

Algebraic Transformations of Objective Functions

Eric Mjolsness and Charles Garrett

Research Report YALEU/DCS/RR-686

March 1989

Algebraic Transformations of Objective Functions

Eric Mjolsness and Charles Garrett
Department of Computer Science
Yale University

March 7, 1989

Abstract

A standard neural network design trick reduces the number of connections in the winner-take-all (WTA) network from $\mathcal{O}(N^2)$ to $\mathcal{O}(N)$. We explain the trick as a general fixpoint-preserving transformation applied to the particular objective function associated with the WTA network. The key idea is to introduce new interneurons which act to maximize the objective, so that the network seeks a saddle point rather than a minimum. A number of fixpoint-preserving transformations are derived, allowing the simplification of such algebraic forms as products of expressions, functions of one or two expressions, and sparse matrix products. The transformations may be applied to reduce or simplify the implementation of a great many structured neural networks, as we demonstrate for inexact graph-matching, convolutions and coordinate transformations, and sorting. Simulations show that fixpoint-preserving transformations may be applied repeatedly and elaborately, and the example networks still robustly converge. We discuss implications for circuit design.

1 Introduction

Objective functions have become important in the study of artificial neural networks, for their ability to concisely describe a network and the dynamics of its neurons and connections. For neurons with objective-function dynamics, the now standard procedure [Hopfield 1984, Hopfield & Tank 1985, Koch Marroquin & Yuille 1986] is to formulate an objective function (called the "objective" in what follows) which expresses the goal of the desired computation, then to derive a local update rule (often a simple application of steepest descent) which will optimize the dynamic variables, in this case the artificial analog neurons. The update rule should ultimately converge to a fixpoint which minimizes the objective and should be interpretable as the dynamics of a circuit.

In this paper we propose to modify the standard procedure by interpolating an extra step: the objective should be formulated, then algebraically manipulated in a way which preserves its meaning but improves the resulting circuits, e.g. by decreasing some measure of their cost. Then the improved circuit is derived from the modified objective. We do not require that the algebraic manipulation be done automatically; we just ask whether and how it can be done. The answer is in the form of a short, non-exhaustive list of very general transformations which can be performed on summands of an objective, provided they are of the requisite algebraic form, without altering the fixpoints of the resulting network. Often it will be clear from the algebra alone that some savings in number of connections or neurons will occur, or that a previously non-implementable objective has been transformed to a form (e.g. single-neuron potentials plus polynomial

This work was supported in part by AFOSR grant AFOSR-88-0240.

interactions [Hopfield 1984]) whose minimization may be directly implemented as an analog circuit in a given technology. It also seems likely that many useful and general algebraic transformations await discovery.

Although we use the terminology appropriate to the optimization of objectives for dynamic neurons with fixed connections, much of the theory may apply also to "learning" considered as the optimization of dynamic connections in an unstructured net (e.g. [Rumelhart Hinton & Williams 1986]) or of some smaller set of parameters which indirectly determine the connections in a structured net [Mjolsness Sharp & Alpert 1989].

In the remainder of this section, we will introduce the ideas by rederiving a well-known network simplification: that of the winner-take-all (WTA) network. Section 2 develops the theory of our algebraic transformations, including the reduction of: squares and products of expressions; a broad class of functions of one and two expressions; the minimum and maximum of a set of expressions; and certain matrix forms which retain their sparseness as they are reduced. Implications for circuit design are discussed, and a major unsolved problem related to the handling of sparse matrices is stated. The fixpoint-preserving transformations are cataloged in section 2.6. In section 3, some of the transformations are exercised in design examples. Experimental results are available for a graph-matching network, which is guaranteed to work if run slowly enough but also works at reasonable speed, and for a sorting network, for which an illegitimate transformation step (encoding a sparse matrix, as described in Appendices B and C) introduces spurious local minima which merely decrease the accuracy of the final answer. The sorting network also involves a series of legitimate fixpoint-preserving transformations which are guaranteed to work if the network is run slowly enough but which actually work at reasonable speed. Finally, a discussion follows in section 4.

1.1 Reversed Linear Neurons in the WTA Network

Consider the ordinary winner-take-all analog neural network. Following [Hopfield & Tank 1985], such a network can be obtained from the objective

$$E_{\text{wta}}(\vec{v}) = \frac{c_1}{2} \left(\sum_i v_i - 1 \right)^2 + \frac{c_2}{2} \sum_i v_i (1 - v_i) + \sum_i \phi(v_i) \quad (c_1 > 0) \quad (1)$$

where [Hopfield 1984, Grossberg 1988]

$$\phi(v_i) = \int^v dx g^{-1}(x), \quad (2)$$

using steepest-descent dynamics

$$\dot{v}_i = -\partial E / \partial v_i \quad (3)$$

or Hopfield-style dynamics

$$\dot{u}_i = -\partial E / \partial v_i, \quad v_i = g(u_i). \quad (4)$$

The resulting connection matrix is

$$T_{ij} = -c_1 + c_2 \delta_{ij},$$

which implies global connectivity among the neurons: if there are N neurons, there are N^2 connections. It is well known that the winner-take-all circuit requires only $\mathcal{O}(N)$ connections if one introduces a linear neuron σ whose value is always $\sum_i v_i$. It is not so well known that this can be done entirely within the objective function, as follows:

$$\hat{E}_{\text{wta}}(\vec{v}) = c_1 \left(\sum_i v_i - 1 \right) \sigma - \frac{c_1}{2} \sigma^2 + \frac{c_2}{2} \sum_i v_i (1 - v_i) + \sum_i \phi(v_i) \quad (5)$$

where the steepest-descent dynamics, for example, is modified to become

$$\begin{aligned} \dot{v}_i &= -r_i \partial \hat{E} / \partial v_i \\ &= -\partial \hat{E} / \partial v_i \quad (\text{specialize to } r_i = 1) \end{aligned} \quad (6)$$

$$\begin{aligned}
\dot{\sigma} &= +r_\sigma \partial \hat{E} / \partial \sigma \quad (r_\sigma > 0) \\
&= c_1 r_\sigma (-\sigma + \sum_i v_i - 1) \\
&= 0 \text{ if } \sigma = \sum_i v_i - 1.
\end{aligned} \tag{7}$$

But if $\sigma = \sum_i v_i - 1$ then one can calculate that $\partial E / \partial v_i = \partial \hat{E} / \partial v_i$; thus equations (1) and (5) have the same fixpoints. The connectivity implied by counting the monomials in equation (5) is $\mathcal{O}(N)$ connections, the minimum possible for this problem.

Note that the σ linear neuron actually behaves so as to *increase* the objective $E(V, \sigma)$, while the v_i neurons act to decrease it; σ may be called a *reversed* neuron. Reversed neurons introduce a new element of competition into a network; indeed, two-person zero-sum games are usually modelled using objectives which one player increases and the other decreases [von Neumann & Morgenstern 1953]. So in this network, and in the others we will introduce, minimization is replaced with finding a saddle point and the problem becomes hyperbolic. This follows immediately from the sign of σ^2 in (5), which eliminates all local minima. Fortunately there are hyperbolic versions of such efficient optimization procedures as the conjugate gradient method; Luenberger [Luenberger 1984] gives two examples.

For finite r_σ , σ is a delayed version of the sum $\sum_i v_i - 1$ and although the network dynamics are different from equation (3), the fixed point is the same. Alternatively, the rate parameter r_σ may be adjusted to make σ move at a different time scale from the rest of the neurons. As r_σ approaches infinity, σ becomes an *infinitely fast* neuron whose value is always

$$\sigma = \sum_i v_i - 1$$

and the other neurons see an effective objective

$$\hat{E}_{\text{wta}}(\vec{v}, \sigma(\vec{v})) = E_{\text{wta}}(\vec{v}) \tag{8}$$

so that their dynamics become identical to that of the original fully connected winner-take-all network. There are very efficient serial and parallel implementations for networks with $r_\sigma \rightarrow \infty$, which update the infinitely fast neuron whenever its ordinary neighbors change; this is a standard trick in neural network and Monte Carlo physics simulations. When it is applied to the simplify the simulation of the WTA equations of motion, we arrive at the standard WTA trick.

Moody [Moody 1989] independently discovered an objective function equivalent to (5) and first simulated its delayed Hopfield-style equations of motion (see (4)). But he remained unaware that the σ neuron acts to maximize rather than minimize the objective, driving the system to a saddle point. Platt and Barr [Platt & Barr 1988, Platt & Barr 1987] performed constrained optimization (including multiple WTA constraints) by using a subset of neurons that explicitly increase the objective, and hence a network that seeks saddle points rather than local minima. Indeed reversed neurons are a generalization of their analog Lagrange multiplier neurons, which were found earlier in a non-neural context by Arrow [Arrow 1958].

Both reversed neurons and Lagrange multiplier neurons act to maximize an objective which other neurons act to minimize. The difference is that Lagrange multiplier neurons must appear linearly in the objective. General reversed neurons can have self-interactions; in particular, the potential of a linear reversed neuron like σ in the winner-take-all network is

$$\begin{aligned}
-\frac{1}{2g_0} \sigma^2 &= \int^\sigma dx g^{-1}(x) \\
\text{where } g(x) &= -g_0 x
\end{aligned} \tag{9}$$

i.e. linear reversed neurons have linear transfer functions with negative gain. Thus in circuit language they are just inverters, which happen to occur in a network with an objective function and to act so as to increase E . Lagrange multiplier neurons, on the other hand, have no potential term and are not inverters. It is also

worth noting that Lagrange multiplier neurons work best in conjunction with an additional penalty term $h(\vec{v})^2/2$ (where $h(\vec{v}) = 0$ is the constraint) and the penalty term can be efficiently implemented using one reversed neuron per constraint, as we will see.

If in addition to being reversed, a neuron is also infinitely fast, then it may be necessary to restrict its connectivity in order to efficiently simulate or implement the network. One possible design rule is to entirely prohibit connections between infinitely fast neurons; this prevents one from having to solve a system of linear equations in order to update a set of infinitely fast neurons. We will *not* generally assume that reversed neurons are infinitely fast.

