

Analysis of a Cooperative Stereo Algorithm

D. Marr*, G. Palm** and T. Poggio**

Abstract. Marr and Poggio (1976) recently described a cooperative algorithm that solves the correspondence problem for stereopsis. This article uses a probabilistic technique to analyze the convergence of that algorithm, and derives the conditions governing the stability of the solution state. The actual results of applying the algorithm to random-dot stereograms are compared with the probabilistic analysis. A satisfactory mathematical analysis of the asymptotic behaviour of the algorithm is possible for a suitable choice of the parameter values and loading rules, and again the actual performance of the algorithm under these conditions is compared with the theoretical predictions. Finally, some problems raised by the analysis of this type of “cooperative” algorithm are briefly discussed.

1. Introduction

The extraction of stereo-disparity information from two images depends upon establishing a correspondence between them. In a recent article, Marr and Poggio (1976) analyzed the nature of the correspondence computation and derived a cooperative algorithm that implements it. Although several examples were given of the performance of the algorithm on random-dot stereograms (Marr and Poggio 1976, Figs. 3–6), space did not permit a thorough analysis of the fixed points of the algorithm, or of its convergence. In this article, we shall examine these issues in detail.

1.1. Computational Structure of the Correspondence Problem

Marr and Poggio (1976) argued that the stereo problem may be reduced to that of matching two primitive

* Massachusetts Institute of Technology, Department of Psychology, Cambridge, MA, USA

** Max-Planck-Institut für Biologische Kybernetik, Tübingen, FRG

descriptions, one from each eye. They showed that the central problem is to find a correspondence between the left and right descriptions, that satisfies the two rules (p. 284 and Marr, 1974):

(R1) *Uniqueness.* Each item from each image may be assigned at most one disparity.




Fig. 1a—c. **a** shows the explicit structure of the two rules *R1* and *R2* for the case of a one-dimensional image, and it also represents the structure of a network for implementing the algorithm described by (1). Solid lines represent “inhibitory” interactions, and dotted lines represent “excitatory” ones. **b** Gives the local structure at each node of the network **a**. This algorithm may be extended to two-dimensional images, in which case each node in the corresponding network has the local structure shown in **c**. (Marr and Poggio, 1976, Fig. 2)




Fig. 2. The excitatory neighborhood (Fig. 1c) used in our implementation has a diameter of 5, and contains 13 cells. The central cell, marked by a square, receives at most 12 excitatory inputs from its neighbours




Fig. 3a—e. The total excitatory contribution for various configurations of “on” cells. The excitatory neighborhood (Fig. 2) is shown with open circles, except for the central cell which is indicated by a square because it makes no contribution to the total excitation. With a threshold of 4.0: **a** shows that a flat border will grow in the absence of inhibition, **b** exhibits the smallest stable configuration, **c** the sharpest stable convexity, and **d** and **e** show concavities that fill in

(R2) *Continuity.* Disparity varies smoothly almost everywhere.

By constructing an explicit geometrical representation of these two rules (Fig. 1c), they were able to derive a cooperative algorithm that implements them. If one thinks of Figure 1a as a network, with a cell at each node, the uniqueness rule *R1* means that only one cell is “on” along each vertical or horizontal line (the line of sight from the left and right eyes); and the continuity rule *R2* implies that solutions (its asymptotic states) tend to spread along the dotted diagonals (lines of constant disparity).

In order to implement these rules, each cell sends “inhibitory” connections to all other cells along the same vertical and horizontal lines, and excitatory connections along its diagonal. This gives the local network geometry shown in Figure 1b. For a two-dimensional image, the only change needed is to make the excitatory neighborhood two-dimensional, which gives the local geometry shown in Figure 1c.

Let $C_{x,y;d}^t$ denote the state at time t of the cell corresponding to coordinate (x, y) on the left retina, matching position $(x + d, y)$ on the right retina. Let $S(x,y,d)$ denote its excitatory neighborhood (the disc in Fig. 1c), and $O(x,y,d)$ its inhibitory neighborhood (the horizontal and vertical lines in Figure 1c). The algor-

ithm implemented by the network may be written (Marr and Poggio 1976, Equation (2))

$$C_{x,y;d}^{t+1} = \sigma \left\{ \sum_{x',y',d' \in S(x,y,d)} C_{x',y';d'}^t - \varepsilon \sum_{x',y',d' \in O(x,y,d)} C_{x',y';d'}^t + C_{x,y;d}^0 \right\} \quad (1)$$

where σ is a threshold function that takes values 0 or 1, and ε is an “inhibition” constant.

This article is concerned with the properties of the algorithms defined by (1) or, equivalently, with the behavior of the corresponding network (Fig. 1). The two inputs to the algorithm or network, from which the initial state of the network is determined, are usually two matrices whose entries consist of 0's and 1's. The second matrix is constructed from the first by x -translations of regions of it. As we shall discuss later the algorithm defined by (1) has some analogies with games like “life”.

The plan of the paper is as follows: Section 2 describes the loading rules, which determine the initial state from the input stereograms, and also defines the algorithm precisely. The relations between the fixed points of the algorithm and the states that satisfy the two conditions *R1* and *R2* are then discussed (Section 3). A probabilistic approach to the convergence of the algorithm is outlined in Section 4. Actual computer simulations of the algorithm are compared with the probabilistic analysis, and the range of parameter values that yield a “nice” convergence is discussed. Some special situations are also analyzed (Section 5). A suitable (and restrictive) choice of the parameter values in (1) allows a satisfactory mathematical analysis of the algorithm: Section 6 is devoted to such an approach. Finally, we briefly discuss the mathematical problems raised by the analysis of this type of “cooperative” algorithm.

2. The Algorithm

2.1. Loading Conditions

Let the positions on the left and right retinas be denoted by $L_{x,y}$ and $R_{x,y}$ respectively. These arrays take the values 0, indicating the absence of a feature, or 1, indicating the presence. The initial condition of the network, for stereogram L, R is given by

$$C_{x,y;d}^0 = L_{x,y} \cdot R_{x+d,y} \quad (2)$$

within the appropriate range d of disparity.

2.2. The Algorithm

The relation between states at times t and $t + 1$, is given by the recurrence relation (1), where σ is a sigmoid

function in general, and here is taken to be the threshold function

$$\sigma(u) = \begin{cases} 1 & \text{if } u \geq \theta, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

ε is a constant, known as the “inhibition constant”. The number of disparity layers d we shall denote by D , and we shall let M be the diameter of the excitatory neighborhood $S(x, y, d)$. In the example shown in Figure 2, $M=5$, and the total number of cells in an excitatory neighborhood is 13. The number less the cell itself is 12, which we shall denote by E . The number of cells in an inhibitory neighborhood of a given cell is $2D-2$, excluding the cell itself.

2.3. Parameter Values and Some Facts

The parameter values chosen for our original algorithm¹ (Marr and Poggio, 1976) were $E=12$, $D=7$, $\varepsilon=2$, $\theta=4$, with the excitatory neighborhood shown in Figure 2. Among other constraints, these parameter values were chosen to satisfy the following conditions:

2.3.1. in the absence of inhibition and of a contribution from the term C^0 , straight line borders should fill in as shown in Figure 3a. This is true when $\theta \leq 4$.

2.3.2. Straight line borders between two “filled-in” planes at different disparities should not grow. This requires that $4 - 2\varepsilon < \theta$.

2.3.3. With the particular values chosen: — A pattern of 5 connected points is the smallest configuration that can survive (see Fig. 3b). It will not grow unless one other point is added (e.g. at P in Fig. 3b).

— The sharpest convexity capable of surviving against one inhibition, with the help of a contribution from C^0 is a right-angle. Figure 3c shows that the condition is $6 - \varepsilon \geq \theta$.

— A convex or flat border cannot grow against one inhibition; it can grow only into scattered active cells.

