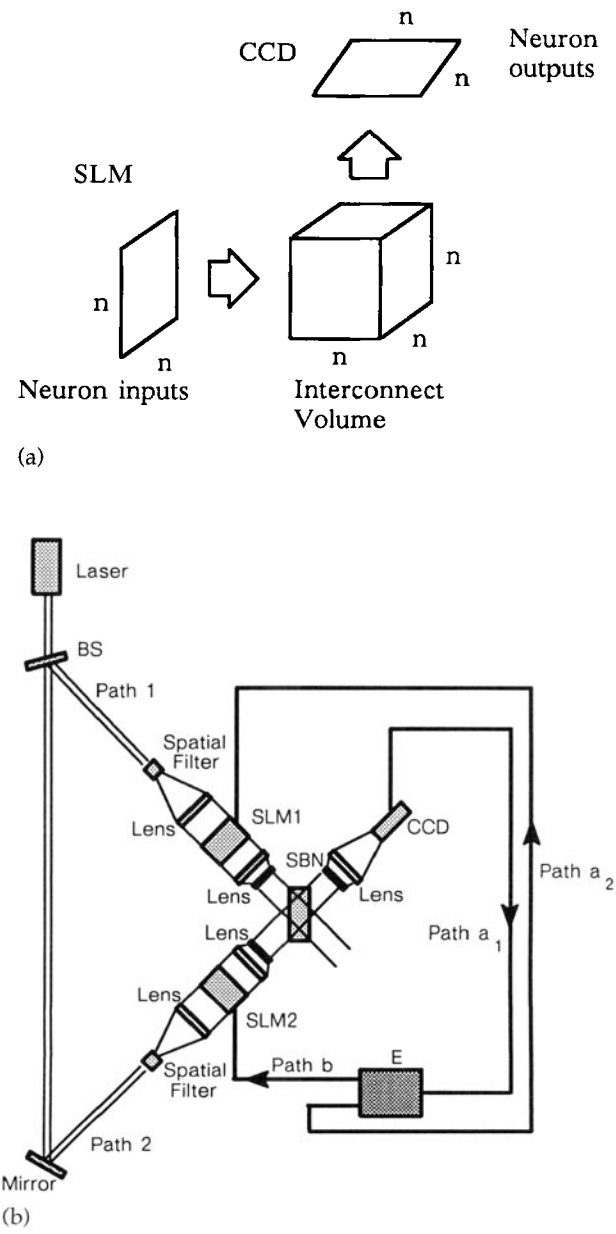

Figure 4: System configuration.

and loading of the SLMs take place in E electronically. Both mean-field-theory and backpropagation learning algorithm implementations have distinctive read (a_1 plus a_2) and write (a_2 plus b) phases that use this generic architecture.

References

Anderson, D. 1987. Adaptive interconnects for optical neuromorphs: Demonstrations of a photorefractive projection operator. In *Proceedings of the IEEE First International Conference on Neural Networks* **III**, 577.

Kanerva, P. 1986. *Sparse Distributed Memory.* MIT Press/Bradford Books, Cambridge, MA.

Keeler, J. D. 1988. Comparison between Kanerva's SDM and Hopfield-type neural networks. *Cognitive Science* **12**, 299.

Peterson, C. and Hartman, E. 1988. Explorations of the mean field theory learning algorithm. *Neural Networks* **2**, 475–494.

Peterson, C. and Redfield, S. 1988. A novel optical neural network learning algorithm. In *Proceedings of the ICO Topical Meeting on Optical Computing,* J. W. Goodman, P. Chavel, and G. Roblin, eds. SPIE, Bellingham, WA, 485–496.

Peterson, C., Redfield, S., Keeler, J. D., and Hartman, E. 1989. Optoelectronic implementation of multilayer neural networks in a single photorefractive crystal. MCC Tech. Rep. No. ACT-ST-146-89 (to appear in *Optical Engineering*).

Psaltis, D. and Farhat, N. 1985. Optical information processing based on associative memory model of neural nets with thresholding and feedback. *Optics Letters* **10**, 98.

Psaltis, D., Wagner, K., and Brady, D. 1987. Learning in optical neurocomputers. *Proceedings of the IEEE First International Conference on Neural Networks,* III–549; Wagner, K., and Psaltis, D. 1987. Multilayer optical learning networks. *Applied Optics* **3**, 5061.

Redfield, S. and Hesselink, B. 1988. Enhanced nondestructive holographic readout in SBN. *Optics Letters* **13**, 880.

Rumelhart, D. E. and McClelland, J. L., eds. 1986. *Parallel Distributed Processing: Explorations of the Microstructure of Cognition. Vol. 1: Foundations.* MIT Press, Cambridge, MA.

Soffer, B. H., Dunning, G. J., Owechko, Y., and Marom, E. 1986. Associative holographic memory with feedback using phase conjugate mirrors. *Optics Letters* **11**, 118.

Received 18 August 1989; accepted 21 December 1989.