2 Theory

The reversed linear neuron is applicable in many circumstances beyond the winner-take-all network. We can begin to see its generality by considering objectives of the form

$$E(\vec{v}) = E_0(\vec{v}) + \frac{c}{2}X^2(\vec{v})$$

where E_0 and X are any algebraic expressions, and c is a constant of either sign. This may be transformed to

$$\hat{E}(\vec{v}, \sigma) = E_0 + cX\sigma - \frac{c}{2}\sigma^2$$

and if X is a polynomial, this represents a reduction in the number and order of the monomials that occur in E . The transformation technique used here is simple to state: find squared expressions $(c/2)X^2$ as summands in the objective function, and replace them with $cX\sigma - (c/2)\sigma^2$. c must be a constant and σ is a new linear interneuron, reversed if c is positive. Most, but not all, of our reversed linear interneurons will be introduced this way. The transformation may be abbreviated as

$$\frac{1}{2}X^2 \rightarrow X\sigma - \frac{1}{2}\sigma^2. \quad (10)$$

We employ the steepest-ascent-descent dynamics

$$\begin{aligned} \dot{v}_i &= -\partial \hat{E} / \partial v_i \\ &= -\partial E_0 / \partial v_i - c\sigma \partial X / \partial v_i, \\ \dot{\sigma} &= +(\tau_\sigma / c) \partial \hat{E} / \partial \sigma \\ &= \tau_\sigma (X - \sigma) \end{aligned}$$

which, at a fixpoint, has $X = \sigma$ and $\partial E_0 / \partial v_i + cX \partial X / \partial v_i = \partial E / \partial v_i = 0$. Likewise a fixpoint of E can be extended, by setting $\sigma = X$, to a fixpoint of \hat{E} . So fixpoints are preserved by the transformation (10). The argument also works if some of the v_i are already reversed neurons, so that $\dot{v}_i = \pm \partial \hat{E} / \partial v_i$.

As an example of the transformation (10), one can robustly implement a constraint $h(\vec{v}) = 0$ using both a penalty term $ch^2/2$ and a Lagrange multiplier neuron. The objective becomes

$$E_{\text{constraint}}(\vec{v}, \sigma, \lambda) = c\sigma h(\vec{v}) - c\sigma^2/2 + \lambda h(\vec{v}). \quad (11)$$

2.1 Products and Order Reduction

From the transformation (10) we may deduce two others, which transform a summand which is a product XY of expressions X and Y . One transformation is

$$\begin{aligned} XY &= \frac{1}{2}(X+Y)^2 - \frac{1}{2}X^2 - \frac{1}{2}Y^2 \\ &\rightarrow (X+Y)\sigma - X\tau - Y\omega - \frac{1}{2}\sigma^2 + \frac{1}{2}\tau^2 + \frac{1}{2}\omega^2 \\ &\rightarrow X(\sigma - \tau) + Y(\sigma - \omega) - \frac{1}{2}\sigma^2 + \frac{1}{2}\tau^2 + \frac{1}{2}\omega^2 \end{aligned} \quad (12)$$

Here σ is a reversed linear neuron, and τ and ω are ordinary linear neurons. If all three linear interneurons are infinitely fast, which is easy to simulate since they are not directly connected, then the transformation does not change the dynamics of the rest of the variables in the network. Otherwise, the dynamics and the basins of attraction change, but the network fixed points remain the same.

This transformation may be simplified,¹ to

$$\begin{aligned} XY &= \frac{1}{4} [(X+Y)^2 - (X-Y)^2] \\ &\rightarrow \frac{1}{2}(X+Y)\sigma - \frac{1}{2}(X-Y)\tau - \frac{1}{4}\sigma^2 + \frac{1}{4}\tau^2 \\ &= \frac{1}{2}X(\sigma - \tau) + \frac{1}{2}Y(\sigma + \tau) - \frac{1}{4}\sigma^2 + \frac{1}{4}\tau^2. \end{aligned} \quad (13)$$

Compared to equation (12) this transformation results in the same number of monomial interactions and one less neuron, which may be useful on occasion.

Reversed neurons allow one to transform a high-order polynomial objective, monomial by monomial, into a third-order objective. Similar transformations on the neural networks or analog circuits are well known. But it is easier to do theoretical work with the objective, and by transforming the objective first, and then translating to a neural net, one can obtain novel third order neural nets.

We may expand a high-order polynomial objective into monomials, each of which corresponds to one connection or "synapse" of the associated neural network. We may reduce the order of an entire objective by reducing the order of each monomial. Consider, then, a single fourth-order monomial:

$$E_{\text{mono}}(x, y, z, w) = -Txyzw \quad (14)$$

which by (12) may be transformed to

$$\hat{E}_{\text{mono}}(x, y, z, w, \sigma, \tau, \omega) = T[xy(\tau - \sigma) + zw(\omega - \sigma) + \frac{1}{2}\sigma^2 - \frac{1}{2}\tau^2 - \frac{1}{2}\omega^2]. \quad (15)$$

Here σ , τ , and ω are linear neurons with gain $1/T$. The order-reducing transformation is illustrated in network language in figure 1.

The same technique may be used to recursively transform a monomial of any order m to a sum of third order monomials, plus potentials for the new linear interneurons. The resulting number of new monomials and interneurons is $< \mathcal{O}(m(\log m)^2)$ if the reduction is done in a balanced way and if expressions like $X(\sigma - \tau)$ are not expanded out to $X\sigma - X\tau$ during the reduction.

Another order-reduction transformation, superior in some circumstances, will be developed in section 2.3.

2.2 Reducing $F(X)$

To build facility in the algebraic manipulation of objective functions, we accumulate a catalog of fixpoint-preserving transformations which are recorded in section 2.6. In addition to transformation (10) and its consequences, there is a major class of transformations associated with functions of a single expression, $F(X)$:

$$\int^X f(u)du \rightarrow X\sigma - \int^\sigma f^{-1}(u)du \quad (f \text{ invertable}) \quad (16)$$

Note that X no longer appears inside the function $F = \int f$. The validity of this transformation (equation 16) may be proven by noting that the optimal value of σ is $f(X)$, and then either integrating by parts the expression

$$\int^{f(X)} f^{-1}(u)du = \int^X v f'(v)dv,$$

¹as pointed out to us by P. Anandan




Figure 1: **Order Reduction.** Arbitrary fourth to third order reduction, using linear interneurons. Open circles: original neurons. Open squares: ordinary linear interneurons. Closed squares: reversed linear interneurons. Dots: connections, with strengths as indicated. Equilibrium values of the interneurons are indicated. After the transformation, neuron w receives input $xyz + z^2w$ from the left and a compensating input of $-z^2w$ from the right.

or else differentiating both pre- and post-transformation objectives with respect to X .

In this way one can treat functions $\exp X$, $|X| \log |X|$, $\log |X|$, and $|X|^p$ of arbitrary algebraic expressions X :

$$\begin{aligned}
 e^X &\rightarrow (X + 1)\sigma - \sigma \log \sigma \\
 |X| \log |X| &\rightarrow |X|(\sigma + 1) - e^\sigma \\
 \log |X| &\rightarrow X\sigma - \log |\sigma| \\
 \frac{1}{1+p}|X|^{1+p} &\rightarrow X\sigma - \frac{1}{1+1/p}|\sigma|^{1+1/p} \quad (p \neq -1, 0).
 \end{aligned} \tag{17}$$

Thus, neural nets may be constructed from some highly nonpolynomial objectives.

The interneurons may still be reversed, but are no longer linear, in this kind of transformation. The potential $\phi(\sigma)$ permits a possible efficiency technique. The dynamics of equation (4) is expected to be more efficient than equation (3), since it may be viewed as a quasi-Newton method which takes into account the potential but not the interaction part of a neural net objective (as shown by J. Utans [Utans et. al. 1989]). A related update scheme for the σ reversed interneuron is

$$\begin{aligned}
 \sigma &= f(s), & \dot{s} &= r_\sigma \partial E / \partial \sigma = r_\sigma (X - s) \\
 v_i &= g_i(u_i), & \dot{u}_i &= -\partial E(\vec{v}, \sigma) / \partial v_i
 \end{aligned} \tag{18}$$

which is an alternative to direct steepest-ascent-descent. This dynamics has the distinct advantage of a simple interpretation in terms of analog electrical circuits [Hopfield 1984]. For example, $F(X) = X(\log X - 1)$ requires a special neuron whose transfer function is logarithmic. This can be provided, approximately and within a mildly restricted domain of the input values, in analog VLSI [Sivilotti Mahowald & Mead 1987]. Similarly $F(X) = \log X$ would require a transfer function $f(s) = 1/s$, and an exponential transfer function would lead to $F(X) = \exp X$. It might be possible to characterize a particular technology by a list of the basic forms

of objectives it makes available, with their respective costs and restrictions, and to compile general networks into the desired forms by using a catalog of algebraic transformations.

It may also be worth noting that the linear programming network of [Tank & Hopfield 1986] can be interpreted as an application of transformation (16) with $r_\sigma \rightarrow \infty$ in the dynamics of equation (18).

Equation (16) generally may be used to transform a term $F(X)$ in an objective by transferring the nonlinearity due to F from $F(X)$ to the single-neuron potential $\phi(\sigma) = \int^\sigma f^{-1}(u)du$. By transferring the nonlinearity from an interaction term (i.e. a summand of the objective which involves several dynamical variables) to a single-neuron potential term, one can not only decrease the cost of implementing a network which uses gradient methods for optimization, but one can transform unimplementable objectives into implementable ones. For example, one might regard the class of multi-variable polynomials as the "implementable" interactions in a certain technology [Rumelhart Hinton & McClelland 1986]. (In this case the word "interaction" is usually reserved for a multivariable monomial, out of which polynomials are built by addition of objectives.) Then one might use equation (16) to reduce other, far more general interaction objectives to the implementable form.