— The least concave patterns capable of growing under two inhibitions are shown in Figures 3d and e. They fill in by one or two cells and then are no longer concave enough to grow under two inhibitions.

3. Invariant States and the Matching Rules

The matching rules for stereopsis that were given in the introduction take the following form for the algorithm discussed here:

(1) *Uniqueness.* Each item from each image may be assigned at most one disparity value.

¹ In Marr and Poggio (1976), the value of θ was given as 3.0, whereas here it is 4.0. The reason for the discrepancy is that the algorithm used to produce the stereograms for that article essentially used the condition $> \theta$, whereas here, we use the condition $\geq \theta$




Fig. 4. The solid lines indicate solution planes (cf. Fig. 1a). Lines of sight PQ_R , PQ_L intersect solution planes at only one point P , except possibly near the (rare) disparity boundaries like A . Thus configurations that obey rule $R1$ are invariant

(2) *Continuity.* Disparity does not change almost everywhere.

Comment. $R2$ has now taken a slightly different form. This is because disparity takes only discrete value in this algorithm. Images containing smoothly varying disparities may be handled by a modified version of the algorithm, which will be discussed in Section 5.

We now show that the states in which these two rules are obeyed are for all practical purposes *invariant*, i.e. they are fixed points of (1), and once achieved, do not change in subsequent iterations.

3.1. Configurations that Satisfy the Matching Rules are Invariant

The continuity and uniqueness conditions mean that, for each value of y , a cross-section of the network has the appearance shown in Figure 4 (the continuity condition also requires that the active segment has some extension in the y direction). That is, the “on” cells in the network form extended segments like that shown as AB (continuity), and most lines of sight (e.g. PQ_L , PQ_R) intersect only one of these extended segments (uniqueness). Some lines of sight (e.g. to D) may intersect two planes: this occurs only at the (rare) boundaries at which disparity changes. The physical situation is that one surface is obscuring the other.

We show now that these configurations are invariant if the parameter values are appropriate.

(i) Interior points like P are certainly invariant if

$$\sum_{x', y', d' \in S(x, y, d)} C_{x', y'; d}^t \geq \theta \quad (4)$$

if P is interior in both x and y .

(ii) Equation (4) implies that boundary points like A (Fig. 4) on a straight boundary (in the $x-y$ plane) will not grow into the interior of an existing segment at another disparity provided that

$$E/2 + 1 - 2\varepsilon < \theta. \quad (5)$$

Concave pieces of boundaries can in principle grow, but not much for two reasons. Firstly, boundaries


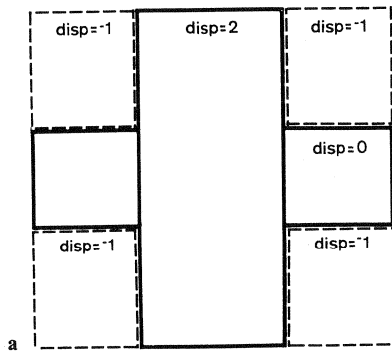



Fig. 5. The two possible stable edges for flat boundaries. Depending on the initial conditions, edges can occur that are defined by the line where cells begin receiving one (A) or two (B) inhibitions from the other surface

cannot be everywhere concave, and secondly, with our particular excitatory neighborhood and parameter values (see Figs. 3d and e) the amount a concave border can fill in is limited to at most two elements. Figure 5 shows the two possible stable edges for flat boundaries.



a




b

3.2. Not All Invariant Configurations Satisfy the Matching Rules

Strictly speaking, the converse result to that of the last section is not true. A counter-example to the uniqueness condition that is stable with our parameters appears in Figure 6. Interior points of a plane, wholly surrounded by other points in the same sheet, can survive inhibition from two other cells and so can boundary points where the boundary is straight. In Figure 6b, points of these two types are the only ones that occur. A counter-example to the continuity condition appears in Figure 7, and it is left as an exercise to show that this pattern is invariant. In practice, neither of these configurations can actually develop from a random-dot stereogram.

When the input consists of two stereograms portraying a single surface, the probabilistic analysis of the next section shows that with high probability, the solutions will in fact obey the uniqueness condition.

Fig. 6a—c. A stable geometrical configuration that violates the uniqueness condition a. The central square consists of two planes, one at disparity 2 and one at disparity 0. This configuration is a stable state of the algorithm, in the sense that if it is loaded directly into the network, an invariant configuration is quickly reached in which both planes are represented; b demonstrates this. The stereogram is marked Left and Right, and 5 iterations of the algorithm are shown. If the network is loaded in the usual way, however, the algorithm develops a solution that is a mosaic of patches from the two levels c



c

If the input stereograms portray a transparent surface in front of another surface, the algorithm with our parameter values will usually fail to represent the input accurately, tending instead to develop a solution that obeys the two conditions and consists of a mosaic of patches from the two levels (Fig. 6c). With the parameters we chose, there seems to be no convenient and precise definition of the stability of configurations that forces the uniqueness and continuity of solutions. For instance, even if one requires in addition to invariance some kind of *spatial stability*², the counter-example of Figure 6 cannot be avoided, although a reasonable "spatial stability" condition would exclude the counter-example of Figure 7.

If one could exclude significant overlaps between surfaces lying at different disparities, it appears that one can derive the continuity conditions for invariant configurations. The argument is based here on the notion of a *hole*³, and shows by straightforward geometry that holes are not invariant.

In one dimension (in which the network consists *only* of the part shown in Fig. 1a) the problem of this section becomes easier. Apparently, the only way of reducing the 2-dimensional problem to a satisfactory state is by changing the parameter values (see Section 6).

4. Probabilistic Analysis of the Algorithm

We have been unable to obtain general results about the convergence of this type of algorithm. Standard approaches — e.g. Liapunov-type methods and the usual fixed point theorems — apparently fail in this situation for reasons that we shall mention in the discussion.

The probabilistic analysis given here, although not completely satisfactory, nevertheless provides useful information about the algorithm's convergence for random-dot stereograms. Strictly speaking its application is restricted to inputs with a random structure.

The idea behind our analysis is that the cells in the network can be divided into populations on which the excitatory and inhibitory inferences are statistically homogeneous (cf. Marr, 1971). Our analysis is very specific to the algorithm of (1) because the way in which the cells are divided into populations depends critically on the geometry of the algorithm and on our a priori knowledge of its invariant state.

² A configuration is "spatially stable" if it is in some sense invariant under small perturbations (for instance each active point can be required to belong to a 3×3 neighborhood of points with the same disparity)

³ There is a hole in the network for a given y if there exist two intersecting lines of sight neither of which contains an "on" cell




Fig. 7a and b. A stable geometrical configuration that violates the continuity condition. At each of two disparity values, the "on" cells form a checkerboard pattern, but they are arranged in such a way that neither level can fill in, because of inhibition from the other

4.1. Assumptions and Notation

The algorithm has the structure shown in Figure 1 and the network is loaded from the input as specified by (2). We shall assume that the inputs have the following properties.

4.1.1. The 1's in each image occur randomly with probability ν , and the autocorrelation of each input sequence (for any given y) is a Kronecker δ .

4.1.2. The input admits a unique solution surface that is large enough to neglect boundary effects.

Condition 4.1.2 means that the left input is equal to the right one, modulo x -translation. Condition 4.1.1 implies that in the initial state of the network C , the density of 1's on the solution layer equals ν , and elsewhere it is ν^2 . We subdivide the cells into five populations, by classifying them in two ways:

- (i) according to whether or not they are a "on" in the initial state C^0 , and
- (ii) according to the number of active inputs from the images.

We draw both the populations 0 and 1 from cells that lie on the solution layer; population 0 is defined to receive no active inputs from the image, and population 1 receives two. Notice that there are no cells in the solution layer that receive exactly one active input.

The other three populations that we define refer to cells that lie off the solution layer; population 11 receives two active inputs from the image, population 10 receives one, population 00 receives none. The five populations $\{0, 1, 11, 10, 00\}$ are exclusive and exhaustive.

We denote by $p_0(t)$, $p_1(t)$, etc. the probability that a cell in the respective population is "on" at time t . The goal of our analysis is to express the values of the $p_\pi(t)$ in terms of $p_\pi(t-1)$ for the various populations π . This allows us to examine the convergence numerically, and

Table 1a—d. The behavior of the algorithm compared with the probabilistic theory of the algorithm, for the stereograms having four different densities that are exhibited in Figure 8

a. $v=0.5, E=12, D=7, \epsilon=2, \theta=4.0$ (SQ 50%)

Iteration		p_r	p_w	p_0	p_1	p_{00}	p_{01}	p_{11}
1	Algorithm	0.46	0.07	0.93	0.01	0.29	0	0
	Theory	0.47	0.087	0.92	0.01	0.35	0	0
2	Algorithm	0.70	0.04	0.43	0.97	0	0	0.16
	Theory	0.62	0.02	0.26	0.97	0.04	0	0.08
3	Algorithm	0.90	0.01	0.99	0.81	0.04	0.001	0.001
	Theory	0.97	0	0.99	0.96	0	0	0
4	Algorithm	0.98	0.002	0.96	1.0	0	0	0.01
	Theory	1.0	0	1.0	1.0	0	0	0
5	Algorithm	0.99	0	1.0	0.99	0	0	0
	Theory	1.0	0	1.0	1.0	0	0	0

b. $v=0.25, E=12, D=7, \epsilon=2, \theta=4.0$ (SQ 25%)

Iteration		p_r	p_w	p_0	p_1	p_{00}	p_{01}	p_{11}
1	Algorithm	0.24	0.001	0.31	0.02	0.002	0	0
	Theory	0.27	0.003	0.35	0.04	0.005	0	0
2	Algorithm	0.39	0	0.26	0.81	0	0	0
	Theory	0.475	0	0.41	0.68	0	0	0
3	Algorithm	0.58	0	0.52	0.78	0	0	0
	Theory	0.92	0	0.90	0.97	0	0	0
4	Algorithm	0.74	0	0.70	0.87	0	0	0
	Theory	1.0	0	1.0	1.0	0	0	0

c. $v=0.1, E=12, D=7, \epsilon=2, \theta=3.0$ (SQ 10%)

Iteration		p_r	p_w	p_0	p_1	p_{00}	p_{01}	p_{11}
1	Algorithm	0.10	0	0.11	0.05	0	0	0
	Theory	0.11	0	0.11	0.106	0	0	0
2	Algorithm	0.20	0	0.15	0.72	0	0	0
	Theory	0.17	0	0.14	0.39	0	0	0
3	Algorithm	0.36	0	0.32	0.68	0	0	0
	Theory	0.35	0	0.32	0.61	0	0	0
4	Algorithm	0.53	0	0.51	0.77	0	0	0
	Theory	0.86	0	0.85	0.96	0	0	0
5	Algorithm	0.71	0	0.69	0.83	0	0	0
	Theory	1.0	0	1.0	1.0	0	0	0

d. $v=0.05, E=12, D=7, \epsilon=2, \theta=2.0$ (SQ 05%)

Iteration		p_r	p_w	p_0	p_1	p_{00}	p_{01}	p_{11}
1	Algorithm	0.13	0	0.12	0.22	0	0	0
	Theory	0.13	0	0.12	0.26	0	0	0
2	Algorithm	0.38	0	0.35	0.88	0	0	0
	Theory	0.48	0	0.46	0.81	0	0	0
3	Algorithm	0.67	0	0.65	0.88	0	0	0
	Theory	1.0	0	1.0	1.0	0	0	0
4	Algorithm	0.90	0	0.89	0.96	0	0	0
	Theory	1.0	0	1.0	1.0	0	0	0
5	Algorithm	0.98	0	0.98	0.98	0	0	0
	Theory	1.0	0	1.0	1.0	0	0	0

we say that a solution is achieved at time T when

$$p_0(t)=p_1(t)=1, \text{ and} \\ p_{00}(t)=p_{10}(t)=p_{11}(t)=0, \text{ for every } t \geq T.$$

The critical assumption here is that the quantity $p_\pi(t)$ completely describes the structure of active cells in the respective population π . This assumption is true for the initial iteration and only approximate thereafter. We shall discuss this point at the end of the section.

4.2. Formulae

The state of a cell (x, y, d) at time $(t + 1)$ depends upon the number of active cells in its excitatory $S(x, y, d)$ and inhibitory $O(x, y, d)$ neighborhoods at time t .

If we denote the populations to which the cell belongs by π , (π running through the five populations 0, 1, 00, 11, 01), let us define:

$e_\pi(r)$ to be the probability that exactly r cells are "on" in the excitatory neighborhood $S(x, y, d)$ at time t and

$i_\pi(r)$ to be the probability that exactly r cells are "on" in the inhibitory neighborhood $O(x, y, d)$ at time t . It is convenient to introduce some further quantities:

$q_s(t)$ is the probability that a given cell on the "solution" plane is active at time t .

$q_w(t)$ is the probability that a given cell elsewhere in the network is active.

$q_-(t)$ is the probability that a given cell is active in the inhibitory neighborhood of a cell in the population 0.

$q_+(t)$ is the probability that a given cell is active in the inhibitory neighborhood of a cell in the population 1.

Then

$$q_s(t) = p_0(t) \cdot (1 - v) + p_1(t) \cdot v \\ q_w(t) = p_{00}(t) \cdot (1 - v)^2 + p_{01}(t) \cdot 2v(1 - v) \\ + p_{11}(t) \cdot v^2 \\ q_-(t) = p_{00}(t) \cdot (1 - v) + p_{01}(t) \cdot v \\ q_+(t) = p_{11}(t) \cdot v + p_{10}(t) \cdot (1 - v).$$

Writing $B(n, f; m) = {}_m C_n \cdot f^n (1 - f)^{m-n}$, where ${}_m C_n$ is the binomial coefficient, we have immediately

$$e_1(r) = e_0(r) = B(r, q_1(t); E) \\ i_1(r) = B(r, q_+(t); 20 - 2) \\ i_0(r) = B(r, q_-(t); 20 - 2) \\ e_{11}(r) = e_{00}(r) = e_{10}(r) = B(r, q_w(t); E).$$

The remaining i_π are more difficult to obtain, since the inhibitory contributions to cells lying off the solution plane come from cells lying on the solution plane and from cells lying off the solution plane, and these two

populations obey different statistics. In fact

$$\begin{aligned} i_{11}(r) = & [p_1(t)]^2 \cdot B(r-2, q_+(t); 2D-4) \\ & + 2p_1(t) \cdot (1-p_1(t)) \cdot B(r-1, q_+(t); 2D-4) \\ & + [1-p_1(t)]^2 \cdot B(r, q_+(t); 2D-4) \end{aligned} \quad (8)$$

$$\begin{aligned} i_{00}(r) = & [p_0(t)]^2 \cdot B(r-2, q_-(t); 2D-4) \\ & + 2p_0(t) \cdot (1-p_0(t)) \cdot B(r-1, q_-(t); 2D-4) \\ & + [1-p_0(t)]^2 \cdot B(r, q_-(t); 2D-4). \end{aligned}$$

The final case i_{10} is especially awkward, because along one of the inhibitory lines the probability of a cell being "on" is q_+ and along the other diagonal it is q_- .

$$\begin{aligned} i_{10}(r) = & \Sigma \{ p_1(t) B(k-1, q_+(t); D-2) \\ & + (1-p_1(t)) B(k, q_+(t); D-2) \} \\ & \cdot \{ p_0(t) \cdot B(r-k+1, q_-(t); D-2) \\ & + (1-p_0(t)) B(r-k, q_-(t); D-2) \}. \end{aligned} \quad (9)$$

We now need to relate the $p_\pi(t+1)$ and the $p_\pi(t)$ in terms of the e_π and i_π . For each cell population we know the distributions of incoming excitation and inhibition, and we know that a cell will be on whenever the excitations exceed the inhibitions by at least θ . Hence:

$$p_\pi^{t+1} = \sum_{\substack{n=\theta \text{ to } E \\ M=\theta \text{ to } 2D-2 \\ n-em \geq \theta_\pi}} e_\pi^t(n) \cdot i_\pi^t(m) \quad (10)$$

where

$$\theta_\pi = \theta - 1 \quad \text{for } \pi=1, \pi=11$$

$$\theta_\pi = \theta \quad \text{otherwise.}$$

If the input term $C_{x,y,d}^0$ of (1) is neglected, $\theta_\pi = \theta$ for all π .