Of course the required potential $\phi(\sigma)$ may itself be "unimplementable", but approximating its gradient with a small circuit is likely to be far more tractable than approximating $\nabla F(X)$ because ϕ , unlike F , is a function of just one dynamic variable. An approximation of $\phi'(\sigma)$ might be formulated as

$$\phi(\sigma) \rightarrow \hat{\phi}(\sigma) = \sum_{\alpha} c_{\alpha} \hat{\phi}_{\alpha}(\sigma)$$

where each $\hat{\phi}_{\alpha}(\sigma)$ is regarded as an implementable self-interaction and c_{α} are adjustable coefficients.

2.3 Interacting Expressions: Reducing $G(X, Y)$

Until now we have attempted to reduce all interactions to the forms xy and xyz , but those may not be the only cost-effective interactions allowed by a given physical technology.

From equation (16) we may derive a generalization to functions of two expressions, $G(X, Y)$:

$$\begin{aligned} \int^X du \left[\int^Y dv g(u, v) \right]^{-1} (u) & \quad \text{[note function inverse]} \\ \rightarrow X\sigma - \int^\sigma du \int^Y dv g(u, v) & \quad \text{[by (16)]} \\ = X\sigma - \int^Y dv \int^\sigma du g(u, v) & \\ \rightarrow X\sigma - Y\tau + \int^\tau dv \left[\int^\sigma du g(u, v) \right]^{-1} (v) & \quad \text{[by (16)].} \end{aligned}$$

Thus,

$$\int^X du \left[\int^Y dv g(u, v) \right]^{-1} (u) \rightarrow X\sigma - Y\tau + \int^\tau dv \left[\int^\sigma du g(u, v) \right]^{-1} (v) \quad (19)$$

Of course, there are complicated restrictions regarding the existence of the inverse functions used.

Taking

$$g(u, v) = \frac{1}{2\sqrt{uv}}$$

we can derive the transformation $-Y/X \rightarrow X\sigma - Y\tau - \sigma/\tau$. Rescaling Y and σ by -1 , then switching X for Y and σ for τ , this is equivalent to

$$X/Y \rightarrow X\sigma - Y\tau + \tau/\sigma \quad (20)$$

which is linear in τ but effectively nonlinear due to the optimization of σ .

From equation 19, Appendix A derives two transformations for the special form $YF(X)$:

$$YF(X) \rightarrow -X\sigma + Y\tau + \sigma F^{-1}(\tau) \quad (21)$$

(which implies (20)) and

$$Y \int^X f(u) du \rightarrow XY\sigma - Y \int^\sigma du f^{-1}(u) \quad (22)$$

assuming $f = F'$ is invertible and $f^{-1} = (F')^{-1}$ is differentiable.

Monomial order reduction can sometimes be accomplished more cheaply using $x \log y$ interactions than third-order ones. If v_i are all restricted to be positive, then

$$\begin{aligned} \prod_{i=1}^m v_i &= \exp \sum \log v_i \\ &\rightarrow \sigma (\sum_i \log v_i + 1) - \sigma \log \sigma \\ &= \sum_i \sigma \log v_i + \sigma (1 - \log \sigma). \end{aligned} \quad (23)$$

This objective has $\mathcal{O}(m)$ interactions of the new type. The fixpoint value of σ is $\prod_i v_i$ at which point the steepest-descent input to v_i is $-\sigma/v_i = -\prod_{j \neq i} v_j$.

A product of expressions could be further reduced using equation (21) and xe^y interactions:

$$\prod_{\alpha=1}^m |X_\alpha| \rightarrow \sum_{\alpha} (\sigma \tau_\alpha - |X_\alpha| \omega_\alpha + \omega_\alpha e^{\tau_\alpha}) + \sigma (1 - \log \sigma). \quad (24)$$

2.4 Min and Max

The minimum or maximum of a set of expressions $\{X_\alpha\}$ can be implemented using a winner-take-all network in which each expression is represented by a neuron σ_α which competes with the others and is constrained to lie between 0 and 1. Indeed, $\sum_{\alpha} X_\alpha \sigma_\alpha$ attains the value $\min_{\alpha} X_\alpha$ when the correct representative wins, and also provides inhibition to the representatives in proportion to the values of the expressions they represent, so that the correct representative will win. The potential $\phi(\sigma)$ that occurs in the WTA network must have only one minimum, so that there is no hysteresis in the circuit, and must closely approximate a square well i.e. must have high gain. Under these circumstances, we can transform

$$E = \min_{\alpha} X_\alpha \rightarrow \sum_{\alpha} X_\alpha \sigma_\alpha + C (\sum_{\alpha} \sigma_\alpha - 1) \lambda - \frac{C}{2} \lambda^2 + \sum_{\alpha} \phi(\sigma_\alpha) \quad (25)$$

(where C and the gain of ϕ are sufficiently large) and at any fixpoint of $\vec{\sigma}$ the derivatives of E with respect to all other dynamical variables will be preserved. So, fixpoints will be preserved.

Likewise

$$\max_{\alpha} X_\alpha \rightarrow \sum_{\alpha} X_\alpha \sigma_\alpha - C (\sum_{\alpha} \sigma_\alpha - 1) \lambda + \frac{C}{2} \lambda^2 - \sum_{\alpha} \phi(\sigma_\alpha). \quad (26)$$

An alternate transformation for max proceeds through the identity

$$\max_{\alpha} X_\alpha = \lim_{p \rightarrow \infty} \left(\sum_{\alpha} X_\alpha^p \right)^{1/p} \quad (X_\alpha > 0). \quad (27)$$

These transformations have application in the design of certain "content addressable memories" for which the ideal objective function, which must be translated to a form polynomial in its interactions, may be taken as

$$E_{CAM}(\vec{s}) = - \max_{\text{memories } m} \vec{s} \cdot \vec{s}^m - \epsilon \vec{s} \cdot \vec{h}_{\text{input}} + \sum_i \phi_{\pm 1}(s_i)$$

with $-1 \leq s_i \leq 1$. Using the transformation (26) yields a CAM with one "grandmother neuron" (representative neuron) per memory, and a WTA net among them (a network described in [Moody 1989]). Another efficient CAM design [Gindi Gmitro & Parthasarathy 1987, Moody 1989] is essentially an application of transformations (27) and (17d), omitting the power $1/p$ and adjusting ϵ and ϕ to compensate.

2.5 Matrices, Graphs, and Pointers

One can apply order reduction to objectives containing polynomials of matrices. For dense $N \times N$ matrices, a typical term like $\text{tr} \prod_{l=1}^L A^{(l)}$ contains N^L scalar monomial interactions, but this can be reduced to $\mathcal{O}(N^3 L (\log L)^2)$. (Here $\text{tr} A \equiv$ trace of $A \equiv \sum_i A_{ii}$.) To show this we need only establish the matrix analog of transformation (12), which upon iteration can reduce an L -th order matrix monomial to $\mathcal{O}(L (\log L)^2)$ third-order matrix monomials like ABC . Each of these involves N^3 scalar monomial interactions.

Using equation (12) (one could also use (13)), one can reduce $\text{tr} XY$, where X and Y are now matrix-valued expressions. This form has exactly the same generality as $\text{tr} XY^T$ (where $Y^T \equiv$ transpose of Y).

$$\begin{aligned} \text{tr} XY^T &= \sum_{ij} X_{ij} Y_{ij} \\ &\rightarrow \sum_{ij} \left(X_{ij} (\sigma_{ij} - \tau_{ij}) + Y_{ij} (\sigma_{ij} - \omega_{ij}) - \frac{1}{2} \sigma_{ij}^2 + \frac{1}{2} \tau_{ij}^2 + \frac{1}{2} \omega_{ij}^2 \right) \\ &= \text{tr} X (\sigma - \tau)^T + \text{tr} Y (\sigma - \omega)^T - \frac{1}{2} \text{tr} \sigma \sigma^T + \frac{1}{2} \text{tr} \tau \tau^T + \frac{1}{2} \text{tr} \omega \omega^T \end{aligned} \quad (28)$$

This transformation preserves the sparseness of X and Y in the following sense: if $X_{ij} = Y_{ij} = 0$ at a fixed point, then $\sigma_{ij} = \tau_{ij} = \omega_{ij} = 0$ and the contribution of these neurons to the gradient of the objective is also zero.

A major problem in neural network research (c.f. [Feldman 1982]) is to reduce the cost of networks which manipulate graphs. Usually [Hopfield & Tank 1985, Hopfield & Tank 1986, Mjolsness Gindi & Anandan 1989] objectives for such problems involve dense matrices of neurons representing all the possible links in a graph. But the graphs that arise in computer science and in computer programming usually have a relatively small number of links per node, and are therefore representable by sparse matrices. (If a sparse matrix's entries are all zero or one, it is equivalent to a set of "pointers" in many current computer languages. Pointers are used ubiquitously, wherever some fluidity in data representation is required.) Since we have just shown how to reduce a wide class of matrix objective functions to summands of the form $\text{tr} ABC$ while retaining sparseness, it becomes important to reduce this form further by exploiting the sparseness of the matrices involved:

$$\text{tr} ABC \rightarrow ?$$

where A , B , and C are sparse-matrix-valued dynamical variables. We do not yet know how to do this correctly.

One approach to this problem is through the use of codes, like the binary code, which can concisely name the pair of nodes connected by each nonzero matrix element. Zero matrix elements are not explicitly encoded and this is the advantage of the method. A disadvantage of such codes, for objective functions, is that the configuration space is altered in such a way that new local minima may be introduced, though the old ones will be preserved. A code which allows order reduction to proceed most advantageously is used in the sorting networks of section 3.3.

Another approach is used by Fox and Furmanky [Fox & Furmanky 1988]. Their load-balancing network involves binary encoding, but the network evolution is divided into a number of phases in which different classes of neurons are allowed to vary while most neurons are held constant. The connections are different from one phase to the next, and do not recur, so that the network is not directly implemented in terms of a circuit but rather requires "virtual" neurons and connections. A virtual neural network can be provided by suitable software on general-purpose computers, or perhaps by further objective function transformations leading to a real (and efficient) neural circuit; the latter alternative has not been achieved.