The Equations (10) are too complex to be solved analytically. Numerical solutions were however obtained for various values of the parameters and some of the results are given in Table 1 and Figure 8.

4.3. Range of Parameter Values and Comparison with Actual Runs

Figure 8 exhibits the performance of the algorithm for stereograms having densities of from 0.5 to 0.05. Table 1 gives the statistics that were measured from these runs, and also the parameters predicted by the probabilistic theory. The values obtained from the theory match those from the algorithm quite well for the first iteration, but except for the case $\nu=0.05$, they diverge quite rapidly thereafter, and even this case diverges by the third iteration.

We have already noted the main reason for the discrepancy. The assumption that the statistical structure of various populations is purely random (inside each population and between populations) holds exactly for the first iteration but only approximately

thereafter, because the operator of Figure 1c has a local structure which can preserve local clusters of active cells. There are two ways in which this affects our probabilistic description for the second and subsequent iterations. The first is that clusters are more stable than the assumption of randomness would predict. Thus clusters forming on the solution layer will in certain circumstances change the rate of convergence predicted by the randomness assumption.

The difficulties arise where clusters form off the solution layer. These will again tend to be more stable than our analysis assumes, but their effect acts against convergence. However, we shall argue that the probability of large "wrong" clusters is small for most patterns. In fact, the typical value of the probability that a wrong cell is "on" after the first iteration lies around 0.1. The probability (after the first iteration) of a self-supporting 3×3 cluster at a given position in a wrong layer (assuming that the cluster was absent in the initial state and accepting the oversimplified assumption of randomness after the first iteration) is about 10^{-9} , and hence less than 10^{-4} that one exists off the solution plane somewhere in the network.

A cluster of this size may survive permanently, because every element in it has at least 6 cells in its excitatory neighborhood, and this is enough to resist 1 inhibition. The probability of this or something larger arising by chance is so small that if it occurs it is likely to be a consequence of the particular image. In fact, some small "wrong" patches do sometimes occur (inspect Marr and Poggio, 1976, Fig. 5d) but such instances can usually be traced to an accidental correlation in the image. In this sense, extended patches are "correct" solution regions.

The second effect that leads to discrepancies between the theory and the behavior of the algorithm is also a side-effect of clustering, since as well as being stable, the clusters tend to concentrate "on" cells more than the randomness assumption would predict. For example, at iteration 2 of the case $\nu=0.25$ (Fig. 8b), although the overall density of ones on the solution plane is about 0.39, it is far from true that each cell can expect to find $0.39E$ "on" cells in its excitatory neighborhood. Cells in the filled in regions have almost all their neighbors on, whereas those in the interstices have none. Convergence is achieved by a growth outwards that fills in the blank regions, but although it is steady, it is necessarily slower than the theory predicts.

5. Observations

5.1.

There is a wide latitude in the range of parameters for which the network converges. Table 2 shows firstly the

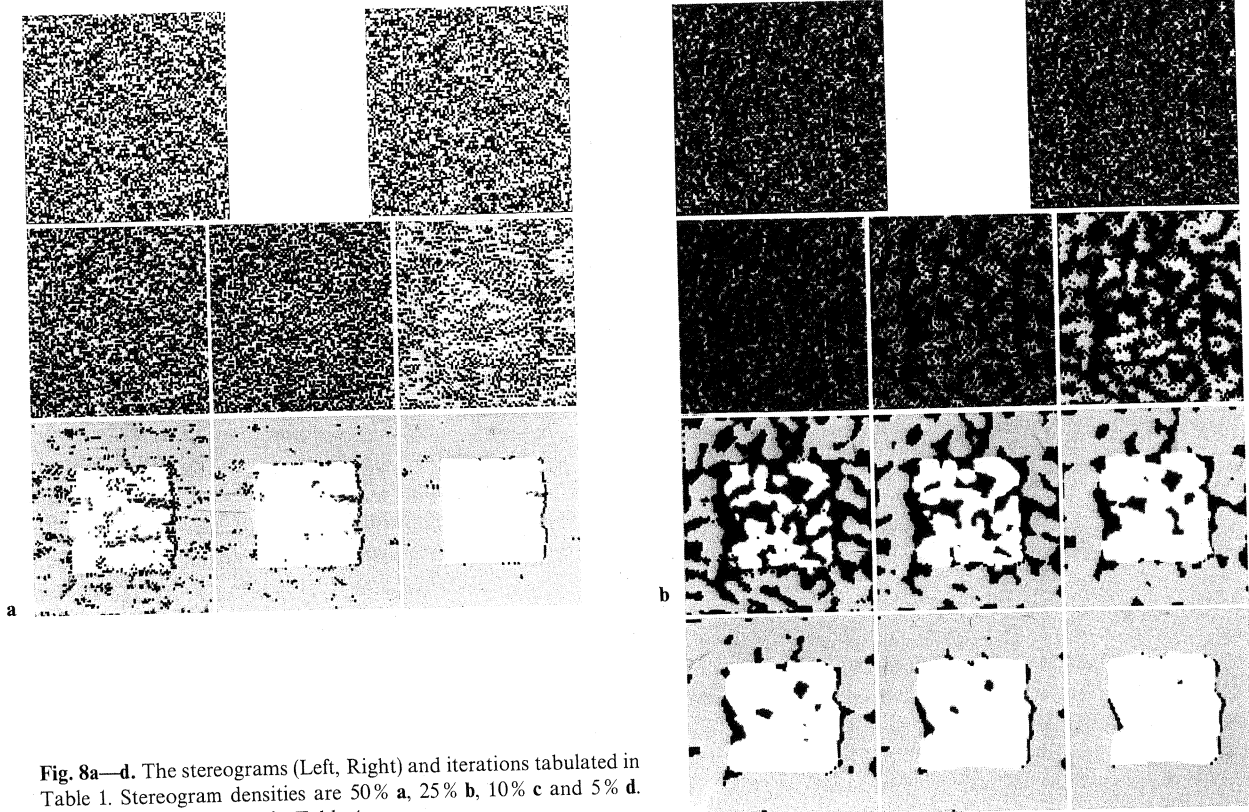



Fig. 8a—d. The stereograms (Left, Right) and iterations tabulated in Table 1. Stereogram densities are 50% a, 25% b, 10% c and 5% d. Parameters are as shown in Table 1



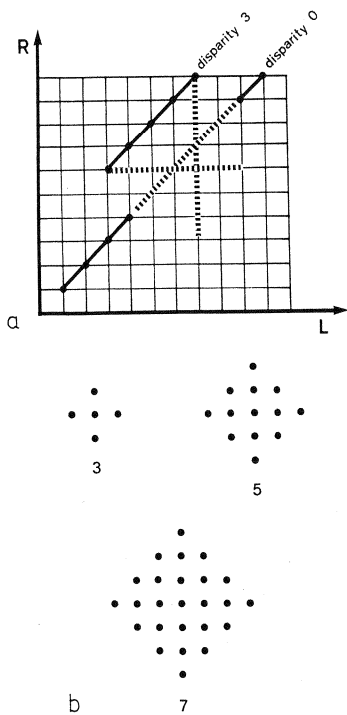


Fig. 9a and b. The minimum resolvable area of a small pattern against a background increases with disparity. To prevent the background from filling in completely, the length of the patch in the x -direction must be at least $|d| + 2a$. **b** Shows the circles of diameters 3, 5 and 7 used in Figure 6 of Marr and Poggio, 1976

wide range in stereogram density v that is tolerated by our parameters (with fixed θ), and secondly, for a fixed value of v ($v=0.5$) gives some idea of the range of the other parameter values for which the network will converge. Note that in the implementation described by Marr and Poggio (1976), the threshold was not fixed, but was determined by the density of "on" cells in the network. This allowed solution to the matching problem over a very wide range of dot densities.