2.6 List of Transformations

Let X and Y be any algebraic expressions, containing any number of variables. We list the following fixpoint-preserving transformations of objectives, or summands thereof:

1.1	$\frac{1}{2}X^2 \rightarrow X\sigma - \frac{1}{2}\sigma^2$
1.2	$XY \rightarrow X(\sigma - \tau) + Y(\sigma - \omega) - \frac{1}{2}\sigma^2 + \frac{1}{2}\tau^2 + \frac{1}{2}\omega^2$
1.3	$XY \rightarrow \frac{1}{2}X(\sigma - \tau) + \frac{1}{2}Y(\sigma + \tau) - \frac{1}{4}\sigma^2 + \frac{1}{4}\tau^2$
2.1	$\int^X f(u)du \rightarrow X\sigma - \int^\sigma f^{-1}(u)du \quad (f \text{ an invertable function})$
2.2	$e^X \rightarrow (X + 1)\sigma - \sigma \log \sigma$
2.3	$ X \log X \rightarrow X (\sigma + 1) - e^\sigma$
2.4	$\log X \rightarrow X\sigma - \log \sigma $
2.5	$\frac{1}{1+p} X ^{1+p} \rightarrow X\sigma - \frac{1}{1+1/p} \sigma ^{1+1/p} \quad (p \neq -1, 0).$
3.1	$\int^X du \left[\int^Y dvg(u, v) \right]^{-1}(u) \rightarrow X\sigma - Y\tau + \int^\tau dv \left[\int^\sigma dug(u, v) \right]^{-1}(v)$ (If function inverses exist.)
3.2	$YF(X) \rightarrow -X\sigma + Y\tau + \sigma F^{-1}(\tau) \quad (\text{If } ((F')^{-1})' \text{ exists.})$
3.3	$Y \int^X f(u)du \rightarrow XY\sigma - Y \int^\sigma du f^{-1}(u) \quad (\text{If } (f^{-1})' \text{ exists.})$
3.4	$X/Y \rightarrow X\sigma - Y\tau + \tau/\sigma$
4.1	$\min_\alpha X_\alpha \rightarrow \sum_\alpha X_\alpha \sigma_\alpha + C(\sum_\alpha \sigma_\alpha - 1)\lambda - \frac{C}{2}\lambda^2 + \sum_\alpha \phi(\sigma_\alpha)$ (Large C , and high-gain hysteresis-free barrier function ϕ which confines σ_α to $(0, 1)$.)
4.2	$\max_\alpha X_\alpha \rightarrow \sum_\alpha X_\alpha \sigma_\alpha - C(\sum_\alpha \sigma_\alpha - 1)\lambda + \frac{C}{2}\lambda^2 - \sum_\alpha \phi(\sigma_\alpha)$ (Same conditions.)
4.3	$\prod_\alpha X_\alpha \rightarrow \sum_\alpha (\sigma\tau_\alpha - X_\alpha \omega_\alpha + \omega_\alpha e^{\tau_\alpha}) + \sigma(1 - \log \sigma).$
5.1	$\text{tr } XY^T \rightarrow \text{tr } X(\sigma - \tau)^T + \text{tr } Y(\sigma - \omega)^T$ $-\frac{1}{2}\text{tr } \sigma\sigma^T + \frac{1}{2}\text{tr } \tau\tau^T + \frac{1}{2}\text{tr } \omega\omega^T$ (All matrices, possibly sparse.)

(29)

The variables σ , τ , ω , and λ are assumed not to occur elsewhere in the original objective. Note that each transformation may have restrictions on its applicability, in addition to the particular form it matches.

We will report experiments only with transformations 1.1, 1.2, and 5.1 on this list. The rest are still theoretical.

These transformations may be iterated, at the expense of creating interactions between the added variables. They can be used to reduce the nonlinearity of the interactions in a neural network, transferring such nonlinearity to single-neuron potentials or distributing it among several simpler interactions.

3 Design Examples

3.1 Convolutions and Coordinate Transformations

Discrete convolutions

$$O_i = \sum_j K_{i-j} I_j$$

(where index subtraction is defined appropriately) and linear coordinate transformations

$$x'_i = \sum_j A_{ij} x_j + b_i$$

can both be expressed as sums of squared penalty terms in an objective:

$$E_{\text{conv}} = \frac{c}{2} \sum_i (O_i - \sum_j K_{i-j} I_j)^2 \quad (30)$$

or

$$E_{\text{coord}} = \frac{c}{2} \sum_i (x'_i - \sum_j A_{ij} x_j + b_i)^2, \quad (31)$$

with $c > 0$. (Equation (31) subsumes equation (30).) Alternatively the convolution or coordinate change could be turned into a hard constraint by using Lagrange multiplier neurons, but those procedures still work best when a penalty term exactly like equation (30) or equation (31) is added to the objective [Platt & Barr 1988, Luenberger 1984]. As they stand, these objectives expand into very expensive networks due to the spurious squaring of the matrix. That is because convolution kernels, which are usually constant but sparse, have their fanout squared; and coordinate transformations, which are usually dense but variable, have an excessive number of new (high-order) interactions created when A is squared.

Of course, equation (31) is of the type which we know how to transform using reversed linear neurons. We obtain the modified objective

$$\hat{E}_{\text{coord}} = c \left[\sum_i (x'_i - \sum_j A_{ij} x_j - b_i) \sigma_i - \frac{1}{2} \sum_i \sigma_i^2 \right] \quad (32)$$

which doesn't square A . If A is constant then there is no order reduction, since both E_{coord} and \hat{E}_{coord} are second order, but there are fewer connections unless A is also dense.

An alternative objective, not using reversed neurons, is also available for convolutions and coordinate system transformations. The objective

$$\hat{\hat{E}}_{\text{coord}} = \frac{c}{2} \sum_{ij} A_{ij} (x'_i - b_i - x_j)^2 \quad (33)$$

is minimal with respect to x' when

$$x'_i = \sum_j A_{ij} x_j / \sum_j A_{ij} + b_i. \quad (34)$$

This type of dynamic normalization may be desirable, or if A is constant and already normalized then it does not hurt. Equation (33) also preserves any sparseness of A , and does not square the matrix.

3.2 Graph Matching and Quadratic Match

Consider the following objective for inexact graph-matching [Hopfield & Tank 1986], c.f. [von der Malsburg & Bienenstock 1986]:

$$\begin{aligned} E_{\text{graph}} = & -c_1 \sum_{\alpha\beta ij} G_{\alpha\beta} g_{ij} M_{\alpha i} M_{\beta j} \\ & + c_2 \sum_{\alpha} (\sum_i M_{\alpha i} - 1)^2 + c_2 \sum_i (\sum_{\alpha} M_{\alpha i} - 1)^2 \\ & + c_3 \sum_{\alpha i} M_{\alpha i} (1 - M_{\alpha i}) \\ & + \sum_{\alpha i} \int^{M_{\alpha i}} dx g^{-1}(x) \end{aligned} \quad (35)$$

where G and g are connection matrices (each entry is zero or one) for two graphs, and $M_{\alpha i}$ is one when node α of G maps to node i of g , and zero otherwise.

The problem may be generalized slightly to "quadratic matching" by replacing the $GgMM$ term with

$$\sum_{\alpha\beta ij} G_{\alpha\beta} g_{ij} A_{\alpha i} B_{\beta j} \quad (36)$$

and altering the other constraints to reflect the fact that $A \neq B$. What we have to say about graph matching will apply equally well to this generalization.

The $GgMM$ term is superficially the expensive one since it involves four sums. If each graph is constant, there are $\mathcal{O}(N)$ nodes in each graph, and both graphs have fanout f , then the number of monomial synapses is $\mathcal{O}(N^2 f^2)$. We can reduce this to $\mathcal{O}(N^2 f)$. Also if one of G or g is variable, with N nodes in the variable graph and m in the constant graph, as in the "Frameville" networks of [Mjolsness Gindi & Anandan 1989], and both graphs are represented densely, then the number of synapses is reduced from $\mathcal{O}(N^2 m f)$ to $\mathcal{O}(N^2 m + N m f)$. The reduction uses linear interneurons, both reversed and normal:

$$\begin{aligned} E_1 &= -c_1 \sum_{\alpha\beta ij} G_{\alpha\beta} g_{ij} M_{\alpha i} M_{\beta j} \\ &= -\frac{c_1}{2} \sum_{\beta i} \left[(\sum_{\alpha} G_{\alpha\beta} M_{\alpha i} + \sum_j g_{ij} M_{\beta j})^2 \right. \\ &\quad \left. - (\sum_{\alpha} G_{\alpha\beta} M_{\alpha i})^2 - (\sum_j g_{ij} M_{\beta j})^2 \right] \\ \hat{E}_1 &= -c_1 \left[\sum_{\alpha\beta i} G_{\alpha\beta} M_{\alpha i} \sigma_{\beta i} + \sum_{\beta ij} g_{ij} M_{\beta j} \sigma_{\beta i} \right. \\ &\quad \left. - \sum_{\alpha\beta i} G_{\alpha\beta} M_{\alpha i} \tau_{\beta i} - \sum_{\beta ij} g_{ij} M_{\beta j} \omega_{\beta i} \right. \\ &\quad \left. - \frac{1}{2} \sigma_{\beta i}^2 + \frac{1}{2} \tau_{\beta i}^2 + \frac{1}{2} \omega_{\beta i}^2 \right] \end{aligned} \quad (37)$$

E_1 and \hat{E}_1 are illustrated in Figure 2.




Figure 2: **Graph Matching Networks.** E_1 is a sum over the indices α , β , i , and j , which are connected by neurons (line segments) in the shape of a "rectangle". This objective can be transformed into \hat{E} , which is a sum of triangles, while preserving fixpoints. The triangles are obtained from the rectangle by introducing linear interneurons along a diagonal, as shown. Only three indices are summed over, resulting in a less costly network.

The reduced graph-matching network works in simulation, for five out of six small ($N = 10$ nodes) hand-designed graphs with low fanout (from 1.8 to 2.5). The sixth case is not solved by the original, untransformed network either. The parameters we used were

$$\begin{array}{cccc}
 N = 10 & c_1 = 1.0 & c_2 = 1.0 & c_3 = 0.5 \\
 g_0(M) = 20 & g_0(\sigma) = 1.0 & g_0(WTA) = 10.0 & r_\sigma = 1 \\
 \Delta t = .004 & \text{sweeps} = 1\ 000 & &
 \end{array}$$

and for the original network:

$$\begin{array}{cccc}
 N = 10 & c_1 = 1.0 & c_2 = 1.0 & c_3 = 0.5 \\
 g_0(M) = 20 & g_0(\sigma) = 1.0 & g_0(WTA) = 10.0 & \Delta t = .004 \\
 \text{sweeps} = 1\ 000 & & &
 \end{array}$$

Here $g_0(M)$ is the gain $g'(0)$ of the transfer function $g(M)$ for M , and similarly $g_0(\sigma)$ is the gain for the linear neurons that were introduced through transforming E_1 , namely σ , τ , and ω . Also $g_0(WTA)$ is the gain for the infinitely fast linear neurons which were used in both reduced and control experiments to implement the WTA constraints; this parameter effectively multiplies c_1 . *sweeps* is the number of iterations of the forward Euler method used in simulating the continuous update equations. Each iteration advanced the time coordinate by Δt .

There is only a little parameter-tuning involved here, concentrated on r_σ , Δt , and *sweeps*. The product $\Delta t \times \max(r_\sigma, r_M = 1)$ should be held fixed to maintain constant resolution in the discrete simulation of continuous update equations. But holding Δt and the other parameters fixed at the quoted values, the rate parameter r_σ can be varied from unity to 100 without altering the network convergence time, measured in sweeps, by more than 30% or so; network performance remains the same in that the same 5 out of 6 graphs are correctly matched. This would suggest, and other experiments confirm, that for low r there is some room for increasing Δt and decreasing sweeps ($r = 10$, $\Delta t = .016$, and *sweeps* = 300 respectively); this saves time whether time is measured in simulation sweeps or in circuit time constants.

3.3 Sorting

Sorting may be described in a manner very similar to graph matching. One requires a permutation matrix M which sorts the inputs into increasing order, so all terms in the objective remain the same except for $GgMM$. The objective becomes

$$\begin{aligned}
 E_{\text{sort}} = & -c_1 \sum_{ij} M_{ij} x_i y_j \\
 & + c_2 \sum_i (\sum_i M_{ij} - 1)^2 + c_2 \sum_j (\sum_j M_{ij} - 1)^2 \\
 & + c_3 \sum_{ij} M_{ij} (1 - M_{ij}) \\
 & + \sum_{ij} \int^{M_{ij}} dx g^{-1}(x).
 \end{aligned} \tag{38}$$

Here x_i are a set of input numbers to be sorted, and y_j are a constant set of numbers in increasing order, e.g. $y_j = j$. M will become that permutation matrix which maximizes the inner product of x and y , i.e. maps the largest x_i to the largest y_j , the next largest x_i to the next largest y_j , and so on.

After using the winner-take-all reduction on the row and column constraints of M , this network has $\mathcal{O}(N^2)$ connections and neurons. One cannot do better without reducing the number of match neurons M . But M is sparse at any acceptable answer, so it may be possible to keep it sparse throughout the time the network is running by using a different encoding of the matrix. For example, one might encode indices i, j , or both in a binary representation as would be done in an ordinary computer program. The resulting objectives generally still have $\mathcal{O}(N^2)$ connections (monomial interactions), but a well-chosen matrix encoding, supplemented by suitable reversed neurons, can drastically reduce the number of connections.

In Appendix B it is shown that any permutation matrix M of size $L^2 \times L^2$ can be represented in the following form:

$$M_{i_1 i_2, j_1 j_2} = \sum_{k_1 k_2} A_{i_1 i_2, k_1}^{(1)} A_{i_2, k_1 k_2}^{(2)} \tilde{A}_{j_1 j_2, k_1}^{(1)} \tilde{A}_{j_2, k_1 k_2}^{(2)} \tag{39}$$

where $i \in \{1, \dots, N = L^2\}$, $i_k \in \{1, \dots, \sqrt{N} = L\}$, and $i = i_1 L + i_2 \equiv (i_1, i_2)$, and where two constraints apply to each nonsquare matrix:

$$\begin{aligned}
 \sum_{i_1} A_{i_1 i_2, j_1}^{(1)} = 1 \quad & \left| \quad \sum_{i_2} A_{i_2, j_1 j_2}^{(2)} = 1 \quad \right\| \quad \left. \sum_{i_1} \tilde{A}_{i_1 i_2, j_1}^{(1)} = 1 \quad \right| \quad \sum_{i_2} \tilde{A}_{i_2, j_1 j_2}^{(2)} = 1 \\
 \sum_{j_1} A_{i_1 i_2, j_1}^{(1)} = 1 \quad & \left| \quad \sum_{j_2} A_{i_2, j_1 j_2}^{(2)} = 1 \quad \right\| \quad \left. \sum_{j_1} \tilde{A}_{i_1 i_2, j_1}^{(1)} = 1 \quad \right| \quad \sum_{j_2} \tilde{A}_{i_2, j_1 j_2}^{(2)} = 1
 \end{aligned} \tag{40}$$

The matrix form (39) contains only $4N^{3/2}$ variables, and is our proposed encoding of M .

As explained in Appendix C, equation (39) is a coarse version of another expression for M which contains $\mathcal{O}(N \log N)$ variables. That expression codes index pairs (i, j) using the ‘‘Butterfly’’ connection topology that arises in the Fast Fourier Transform (FFT) and in many other parallel algorithms. The advantage of the Butterfly is that it allows one to make a gradual transition from one space (e.g. index i) to another (index j). There has been little success in transforming objectives based on the much less gradual binary or base- b encoding of i and M_{ij} .² An example similar to equation (39) is the base \sqrt{N} code obtained by listing all N links in the permutation matrix, indexed by k , and encoding their starting and ending locations as $i = (i_1, i_2)$ and $j = (j_1, j_2)$. Then

$$M_{i_1 i_2, j_1 j_2} = \sum_k A_{i_1, k}^{(1)} A_{i_2, k}^{(2)} \tilde{A}_{j_1, k}^{(1)} \tilde{A}_{j_2, k}^{(2)} \tag{41}$$

²A partial exception is the load-balancing network of Fox and Furmanky [Fox & Furmanky 1988], in which the crucial ‘‘histogram’’ may be understood as a set of reversed linear interneurons which simplify their load-balancing objective. But the result is a virtual neural net, not a statically connected circuit.

subject to obvious constraints on the A 's.

Since any permutation matrix can be expressed using equation (39), any (approximately quadratic) local minimum of $E_{\text{sort}}(M)$ (equation (38)), for which M is sufficiently close to being a permutation matrix, should be a local minimum of $E_{\text{sort}}(A^{(1)}, A^{(2)}, \tilde{A}^{(1)}, \tilde{A}^{(2)})$; but there may be local minima with respect to A and \tilde{A} which would be unstable in the larger M space. Thus, making the substitution (39) into equation (38) may expand the set of fixed points, or alter it entirely if the original objective is not yet tuned to produce permutation matrices. By contrast, the objective function transformations used heretofore have exactly preserved the set of fixed points. Nevertheless the transformation is worth trying out to decrease cost.

The problem now is to reduce the number of monomial interactions to $\mathcal{O}(N^{3/2})$. This is easy for quadratic penalty terms corresponding to the constraints (40), which consist of $8N$ winner-take-all constraints each involving $N^{1/2}$ variables. The remaining interaction $-c \sum Mxy$ can be reduced in two stages: substitute $M = A\tilde{A}^T$ with no reduction in connection costs; then substitute the $\mathcal{O}(N^{3/2})$ forms for A and \tilde{A} .

Thus we may replace $E_1 = -c_1 \sum_{ij} M_{ij} x_i y_j$ with

$$\begin{aligned} E_1(A, \tilde{A}) &= -c_1 \sum_{ijk} A_{ik} \tilde{A}_{jk} x_i y_j \\ &= -\frac{c_1}{2} \left[\sum_k (\sum_i A_{ik} x_i + \sum_j \tilde{A}_{jk} y_j)^2 \right. \\ &\quad \left. - \sum_k (\sum_i A_{ik} x_i)^2 - \sum_k (\sum_j \tilde{A}_{jk} y_j)^2 \right] \end{aligned} \quad (42)$$

$$\begin{aligned} \hat{E}_1(A, \tilde{A}, a, b, \tilde{b}) &= -c_1 \left[\sum_k (a_k - b_k) \sum_i A_{ik} x_i + \sum_k (a_k - \tilde{b}_k) \sum_j \tilde{A}_{jk} y_j \right. \\ &\quad \left. - \frac{1}{2} \sum_k a_k^2 + \frac{1}{2} \sum_k b_k^2 + \frac{1}{2} \sum_k \tilde{b}_k^2 \right] \end{aligned}$$

which may be interpreted as four interacting sorting problems, with linear interneurons a interpolating between x and y and with reversed neurons b and \tilde{b} cancelling echos as in equation (15). So far there is no reduction in number of neurons or connections.