5.2.

Let us define the probability that a cell on the solution layer is "on" at time t to be

$$p_r(t) = v \cdot p_1(t) + (1 - v)p_0(t)$$

and the probability that a cell off the solution layer is on at time t as

$$p_w(t) = v^2 \cdot p_{11}(t) + 2v(1 - v)p_{10}(t) + (1 - v^2)p_{00}.$$

In a successful run, p_r converges to 1 and p_w to 0. With our particular parameters, convergence is monotonic if it occurs. This is not true, however, for the individual quantities p_1 , p_0 , p_{11} , p_{10} , p_{00} , neither is it true of p_r and p_w for all values of the parameters (see Table 2).

5.3.

We have already seen that the sharpest local corner capable of resisting 1 inhibitory input is about 90° or

Table 2a—c. The algorithm of (1) converges for a wide range of control parameters. Tables 2a,b show convergence for $v=0.5$ and $v=0.1$ with the same parameters. Table 2c shows convergence for an entirely different set of parameters

a. $v=0.5$, $E=12$, $D=7$, $\varepsilon=2$, $\theta=3.0$

Iteration	p_r	p_w	p_0	p_1	p_{00}	p_{10}	p_{11}
1	0.50	0.15	0.98	0.026	0.61	0	0
2	0.57	0.13	0.15	0.997	0	0	0.54
3	0.69	0.039	0.995	0.39	0.16	0	0
4	0.97	0.007	0.935	1.0	0	0	0.029
5	1.0	0	1.0	1.0	0	0	0

b. $v=0.1$, $E=12$, $D=7$, $\varepsilon=2$, $\theta=3.0$

Iteration	p_r	p_w	p_0	p_1	p_{00}	p_{10}	p_{11}
1	0.11	0	0.11	0.106	0	0	0
2	0.17	0	0.14	0.39	0	0	0
3	0.35	0	0.32	0.62	0	0	0
4	0.86	0	0.85	0.96	0	0	0
5	1.0	0	1.0	1.0	0	0	0

c. $v=0.5$, $E=2$, $D=7$, $\varepsilon=0.5$, $\theta=1.0$

Iteration	p_r	p_w	p_0	p_1	p_{00}	p_{10}	p_{11}
1	0.40	0.11	0.75	0.058	0.43	0	0.010
2	0.55	0.23	0.11	0.99	0.004	0	0.90
3	0.45	0.083	0.78	0.11	0.32	0	0.006
4	0.59	0.20	0.20	0.99	0.003	0	0.80
5	0.51	0.063	0.82	0.20	0.24	0	0.009
6	0.65	0.17	0.32	0.99	0.003	0	0.66
7	0.62	0.042	0.87	0.36	0.15	0.001	0.014
8	0.76	0.11	0.54	0.97	0.002	0	0.43
9	0.82	0.021	0.94	0.71	0.053	0.002	0.026
10	0.94	0.027	0.88	0.995	0	0	0.11
11	0.995	0.009	0.996	0.995	0.001	0	0.031
12	1.0	0.004	1.0	1.0	0	0	0.014
13	1.0	0.002	1.0	1.0	0	0	0.007
14	1.0	0	1.0	1.0	0	0	0.003

more, hence thin, sharp regions will tend to be rounded off locally (see Marr and Poggio, 1976, Fig. 5c). The exact shape of the input pattern is preserved only up to this limit.

5.4. Minimum Size vs. Disparity

A natural consequence of the structure of the algorithm is that the minimum resolvable area of a small pattern against a background increases with disparity (see Marr and Poggio, 1976, Fig. 6). We give an estimate of the dependence of minimum patch size on disparity difference. Consider a section for some fixed y of the network (Fig. 9). Assume that the patch and background regions are filled in. The condition for growth at a point (x, y, d) under 1 inhibition is that the number of "on" cells in an excitatory neighborhood should be not less than $\theta + \varepsilon - C^0 = 5$ or 6, depending




Fig. 10a—c. Thin vertical and horizontal stripes of various disparities. The left and right stereograms are shown with the stable network solution to them. The stereograms are 100 by 100, and consist of stripes with the following coordinates (x or y), thicknesses s and disparities d :

a			b			c		
x	s	d	y	s	d	y	s	d
15	2	-1	15	2	-1	15	2	+1
30	3	-1	30	3	-1	30	3	+1
45	3	-2	45	2	-2	45	2	+2
60	4	-2	60	3	-2	60	3	+2
75	4	-3	75	2	-3	75	2	+3
90	5	-3	90	3	-3	90	3	+3

on the initial conditions. From Figure 3 we see that flat or convex regions will not grow whereas concave regions will. Hence our small patch will not tend to grow, whereas the background will spread until stopped by two inhibitions. We see from Figure 9a that to prevent the background from filling in completely (which would subsequently destroy the patch because convex borders cannot survive two inhibitions), the length of the patch in the x direction must be at least $|d|+2$. This condition must hold for at least three adjacent lines aligned in the y direction. Figure 6 of Marr and Poggio (1976) illustrates the approximate validity of this relation. Figure 9b shows the sizes of circles of diameter 3, 5, and 7 used in the input for that figure. These precise patterns do not necessarily emerge in the appropriate layer of the network because of the random nature of the borders. The circle of diameter 3 contains no 3×3 subset and therefore does not survive at any disparities. The circle of diameter 5 contains one 3×3 square and survives as expected at

disparity 1; it also survives, apparently accidentally, at disparity 2, but not at disparity 3. The circle of diameter 7 contains one 5×5 square and thus survives at disparity 3.

A trivial consequence of this analysis is that horizontal stripes (parallel to the x axis) are in general more stable than vertical ones (parallel to the y axis). The minimum thickness for horizontal stripes is about 3 and is independent of disparity whereas the minimum thickness for vertical stripes is about $|d|+2$ (see Fig. 10).

5.5. Uncorrelated Areas

If there exists a sufficiently large area in the input where there is no correlation between the two images, the network will detect it (see Figs. 5 and 6 of Science). After the first iteration (with our parameter values and $v=0.5$) only a few cells remain "on" in the uncorrelated region, but provided the region is sufficiently large they will receive no inhibition from the surrounding more organized layers. Hence those cells that are on may act as germs for small regions that have become stable by the time the surround encroaches upon them, e.g. Figure 5d of Marr and Poggio (1976). Relatively small ($\ll d$) uncorrelated areas probably have to develop stable platelets to survive (see Fig. 6d of Marr and Poggio, 1976), and large uncorrelated areas decompose into a random mosaic of stable platelets (see Fig. 11).

Uncorrelated areas can be recognized as such during the read-out from the network, when the 1's that appear in the solution found by the network are used to establish an explicit correspondence between the two images.

5.6. Extension to Images in which Disparity Varies Continuously

The algorithm of (1) with the loading rules of (2) can deal only with images having discrete disparity values. This disparity in natural images commonly varies continuously. There are two approaches to this problem. One is to incorporate the representation of continuous values directly into the algorithm, and the other is to use the same algorithm, but with special rules for loading it and for interpreting its final state.