If we substitute the special forms for A and \tilde{A} , we find

$$\begin{aligned} \hat{E}_1 &= -c_1 \left[\sum_{i_1 i_2} \sum_{k_1 k_2} x_i A_{i_1 k_1}^{(1)} A_{i_2 k_2}^{(2)} (a_k - b_k) + \sum_{j_1 j_2} \sum_{k_1 k_2} y_j \tilde{A}_{j_1 k_1}^{(1)} \tilde{A}_{j_2 k_2}^{(1)} (a_k - \tilde{b}_k) \right. \\ &\quad \left. - \frac{1}{2} \sum_k a_k^2 + \frac{1}{2} \sum_k b_k^2 + \frac{1}{2} \sum_k \tilde{b}_k^2 \right] \\ &= -c_1 \left[\frac{1}{2} \sum_{i_2 k_1} [(\sum_{i_1} A_{i_1 i_2, k_1}^{(1)} x_i + \sum_{k_2} A_{i_2, k_1 k_2}^{(2)} (a_k - b_k))^2 \right. \\ &\quad \left. - (\sum_{i_1} A_{i_1 i_2, k_1}^{(1)} x_i)^2 - (\sum_{k_2} A_{i_2, k_1 k_2}^{(2)} (a_k - b_k))^2 \right] \\ &\quad + \frac{1}{2} \sum_{j_2 k_1} [(\sum_{j_1} \tilde{A}_{j_1 j_2, k_1}^{(1)} y_j + \sum_{k_2} \tilde{A}_{j_2, k_1 k_2}^{(2)} (a_k - \tilde{b}_k))^2 \\ &\quad \left. - (\sum_{j_1} \tilde{A}_{j_1 j_2, k_1}^{(1)} y_j)^2 - (\sum_{k_2} \tilde{A}_{j_2, k_1 k_2}^{(2)} (a_k - \tilde{b}_k))^2 \right] \\ &\quad \left. - \frac{1}{2} \sum_k a_k^2 + \frac{1}{2} \sum_k b_k^2 + \frac{1}{2} \sum_k \tilde{b}_k^2 \right] \end{aligned} \quad (43)$$

and finally

$$\begin{aligned}
\hat{E}_1 = -c_1 & \left[\sum_{i k_1} A_{i_1 i_2, k_1}^{(1)} x_i (\sigma_{i_2 k_1} - \tau_{i_2 k_1}) + \sum_{i_2 k} A_{i_2, k_1 k_2}^{(2)} (a_k - b_k) (\sigma_{i_2 k_1} - \omega_{i_2 k_1}) \right. \\
& - \frac{1}{2} \sum_{i_2 k_1} \sigma_{i_2 k_1}^2 + \frac{1}{2} \sum_{i_2 k_1} \tau_{i_2 k_1}^2 + \frac{1}{2} \sum_{i_2 k_1} \omega_{i_2 k_1}^2 \\
& + \sum_{j k_1} \tilde{A}_{j_1 j_2, k_1}^{(1)} y_j (\bar{\sigma}_{j_2 k_1} - \bar{\tau}_{j_2 k_1}) + \sum_{j_2 k} \tilde{A}_{j_2, k_1 k_2}^{(2)} (a_k - \bar{b}_k) (\bar{\sigma}_{j_2 k_1} - \bar{\omega}_{j_2 k_1}) \\
& - \frac{1}{2} \sum_{j_2 k_1} \bar{\sigma}_{j_2 k_1}^2 + \frac{1}{2} \sum_{j_2 k_1} \bar{\tau}_{j_2 k_1}^2 + \frac{1}{2} \sum_{j_2 k_1} \bar{\omega}_{j_2 k_1}^2 \\
& \left. - \frac{1}{2} \sum_k a_k^2 + \frac{1}{2} \sum_k b_k^2 + \frac{1}{2} \sum_k \bar{b}_k^2 \right] \quad (44)
\end{aligned}$$

which has $\mathcal{O}(N^{3/2})$ neurons and connections.

In Appendix C we show how to extend this result to a series of successively cheaper approximations of the original sorting network, down to $\mathcal{O}(N \log N)$ neurons and connections.

3.4 Sorting: Experiments

The $\mathcal{O}(N^{3/2})$ sorting network only sorts in an approximate way. The reversed neurons work correctly at finite r , which is nontrivial since they are connected to each other, but the encoding scheme is prone to trapping by local minima. We used the following parameter values for the $\mathcal{O}(N^{3/2})$ sorting network:

$N = 16$	$c_1 = 0.6$	$c_2 = 6.0$	$c_3 = 0$
$g_0(A) = 20$	$r_A = 1$	$r_\sigma = 3$	$\Delta t = .01$
<hr style="border: 0.5px solid black;"/>			
$N = 25$	$c_1 = .44$	$c_2 = 6.0$	$c_3 = 0$
$g_0(A) = 20$	$r_A = 1$	$r_\sigma = 3$	$\Delta t = .01$
<hr style="border: 0.5px solid black;"/>			
$sweeps = 20\ 000$			

and for the $\mathcal{O}(N^2)$ network:

$N = 16$	$c_1 = 0.6$	$c_2 = 6.0$	$c_3 = 0$
$g_0(A) = 20$	$\Delta t = .01$	$sweeps = 5\ 000$	
<hr style="border: 0.5px solid black;"/>			
$N = 25$	$c_1 = .44$	$c_2 = 6.0$	$c_3 = 0$
$g_0(A) = 20$	$\Delta t = .01$	$sweeps = 5\ 000$	

As in the graph-matching example, most of the parameter-tuning was concentrated on r_σ , Δt , and $sweeps$. Here, r_σ is the rate parameter appearing in the update equation (7), and it applies to neurons $a, b, \bar{b}, \sigma, \bar{\sigma}, \tau, \bar{\tau}, \omega, \bar{\omega}$. Likewise r_A applies to the update equations for A and \bar{A} . As in equation (38), c_1 multiplies the strength of the permuted inner product of \vec{x} and \vec{y} in the objective. Also c_2 is the strength of the syntax constraints, and c_3 is the strength of a term penalizing intermediate neuron values. $g_0(A)$ is the gain $g'(0)$ of the transfer function $g(A)$ for A and \bar{A} , which obeyed steepest-descent dynamics. The constant y values are $y_1 = -(N-1)/2$, $y_2 = -(N-3)/2, \dots, y_{N-1} = (N-3)/2$, $y_N = (N-1)/2$. $sweeps$ is the number of iterations of the forward Euler method used in simulating the continuous update equations.

For input size $N = 16$ we find an average placement error of 1.4 out of 16 possible output places. Eight would be random. For $N = 25$, which is the first size (with integral \sqrt{N}) for which there are fewer neurons in the asymptotically smaller network, the average placement error is 1.7 out of 25 places. The errors can be characterized by a histogram showing the frequency with which placement errors of different sizes occur. (The size of a placement error is the difference, in the permuted output vector, between the desired and actual positions of an element.) Histograms for $N = 16$ and $N = 25$ are presented in figure 3 and they show




Figure 3: Histogram of placement errors. Sorting network with butterfly encoding, $\mathcal{O}(N^{3/2})$ connections. (a) Size $N = 16$. Average and standard deviation (upper or lower half of an error bar) for 62 runs. (b) Size $N = 25$. Average and standard deviation for 39 runs.

that small mistakes are far more likely than large ones, and that the frequency falls off as roughly the -2.1 (respectively -2.0) power of the size of the placement error, with correlation $r = .90$ (.93). In addition, 17.7% (respectively 21%) of the experimental runs failed to meet our convergence criterion, which was that exactly one element in each row and column of the computed matrix M must have a value greater than .5.

Iterating the sort can improve the score marginally, but not to the perfect sorts achievable with the $\mathcal{O}(N^2)$ network at these sizes.

4 Discussion and Conclusion

There are algebraic transformations of objective functions which not only preserve the set of fixpoints of the resulting analog neural network, but alter the number and nature of interactions requiring physical connections between different neurons. These transformations generally require the network to find a saddle point rather than a minimum of the objective, but that can be arranged. A set of such transformations was derived, together with their conditions of validity.

Several design examples were given, along with experimental results to show convergence with reasonable parameter values. A reduced network for sorting converged to approximately correct answers despite the use of an illegitimate transformation which introduced spurious fixpoints. This design also involved legitimate but repeated, interacting product-reduction transformations. A reduced network for quadratic matching was simulated without difficulty. Reduced-cost designs were presented for convolution and linear coordinate transformation.

Although much research now focuses on expressing new computational problems using objective functions and then deriving neural networks which solve them, we would like to suggest that when such efforts are successful there may be an additional advantage to be obtained. If the solution can be regarded as a novel algebraic transformation of an unremarkable objective function, then the transformation may also be immediately applicable to other objective functions and problems.

Acknowledgements

We wish to acknowledge discussions with P. Anandan, Christian Darken, Wojtek Furmansky, Gene Gindi, John Platt, and Joachim Utans.

Appendices

A Reducing $YF(X)$

The problem is to use equation 19 to reduce expressions of the form $YF(X)$. We will derive two versions: the first, from stopping after the first step in the derivation of (19); the second, from carrying the transformation all the way through.

Now

$$\begin{aligned}
 Y \int^X du f(u) &= \int^X du \left[\int^Y dv g(u, v) \right]^{-1} (u) \\
 \Rightarrow Y f(X) &= \left[\int^Y dv g(u, v) \right]^{-1} (X) \\
 \Rightarrow \int^Y dv g(X, v) &= f^{-1}(X/Y) \\
 \Rightarrow g(u, v) &= (d/dv) f^{-1}(u/v) = -(u/v^2) (f^{-1})' (u/v).
 \end{aligned} \tag{45}$$

To use the first step in the derivation of equation (19), we must evaluate

$$\int^\sigma du \int^Y dv g(u, v) = \int^\sigma du f^{-1}(u/v) = Y \int^{\sigma/Y} du f^{-1}(u)$$

so

$$Y \int^X f(u) du \rightarrow XY\sigma - Y \int^\sigma du f^{-1}(u) \tag{46}$$

since σ may be rescaled by Y . This shows that both sides of the transformation (16) may be multiplied by any expression Y .

To carry through transformation (19) we must calculate

$$\begin{aligned}
 \int^\tau dv \left[\int^\sigma du g(u, v) \right]^{-1} (v) &= \int^\tau dv \left[\int^\sigma du (-u/v^2) (f^{-1})' (u/v) \right]^{-1} (v) \\
 &= \int^\tau dv \left[- \int^{\sigma/v} du u (f^{-1})' (u) \right]^{-1} (v) \\
 &= \int^\tau dv \left[- \int^{f^{-1}(\sigma/v)} dz z f'(z) \right]^{-1} (v) \\
 &= \int^\tau dv \left[-F(f^{-1}(\sigma/v)) \right]^{-1} (v) \\
 &= \int^\tau dv \left[\sigma / f(F^{-1}(-v)) \right] = -\sigma \int^{F^{-1}(-\tau)} dw \\
 &= -\sigma F^{-1}(-\tau)
 \end{aligned} \tag{47}$$

from which we deduce that

$$YF(X) \rightarrow -X\sigma + Y\tau + \sigma F^{-1}(\tau). \tag{48}$$

(τ and σ have been rescaled by -1). The algebra can be checked by optimizing with respect to σ . Note that the derivations of (46) and (48) assume that $f = F'$ is invertible and that $f^{-1} = (F')^{-1}$ is differentiable (i.e. $((F')^{-1})'$ exists).

B Butterfly Networks

B.1 Back-to-back Butterflies

By a recursive induction argument, any permutation matrix of size $N = 2^n$ (n an integer) can be expressed by setting switches in two back-to-back butterfly networks independently, as shown in Figure 4. If we label the corresponding connection matrices A and \bar{A} , then the entire network represents a permutation matrix M

in the form of a matrix product $M_{ij} = (A\tilde{A}^T)_{ij} = \sum_k A_{ik}\tilde{A}_{jk}$, with A and \tilde{A} being of the special “butterfly” form. Butterfly networks are best analyzed by introducing binary notation for all indices, e.g.

$$\begin{aligned} i &\rightarrow (p_1, \dots, p_n) \sim p_1 \dots p_n, \\ j &\rightarrow (q_1, \dots, q_n) \sim q_1 \dots q_n. \end{aligned} \quad (49)$$

In this notation, the outer column of switches in the A butterfly has the form

$$B_{p_1 \dots p_n, q_1}^{(1)} \in [0, 1] \quad (50)$$

with constraints

$$\sum_{p_1} B_{p_1 \dots p_n, q_1}^{(1)} = 1 \text{ and } \sum_{q_1} B_{p_1 \dots p_n, q_1}^{(1)} = 1 \quad (51)$$

as illustrated by the “butterfly” (\bowtie) highlighted in Figure 4. Likewise the l 'th column of switches has the form

$$B_{p_1 \dots p_n, q_1 \dots q_l}^{(l)} \in [0, 1] \text{ with } \sum_{p_l} B_{p_1 \dots p_n, q_1 \dots q_l}^{(l)} = 1 = \sum_{q_l} B_{p_1 \dots p_n, q_1 \dots q_l}^{(l)}, \quad (52)$$

for $1 \leq l \leq n$. With constraints, each of the $\log N$ layers contains $N/2$ bits, which is one reason that \tilde{A} is also needed to specify an entire permutation matrix.

The entire permutation matrix is obtained by finding all the possible paths through the network from i to j , of which there is only one since each stage irrevocably decides one bit of j . This is equivalent to a conjunction of switch settings:

$$A_{ij} = \prod_{l=1}^n B_{p_l \dots p_n, q_1 \dots q_l}^{(l)}. \quad (53)$$

It is easy to check that this is a permutation matrix, using the constraints on $A^{(l)}$. (In what follows the terms of a product \prod do not commute because they contain summations \sum that extend over subsequent terms in the product).

$$\begin{aligned} \sum_j A_{ij} &= \sum_{q_1 \dots q_n} A_{i, q_1 \dots q_n} \\ &= \prod_{l=1}^n \left(\sum_{q_l} B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \right) \\ &= \left[\prod_{l=1}^{n-1} \left(\sum_{q_l} B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \right) \right] \left(\sum_{q_n} B_{p_n, q_1 \dots q_n}^{(n)} = 1 \right) \\ &= \prod_{l=1}^{n-1} \left(\sum_{q_l} B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \right) \\ &= 1, \text{ by induction on } n. \end{aligned} \quad (54)$$

Likewise

$$\begin{aligned} \sum_i A_{ij} &= \sum_{p_1 \dots p_n} A_{i, p_1 \dots p_n} \\ &= \prod_{l=n}^1 \left(\sum_{p_l} B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \right) \\ &= \left[\prod_{l=n}^2 \left(\sum_{p_l} B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \right) \right] \left(\sum_{p_1} B_{p_1 \dots p_n, q_1}^{(1)} = 1 \right) \\ &= \prod_{l=n}^2 \left(\sum_{p_l} B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \right) \\ &= 1, \text{ by induction.} \end{aligned} \quad (55)$$

B.2 Coarse Butterflies

A less restrictive form for A may be derived as follows. Let k be roughly $n/2$. Then

$$i_1 \equiv p_1 \dots p_k, \quad i_2 \equiv p_{k+1} \dots p_n, \quad j_1 \equiv q_1 \dots q_k, \quad j_2 \equiv q_{k+1} \dots q_n;$$




Figure 4: **Butterfly switching networks.** Any permutation of $N = 2^n$ elements can be represented by appropriately setting the switches in a pair of back-to-back butterfly switching networks, as can be shown recursively by induction on n . Highlighted: one 2×2 permutation matrix or “butterfly”, and one path through the the entire switching network.

from equation (53),

$$\begin{aligned}
 A_{ij} &= \prod_{l=1}^n B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \\
 &= \left(\prod_{l=1}^k B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \right) \left(\prod_{l=k+1}^n B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \right) \\
 &= \left(\prod_{l=1}^k B_{p_l \dots p_n, q_1 \dots q_k}^{(l)} \right) \left(\prod_{l=k+1}^n B_{p_{k+1} \dots p_n, q_1 \dots q_n}^{(l)} \right) \\
 &= A_{i_1 i_2, j_1}^{(1)} A_{i_2 j_1 j_2}^{(2)}
 \end{aligned} \tag{56}$$

and as in equation (54) one can derive the constraints

$$\begin{aligned}
 \sum_{j_1} A_{i_1 i_2, j_1}^{(1)} &= \prod_{l=1}^k \left(\sum_{q_l} B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \right) \\
 &= \left[\prod_{l=1}^{k-1} \left(\sum_{q_l} B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \right) \right] \left(\sum_{q_k} B_{p_k \dots p_n, q_1 \dots q_k}^{(k)} = 1 \right) \\
 &= \prod_{l=1}^{k-1} \left(\sum_{q_l} B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \right) \\
 &= 1, \text{ by induction on } k.
 \end{aligned} \tag{57}$$

Likewise

$$\begin{aligned}
 \sum_{i_1} A_{i_1 i_2, j_1}^{(1)} &= \prod_{l=k}^1 \left(\sum_{p_l} B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \right) \\
 &= \left[\prod_{l=k}^2 \left(\sum_{p_l} B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \right) \right] \left(\sum_{p_1} B_{p_1 \dots p_n, q_1}^{(1)} = 1 \right) \\
 &= \prod_{l=k}^2 \left(\sum_{p_l} B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \right) \\
 &= 1, \text{ by induction.}
 \end{aligned} \tag{58}$$

The constraints on $A^{(2)}$ and \bar{A} are similar. Thus the constraints on $B^{(l)}$ imply the less restrictive constraints on $A^{(1)}$ and $A^{(2)}$, which we then adopt as the only constraints operating on A .

This completes the proof that any permutation matrix M can be represented by $\mathcal{O}(N^{3/2})$ variables in the form of equation (39) with constraints as in equation (40).

C Full Butterfly Neural Nets

The arguments of Appendix B can be generalized to yield a series of reduced objectives interpolating between $\mathcal{O}(N^{3/2})$ and $\mathcal{O}(N \log N)$ variables, each having about the same number of connections as variables. In Appendix B the idea was to express each index i in base $\sqrt{N} = 2^{n/2=k}$, by dividing the indices $p_1 \dots p_n$ of i 's binary expansion in two groups. We may instead divide the n binary indices into m groups of size k , with $n = km$, deriving the base 2^k expansion $i = i_1 \dots i_m$. (We have dealt with the special cases $k = n$ and $k \approx n/2$.) Then

$$\begin{aligned} A_{ij} &= \prod_{l=1}^n B_{p_l \dots p_n, q_1 \dots q_l}^{(l)} \\ &= \prod_{l=1}^m \left(\prod_{r=1}^k B_{p_{l,r} \dots p_n, q_1 \dots q_{l,r}}^{(l,r)} \right) \\ &= \prod_{l=1}^m A_{i_l \dots i_m, j_1 \dots j_l}^{(l)} \end{aligned} \quad (59)$$

as in equation (56). The constraints on $A^{(l)}$ are, as usual,

$$\sum_{i_l} A_{i_l \dots i_m, j_1 \dots j_l}^{(l)} = 1 = \sum_{j_l} A_{i_l \dots i_m, j_1 \dots j_l}^{(l)} \quad (60)$$

which are generally less restrictive than the original constraints on the butterfly switches B .

From equation (42), our problem is to reduce

$$E_1(A) = -c_1 \sum_{ij} A_{ij} x_i e_j \quad (61)$$

(where e_j is any expression) upon substituting equation (59).