The first approach would clearly lead to a considerably different algorithm, perhaps more along the lines of the networks studied by Wilson and Cowan (1973), (see also Wilson, 1977). Such an algorithm could not be treated within the framework of this article.

The second approach does not require any changes in the analysis of the algorithm itself. One could, for example, define the loading conditions as follows:

Let Δ be the disparity attached to a possible correspondence between items in the left and right images. For integral d ,

5.6.1. If $d - \eta \leq \Delta < d + \eta$, load the cell corresponding to disparity level d in the network.

For surfaces whose disparity does not oscillate too much or too densely, the value $\eta = 0.5$ will lead to satisfactory results. The final state of the network establishes a correspondence between items in the left and right images, but their associated disparity is read not from the network (i.e. d) but directly from the input (i.e. Δ). Confusions may of course arise in the correspondence established by the network if the value of d spans the disparity range too coarsely.

In order to deal with surfaces that are less well-behaved, one can incorporate some hysteresis into the loading rules. The loading process then consists of the following steps:

5.6.2. Load cells according to 5.6.1 with $\eta = 0.3$ (say).

5.6.3. Moving across the image (x, y) in a spatially ordered way, if a possible match (x, y, Δ) was not loaded by 5.6.2, adopt the following procedure:

Let $d^- = \text{Integral part of } \Delta$, $d^+ = 1 + d^-$. Examine (x, y) neighborhoods of (x, y, d^-) and of (x, y, d^+) in the network as it is loaded so far. Assign the current match to that d whose neighborhood contains more loaded cells, if one of them does. Else load this point according to 5.6.1 with $\eta = 0.5$.

This process will load most images in satisfactory way, and the read-out procedure is similar to that of the previous case.

6. A Mathematically Tractable Version of the Algorithm

A suitable choice of the parameter values and of the loading rules of the algorithm allows a complete mathematical analysis of its asymptotic behavior. In this section we introduce this "strict" version of the algorithm and we characterize rigorously its properties. The actual performance of this version of the algorithm for various random dot stereograms will be then compared with the original algorithm.

6.1. Loading Conditions

The initial state C^0 of the network is loaded from the stereograms L, R in a way similar to the previous case but according to (11) (instead of (2)).

$$C_{x,y;d}^0 = L_{x,y} \cdot R_{x+d,y} \quad (11)$$

where $1 \cdot 1 = 0 \cdot 0 = 1$, $1 \cdot 0 = 0 \cdot 1 = 0$.

This loading rule can be easily extended to cases in




Fig. 11. The central band is uncorrelated. It decomposes into a random mosaic of patches, each of which is eventually stable

which more than two features are present. It is enough to define

$$f_i \cdot f_j = \delta_{ij}, \quad (12)$$

where f_i and f_j , ($i \neq j$) are two different features. The case when only two features are present clearly poses the hardest matching problem. We shall later compare this loading rule with the original (2) and discuss their relative merits for real images.

6.2. The Algorithm

The relation between states at times t and $t+1$ is given by (compare (1))

$$C_{x,y;d}^{t+1} = \sigma \left\{ \inf \left[\sum_{x',y',d' \in S(x,y,d)} C_{x',y';d'}^t, H \right] - \epsilon \sum_{x',y',d' \in 0(x,y,d)} C_{x',y';d'}^t \right\} \quad (13)$$

where H is a number that represents the "saturation" value for the excitation.

6.3. Choice of Parameter Values

In this case the loading rules lead, for random dot stereograms with two features, to a density of 1 for the

Table 3a and b. The behavior of the mathematically tractable version of the algorithm, together with the probabilistic theory of the first iteration, for the two stereograms exhibited in Figures 14a and b

a. $v=0.5, E=13, D=7$

ϵ	θ	H	Iteration	p_r	p_w	p_0	p_1	p_{00}	p_{10}	p_{11}	
0.2	10.75	13.0	1	0.9998	0.0017	0.9998	0.9998	0.0023	0.0011	0.0023	Theory Algorithm
			1	0.99	0.0003	0.99	0.98	0.0008	0	0.0004	
4.0	3.75	7.0	2	0.99	0	0.99	0.99	0	0	0	
			3	1.0	0	1.0	1.0	0	0	0	

b. $v=0.25, E=13, D=7$

ϵ	θ	H	Iteration	p_r	p_w	p_0	p_1	p_{00}	p_{10}	p_{11}	
0.2	10.75	13.0	1	0.976	0.0078	0.968	1.0	0.0039	0.0015	0.082	Theory Algorithm
			1	0.96	0.0002	0.95	1.0	0.002	0.0003	0.002	
4.0	3.5	7.0	2	0.98	0	0.98	0.97	0	0	0	
			3	1.0	0	1.0	1.0	0	0	0	

“on” cells on the “correct” diagonal segments and, correspondingly, to a density of $v^2 + (1-v)^2$ for the “on” cells on the “wrong” diagonal segments (v is the density of 1’s in the input images). When $v=0.5$, the density of the wrong cells is also 0.5; for smaller or larger v the density is higher. The idea behind this approach is to choose parameter values for the first iteration that “kill” most of the “wrong” cells (and of course some of the “right” ones); from the second iteration on, the parameter values are such to ensure “filling-in” of the right diagonal segments, allowing, at the same time, a satisfactory mathematical analysis of the evolution of the network’s state. This approach, which is carried out in the next two sections, leads to the following parameter values:

6.3.1. S and O are as in Figure 2, $D=7$ and $M=5$ as before. Selfexcitation is now included but the C^0 term is omitted. We therefore write $E=13$ instead of 12.

6.3.2. *Iteration 1:*

$H=13$ (so that the inf operation can be neglected)
 $\epsilon=0.2$.
 $\theta=10.75$.

6.3.3. *Second and Subsequent Iterations:*

$H=7$
 $\epsilon=4.0$
 $\theta=3.5$.

6.4. Probabilistic Analysis of the First Iteration

We shall assume that the inputs have the properties 4.1.1 and 4.1.2. As in Section 4.1, we distinguish several populations of cells which are homogeneous with respect to the interaction structure: the populations are again denoted by 0, 1, 11, 10, 00 according to their respective inputs from the two images (see Section 2), and p_0, p_1 , etc. denote the probability that a cell in the

respective population is “on” after the first iteration. In this case the formulae for the solution layer are:

$$p_1 = \sum_{i=0}^{\epsilon} {}_{12}C_i \cdot v^i (1-v)^{12-i}$$

$$p_0 = \sum_{i=0}^{\epsilon} {}_{12}C_i \cdot (1-v)^i \cdot v^{12-i}$$

For the “wrong” layers (writing $\mu = v^2 + (1-v)^2$), the formulae are

$$p_{11} = \sum_{k=0}^{12} {}_{12}C_k \mu^k (1-\mu)^{12-k} \sum_{i=0}^{\frac{k+1-\theta}{\epsilon} - 2} {}_{10}C_k v^i (1-v)^{10-i}$$

$$p_{00} = \sum_{k=0}^{12} {}_{12}C_k \mu^k (1-\mu)^{12-k} \sum_{i=0}^{\frac{k+1-\theta}{\epsilon} - 2} {}_{10}C_k (1-v)^i v^{10-i}$$

$$p_{10} = \sum_{k=0}^{12} {}_{12}C_k \mu^k (1-\mu)^{12-k} \sum_{i=0}^{\inf\{\frac{k-\theta}{\epsilon}, 5\}} {}_5C_i v^i (1-v)^{5-i} \sum_{j=0}^{\inf\{\frac{k-\theta}{\epsilon} - 2, 5\}} {}_5C_j (1-v)^j v^{5-j}$$

Therefore the probability that a cell in the solution layer is “on” after the first step is

$$p_r = p_1 \cdot v + p_0 \cdot (1-v)$$

and the probability that a cell off the solution layer is “on” after the first step is

$$p_w = v^2 p_{11} + 2v(1-v)p_{10} + (1-v)^2 p_{00}$$

These equations can be used to find suitable parameter values. The parameters given in the previous section yield the values for p_r and p_w shown in Table 3.

6.5. Equivalent Rules

The parameter values from the second iteration on imply the following main "rules" for the algorithm:

6.5.1. One "on" cell in the inhibitory neighborhood always suffices to kill an "on" cell.

6.5.2. Without inhibition, at least three excitatory "on" cells are needed for "survival" of an "on" cell and four for its "birth".

6.6. Analysis of the Second Iteration

Table 3 gives the densities p_r and p_w after the first iteration. Only for the first iteration can a probabilistic analysis provide a reliable estimate of the density of "on" cells on the solution surface. As in our earlier analysis (Table 1), it becomes unreliable for the second iteration, because clusters of "on" cells can be expected to form off the solution layer (see Fig. 14 below). Rule 6.5.1 implies, however, that "wrong" clusters will disappear after the second iteration, unless they consist of at least four elements. Moreover, these elements must in practise be very close together for each to support the other three. In addition, according to rule 6.5.2, none of them can lie in the inhibitory neighborhood of other "on" cells (for instance on the solution layer where the density of on cells is relatively high, see Table 3). We argue that the probability of such situations is very small (actually much smaller than in the case considered in Section 4.4). If this occurs it can be attributed to an accidental correlation in the images. In this sense extended clusters are in fact "right" solution regions.

6.7. Asymptotic Analysis

The probabilistic analysis of the first iteration (Table 3) shows that one can assume that, from the second iteration onwards, there are no wrong "on" cells. It remains now to show that the density of "on" cells on the solution layer is high enough to allow asymptotic filling-in of the "right" surfaces. We prove the following:

6.7.1. Filling-in Proposition. Assume that (at some iteration n) there are no "on" cells off a given layer (diagonal), and that the density of "on" cells on this layer exceeds $0.4375 = 7/16$. Then, in the asymptotic configuration, there are no "off" cells on this layer.

Proof. Divide the solution plane into squares of 4 by 4 cells (we neglect boundaries). At least one of these squares must contain 8 "on" cells, for, otherwise, every square would contain at most 7 "on" cells yielding a density of at most $7/16$, in contradiction with the hypothesis. This square will fill up with "on" cells.




Fig. 12. With the modified parameters, cells cannot survive against one inhibition. Hence stable states satisfy the uniqueness condition, because no overlap is possible (compare Fig. 5)




Fig. 13a and b. An oscillating solution with the modified parameters. The state a occurs at iterations $i, i+2, i+4, \dots$, whereas state b occurs at iterations $i+1, i+3, i+5, \dots$

(This can be seen by examining the various possible ways in which the 8 cells can be distributed, and we leave it as an exercise for the reader). Starting from this square, the whole plane will asymptotically be filled by "on" cells (since, by hypothesis, no inhibitory cells need be considered).

6.8. Invariant States and Matching Rules

The matching rules were defined in Section 3. States that satisfy the matching rules with the present parameter values are shown in Figure 12. In view of the rules 6.5.1 and 6.5.2, the following clearly hold:


i) Configurations that satisfy the matching rules (Fig. 12) are invariant.

ii) Conversely, invariant configurations clearly have to obey the uniqueness condition (because of 6.5.1). The probabilistic analysis of the second step, together with the "filling-in" proposition 6.7.1, ensures in practise that there will be no holes² in the asymptotic invariant configurations.


6.9. Asymptotic Liapunov Description

Besides the invariant asymptotic configuration, limit cycles of the type described in Figure 13 may also occur. Thus the previous description of asymptotic invariant states is not complete. We provide here an asymptotic analysis in terms of a Liapunov-like function which also encompasses such non-invariant states.

For a given state C^i , we define $F(C^i)$ to be the number of "on" cells having no "on" cells in their inhibitory neighborhood. We call an "on" cell that has less than three "on" cells in its excitatory neigh-



a




b

neighborhood a “solitary cell”. Observe that solitary cells can never be “born” and that, after a finite number of iterations, all solitary cells will have disappeared.

6.9.1. Growth Proposition. After a finite number of iterations, the function $i \rightarrow F(C^i)$ is non-decreasing.

Proof. After a finite number of iterations i , all solitary cells have died out. Let us consider the transition from C^i to C^{i+1} . If a new cell is born, rule 6.5.1 implies that it cannot lie in the inhibitory neighborhood of an already present “on” cell. Thus F will not decrease (from C^i to C^{i+1}). If a cell dies out, it cannot be a solitary cell. Therefore it must have had an “on” cell in its inhibitory neighborhood at iteration i . Thus F will not decrease.



c

Fig. 14a—c. The behavior of the algorithm with modified parameters. The densities are 50% a and 25% b. The parameters are as stated in Section 6.3, and in Table 3. c Compares the two sets of parameters on a stereogram of a star that contains arms of various angles. The original parameters tend to give a more accurate final configuration

The growth of F adequately describes the filling-in process, respecting at the same time the “uniqueness” matching rule.

The growth proposition implies that :

6.9.2. For any initial configuration C^0 , the limit $\text{Lim} F(C^i) = F(C)$ exists (since F is bounded above by the number of cells in one layer).

6.9.3. After a finite number of iterations, $F(C^i)$ remains constant.

Thus the asymptotic behavior of the system is characterized by the following: Apart from invariant solutions, only those cycles can (asymptotically) occur for which F remains constant. This is a strong restriction on the possible asymptotic oscillatory states, for it means that they have to be of the type shown in Figure 13. The growth proposition by itself does not exclude for instance the “zero” invariant state. The probabilistic analysis of the second step of the algorithm together with the “filling-in” proposition ensures, however, that invariant states as well as limit cycles will in practice have no “holes”.

6.10. Observations

6.10.1. Figure 14 shows the performance of the algorithm in this form for a few different patterns and pattern densities. A comparison with Figure 8 reveals that the type of “strategy” for achieving a successful

matching is different: Firstly, wrong cells are drastically eliminated at the expense of losing many right cells, and then filling-in of the surviving surfaces takes place. This contrasts with the more complicated "strategy" revealed in Figures 3-6 of Marr and Poggio, 1976. It is remarkable how, while the basic structure of the algorithm remains the same, a change of parameter values and loading conditions can bring about so deep a change in the algorithms behavior.

6.10.2. Because of the rules of the present algorithm, especially rule 6.5.1, the sharpest corner capable of surviving (of course under no inhibitions) is limited only by the need for an "on" cell to have at least three excitatory neighboring cells. This allows a 45 degree corner.

6.10.3. *Minimum Size vs. Disparity.* Again because of rule 6.5.1, the minimum resolvable area of a small pattern against a background does not depend on disparity. It is given by the minimum self-supporting configuration (four adjacent cells, from rule 6.5.2).

This contrasts sharply with the property discussed in Section 5.4, where the minimum size has a characteristic dependence on disparity.

6.10.4. *Loading Conditions.* While the present loading rule is characterized by

$$f_i f_j = \delta_{ij}, \quad (14)$$

where f_i is a feature and δ_{ij} is the Kronecker δ ,

the previous loading rule (Section 2) can also be characterized by (12) with the convention that δ_{ij} is also zero when either i or j are zero. In other words the "null" feature has a special status ($f_0 f_j = f_j f_0 = 0$, all j).

In case of two-valued (0,1) random dot stereograms, the choice of either one of the two loading rules is somewhat arbitrary. For densities around 0.5, the straight Equation (14) seems to make more sense, since the black and white dots play equivalent roles. This is not clear, however, at very low densities (nor at very high ones).