If we take $e_j = e_{j_1 \dots j_m}^{(m)}$, $A_{ij} = C_{i_1 \dots i_m, j_1 \dots j_m}^{(m)}$, and $E_1(A) = E_1^{(m)}(C^{(m)})$, then we may use induction on a to reduce

$$E_1^{(a)}(C^{(a)}) = -c_1 \sum_{\substack{i_1 \dots i_m \\ j_1 \dots j_a}} C_{i_1 \dots i_m, j_1 \dots j_a}^{(a)} x_{i_1 \dots i_m} e_{i_{a+1} \dots i_m, j_1 \dots j_a}^{(a)} \quad (a \geq 2), \quad (62)$$

where

$$C_{i_1 \dots i_m, j_1 \dots j_a}^{(a)} \equiv \prod_{l=1}^a A_{i_l \dots i_m, j_1 \dots j_l}^{(l)}. \quad (63)$$

Then

$$C_{i_1 \dots i_m, j_1 \dots j_a}^{(a)} = C_{i_1 \dots i_m, j_1 \dots j_{a-1}}^{(a-1)} A_{i_a \dots i_m, j_1 \dots j_a}^{(a)}$$

whence

$$\begin{aligned}
E_1^{(a-1)} &= -\frac{c_1}{2} \sum_{\substack{i_a \dots i_m \\ j_1 \dots j_{a-1}}} \left[\left(\sum_{i_1 \dots i_{a-1}} C_{i_1 \dots i_m, j_1 \dots j_{a-1}}^{(a-1)} x_{i_1 \dots i_m} \right. \right. \\
&\quad \left. \left. + \sum_{j_a} A_{i_a \dots i_m, j_1 \dots j_a}^{(a)} e_{i_{a+1} \dots i_m, j_1 \dots j_a}^{(a)} \right)^2 \right. \\
&\quad \left. - \left(\sum_{i_1 \dots i_{a-1}} C_{i_1 \dots i_m, j_1 \dots j_{a-1}}^{(a-1)} x_{i_1 \dots i_m} \right)^2 - \left(\sum_{j_a} A_{i_a \dots i_m, j_1 \dots j_a}^{(a)} e_{i_{a+1} \dots i_m, j_1 \dots j_a}^{(a)} \right)^2 \right] \\
&\rightarrow -c_1 \left[\sum_{i_1 \dots i_m} \sum_{j_1 \dots j_{a-1}} C_{i_1 \dots i_m, j_1 \dots j_{a-1}}^{(a-1)} x_{i_1 \dots i_m} \left(\sigma_{i_a \dots i_m, j_1 \dots j_{a-1}}^{(a-1)} - \tau_{i_a \dots i_m, j_1 \dots j_{a-1}}^{(a-1)} \right) \right. \\
&\quad \left. \sum_{i_a \dots i_m} \sum_{j_1 \dots j_a} A_{i_a \dots i_m, j_1 \dots j_a}^{(a)} e_{i_{a+1} \dots i_m, j_1 \dots j_a}^{(a)} \left(\sigma_{i_a \dots i_m, j_1 \dots j_{a-1}}^{(a-1)} - \omega_{i_a \dots i_m, j_1 \dots j_{a-1}}^{(a-1)} \right) \right. \\
&\quad \left. + \frac{1}{2} \sum_{i_a \dots i_m} \sum_{j_1 \dots j_{a-1}} \left[- \left(\sigma_{i_a \dots i_m, j_1 \dots j_{a-1}}^{(a-1)} \right)^2 + \left(\tau_{i_a \dots i_m, j_1 \dots j_{a-1}}^{(a-1)} \right)^2 + \left(\omega_{i_a \dots i_m, j_1 \dots j_{a-1}}^{(a-1)} \right)^2 \right] \right]. \tag{64}
\end{aligned}$$

Note that, upon identifying

$$\sigma^{(a-1)} - \tau^{(a-1)} = e^{(a-1)},$$

we have an induction step $a \rightarrow a - 1, \forall a \geq 2$, which decreases the number of neurons and connections in the

network. By induction on a , one may reduce equation (42) to:

$$\begin{aligned}
E_{\text{sort}} = & -c_1 \left[\sum_{i_1 \dots i_m, j_1} A_{i_1 \dots i_m, j_1}^{(1)} x_{i_1 \dots i_m} \left(\sigma_{i_2 \dots i_m, j_1}^{(1)} - \tau_{i_2 \dots i_m, j_1}^{(1)} \right) \right. \\
& + \sum_{a=2}^{m-1} \sum_{i_a \dots i_m} A_{i_a \dots i_m, j_1 \dots j_a}^{(a)} \left(\sigma_{i_a \dots i_m, j_1 \dots j_a}^{(a-1)} - \omega_{i_a \dots i_m, j_1 \dots j_{a-1}}^{(a-1)} \right) \\
& \quad \times \left(\sigma_{i_{a+1} \dots i_m, j_1 \dots j_a}^{(a)} - \tau_{i_{a+1} \dots i_m, j_1 \dots j_{a-1}}^{(a)} \right) \\
& + \sum_{i_m, j_1 \dots j_m} A_{i_m, j_1 \dots j_m}^{(m)} \left(\sigma_{i_m, j_1 \dots j_{m-1}}^{(m-1)} - \omega_{i_m, j_1 \dots j_{m-1}}^{(m-1)} \right) (a_{j_1 \dots j_m} - b_{j_1 \dots j_m}) \\
& + \frac{1}{2} \sum_{a=1}^{m-1} \sum_{i_{a+1} \dots i_m} \left[- \left(\sigma_{i_{a+1} \dots i_m, j_1 \dots j_a}^{(a)} \right)^2 \right. \\
& \quad \left. + \left(\tau_{i_{a+1} \dots i_m, j_1 \dots j_a}^{(a)} \right)^2 + \left(\omega_{i_{a+1} \dots i_m, j_1 \dots j_a}^{(a)} \right)^2 \right] \\
& + \left\{ \begin{array}{ll} x \rightarrow y & \sigma \rightarrow \tilde{\sigma} \\ \text{same, with } B \rightarrow \tilde{B} & \tau \rightarrow \tilde{\tau} \\ b \rightarrow \tilde{b} & \omega \rightarrow \tilde{\omega} \\ a \rightarrow a & \end{array} \right\} \\
& + \frac{1}{2} \sum_{i_1 \dots i_m} \left[- a_{i_1 \dots i_m}^2 + b_{i_1 \dots i_m}^2 + \tilde{b}_{i_1 \dots i_m}^2 \right].
\end{aligned} \tag{65}$$

The number of variables and connections for this objective is $\mathcal{O}(mN^{1+1/m})$, $1 \leq m \leq n$, which takes on values $N^2, 2N^{3/2}, 3N^{4/3}, \dots, N \log N$.

References

- [Arrow 1958] K. J. Arrow, L. Hurwicz, and H. Uzawa, editors. *Studies in Linear and Nonlinear Programming*. Stanford University Press, 1958.
- [Feldman 1982] J. A. Feldman. Dynamic connections in neural networks. *Biological Cybernetics*, 46:27-39, 1982.
- [Fox & Furmanský 1988] G. C. Fox and W. Furmanský. Load balancing loosely synchronous problems with a neural network. Technical Report *C³P363B*, California Institute of Technology, February 1988.
- [Gindi Gmitro & Parthasarathy 1987] Gene Gindi, Arthur Gmitro, and Kannan Parthasarathy. Winner-take-all networks and associative memory: Analysis and optical realization. In *Proc. of First International Conference on Neural Networks*, volume vol. III, pages 607-614. IEEE, 1987.
- [Grossberg 1988] Stephen Grossberg. Nonlinear neural networks: Principles, mechanisms, and architectures. *Neural Networks*, 1:17-61, 1988.
- [Hopfield 1984] J. J. Hopfield. Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences USA*, vol. 81:3088-3092, May 1984.
- [Hopfield & Tank 1985] J. J. Hopfield and D. W. Tank. 'Neural' computation of decisions in optimization problems. *Biological Cybernetics*, vol. 52:141-152, 1985.
- [Hopfield & Tank 1986] J. J. Hopfield and D. W. Tank. Collective computation with continuous variables. In *Disordered Systems and Biological Organization*, pages 155-170. Springer-Verlag, 1986.
- [Koch Marroquin & Yuille 1986] Christof Koch, Jose Marroquin, and Alan Yuille. Analog "neuronal" networks in early vision. *Proceedings of the National Academy of Sciences USA*, 83, June 1986.
- [Luenberger 1984] David G. Luenberger. *Linear and Nonlinear Programming*. Addison-Wesley, 1984.
- [Mjolsness Gindi & Anandan 1989] Eric Mjolsness, Gene Gindi, and P. Anandan. Optimization in model matching and perceptual organization. *Neural Computation*, 1989. To appear.
- [Moody 1989] John Moody. Optimal Architectures and Objective Functions for Associative Memory. Technical report, Yale University Department of Computer Science, March 1989. In preparation.
- [Mjolsness Sharp & Alpert 1989] Eric Mjolsness, David H. Sharp, and Bradley K. Alpert. Scaling, machine learning, and genetic neural nets. *Advances in Applied Mathematics*, 1989. To appear.
- [Platt & Barr 1987] John C. Platt and Alan H. Barr. Constrained differential optimization. In Dana Z. Anderson, editor, *Neural Information Processing Systems*. American Institute of Physics, 1987.
- [Platt & Barr 1988] John C. Platt and Alan H. Barr. Constraint methods for flexible models. *Computer Graphics*, 22(4), August 1988. Proceedings of SIGGRAPH '88.
- [Rumelhart Hinton & McClelland 1986] D. E. Rumelhart, G. E. Hinton, and J. L. McClelland. A general framework for parallel distributed processing. In *Parallel Distributed Processing*, pages 73-74. MIT Press, 1986.

- [Rumelhart Hinton & Williams 1986] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. In *Parallel Distributed Processing*, pages 318–362. MIT Press, 1986.
- [Sivilotti Mahowald & Mead 1987] Massimo A. Sivilotti, Michelle A. Mahowald, and Carver A. Mead. Real-time visual computations using analog CMOS processing arrays. In *Advanced Research in VLSI: Proceedings of the 1987 Stanford Conference*. MIT Press, 1987.
- [Tank & Hopfield 1986] David W. Tank and John J. Hopfield. Simple 'neural' optimization networks: An a/d converter, signal decision circuit, and a linear programming circuit. *IEEE Transactions on Circuits and Systems*, CAS-33, May 1986.
- [Utans et. al. 1989] Joachim Utans, Gene Gindi, Eric Mjolsness, and P. Anandan. Neural networks for object recognition within compositional hierarchies: Initial experiments. Technical Report Center for Systems Science Report No. 8903, Yale University Department of Electrical Engineering, February 1989.
- [von der Malsburg & Bienenstock 1986] Christoph von der Malsburg and Elie Bienenstock. Statistical coding and short-term synaptic plasticity: A scheme for knowledge representation in the brain. In *Disordered Systems and Biological Organization*, pages 247–252. Springer-Verlag, 1986.
- [von Neumann & Morgenstern 1953] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1953.