In the case of natural images, more than two feature types have to be used (for instance, lines and edges at various orientations). In this case, however, not every point is labelled with a corresponding feature; the absence of any feature at a given point is a common event. The null feature seems to have a basically different role from the other features. These arguments clearly support the loading conditions used in the first part of the paper (see Marr and Poggio, 1976). It is clear, on the other hand, that both loading conditions may work. For both, an increasing number of "feature types" implies of course an increasingly better algorithm convergence. The choice between

them depends in the end on the typical feature densities that one wants to deal with. For natural images, quantitative estimates have only recently become possible (Marr, 1976, 1977).

7. Discussion

7.1. Alternative Algorithms

The algorithm (1) can be modified in various ways. One can adopt alternative loading rules for the network as in Section 6, and one can vary the parameters over a substantial range. Such apparently minor changes can cause considerable changes in the network's behavior, but often without changing the end result (see for instance Section 6), because they still implement the same computational constraints.

If the geometry of the local interactions (i.e. the shape of the excitatory and inhibitory neighborhoods) is changed, the network will in general implement a different computation, because the local constraints will have changed. If only the parameter values are changed, our analysis (Section 3) may still apply. If the geometry is changed, our analysis will in general become irrelevant.

Interestingly, for a specific stereogram density, a non-iterative version of our algorithm can recover disparity satisfactorily (see Fig. 14a iteration 1). John Fairfield (personal communication) suggested an algorithm in which 1) excitation is summed independently within each disparity layer, and 2) for each position, one selects only the most excited of the cells in the different disparity layers. This algorithm performs well for the case $v=0.5$.

7.2. Comments on Analyzing Such Operations

We find the style of analysis that we were forced to adopt to be unsatisfactory for a number of reasons. Firstly, although our arguments appear to provide a qualitatively accurate description of the algorithms' behavior, the arguments are not completely rigorous. The main reasons for this lie in the difficulty of assessing the validity of the randomness assumptions that are necessary for the probabilistic analysis; and, to a lesser extent, in the need to examine a number of special cases in order to establish the stability of various solutions.

Secondly, our analysis is very specific to the particular algorithm and the particular parameters. This style of proof cannot lead to any general results about the convergence of such operators.

In order to overcome the first of these problems one can follow the approach of Section 6. The price one pays is that the analysis is valid for a narrower parameter range, which happens not to include the




Fig. 15. Conway's game "life", which is played on an infinite plane square lattice, may be represented in a manner very similar to that of our stereo operator. The excitatory neighborhood, together with appropriate weights, is shown in 15a, and the threshold function appears in 15b. This combination reproduces the rules of life exactly, and these are; 1. A cell will die at generation $n+1$ if <2 or >3 of its 8 neighbors are alive at generation n (death by starvation or overfeeding). 2. A cell with exactly 2 living neighbors at generation n will be alive at generation $n+1$ if and only if it is alive at generation n . 3. A cell with exactly 3 living neighbors at generation n will be alive at generation $n+1$

original parameter values (see Marr and Poggio, 1976). The difficulty with the assumption of randomness arises because of the constant spatial structure of the operator Ξ (Eq. (1) and Fig. 2c of Marr and Poggio, 1976). It should perhaps be noted that this objection does not apply to the similar analysis given by Marr (1971) of a cooperative associative memory algorithm, because there the local operator had a variable and essentially random structure.

The second of these difficulties seems to be inherent in the nature of this type of cooperative algorithm. No general approach is at present available. Standard approaches⁴ that we have tried have failed up to now. The flavor of the difficulties is the following. A configuration that is stable may be perturbed by changing a large number of cells without affecting its asymptotic state, provided that the perturbed cells are well scattered and interior. On the other hand, one fixed point

⁴ The continuous version of the algorithm (1) cannot be described in terms of a potential dynamics. In fact the dynamical system

$$C = \sigma\{\Xi C\} - C = f(C) \quad (\sigma \text{ a "smooth" threshold})$$

does not admit a scalar potential function $V(C)$ such that

$$[f(C)]_{x,y,d} = \partial V(C) / \partial C_{x,y,d}.$$

A necessary condition for this to be true is that

$$\frac{\partial}{\partial C_{x,y,d}} f_{x',y',d'}(C) = \frac{\partial}{\partial C_{x',y',d'}} f_{x,y,d}(C), \quad \text{all } x, y, d, x', y', d'.$$

This is not true in general, because of the nonlinearity σ (consider the case in which x,y,d and x',y',d' are on the same disparity layer and are reciprocally excitatory)

of the algorithm can be shifted into another by perturbing only a few cells, provided that they have a suitable configuration. Thus, the usual distance between two configurations, namely the number of cells having different states, does not reflect the behavior of the algorithm. Therefore, the problem seems to be how to incorporate the geometry of the interactions into the metric distance between configurations.

It seems unlikely that one can construct a useful general theory of algorithms of the form

$$C^{n+1} = \sigma\{L(C^n)\}, \quad (15)$$

where L is a linear operator on the vector C , and σ is a nonlinear (coordinate-wise) function. J. H. Conway's game "Life" can, for example, be written this way (see Fig. 15) and with an appropriate input pattern is Turing universal (unpublished result discovered independently by J. H. Conway and R. W. Gosper).

This suggests that theories of this type of algorithm must take due account of the structure of the input data and will probably be restricted to very specific forms of (15).

A mathematical understanding of the behavior of (15) would represent a breakthrough of rather general importance. Cooperative phenomena similar to those which can be described by (15) are important in physics (Haken, 1977; Kawasaki, 1972; K. G. Wilson, 1975), in development (Mostow, 1975), and in biology (Eigen, 1971; Marr, 1971; Richter, 1976).

Furthermore, such a theory might also allow one to synthesize in a standard way cooperative algorithms of the form of (15) from an analysis of the constraints on a computation.

Acknowledgements: We thank K. P. Haderl and W. Reichardt for interesting discussions, Karen Prendergast for preparing the illustrations, and the Matlab Group for allowing us to run the probabilistic theory in Macsyma, the M.I.T. symbolic algebraic manipulation system. Pirooz Vatan helped with Macsyma programming. Science kindly gave permission for the reproduction of Figure 1. This work was conducted at the Artificial Intelligence Laboratory, a Massachusetts Institute of Technology research program supported in part by the Advanced Research Projects Agency of the Department of Defense, and monitored by the Office of Naval Research under contract number N00014-75-C-0643. D. M. thanks the Max-Planck-Institut für biologische Kybernetik in Tübingen for its kind hospitality during his visit there.

References

- Eigen, M.: Self-organization of matter and the evolution of biological macromolecules. *Naturwissenschaften* **10**, 465—523 (1971)
- Haken, H.: *Synergetics*. Berlin-Heidelberg-New York: Springer 1977
- Kawasaki, K.: Kinetics of Ising models. In: *Phase transitions and critical phenomena*, Vol. 2, pp. 443—501. Domb and Green eds., New York: Academic Press 1972
- Marr, D.: Simple memory: s theory for archicortex. *Phil. Trans. Roy. Soc. B* **262**, 23—81 (1971)

- Marr, D.: A note on the computation of binocular disparity in a symbolic, low-level visual processor. M.I.T. A.I. Lab. Memo 327, (1974)
- Marr, D.: Early processing of visual information. Phil. Trans. Roy. Soc. B **275**, 483—524 (1976)
- Marr, D.: Representing visual information. AAAS 143rd Annual Meeting, Symposium on Some Mathematical Questions in Biology, February, (in press). Also available as M.I.T. A.I. Lab. Memo 415 (1977)
- Marr, D., Poggio, T.: Cooperative computation of stereo disparity. Science **194**, 283—287 (1976)
- Mostow: Mathematical models for cell rearrangement. Newhaven, Yale University Press 1975
- Richter, P.H.: The network idea and the immune response. In: Theoretical Immunology, Bell, G. I., Perelson, A.S., Pimbley, G.H., eds. New York: M. Dekker 1976
- Wilson, H.R.: Hysteresis in binocular grating perception: contrast effects. Vision Res. (1977) (in press)
- Wilson, H.R., Cowan, J.D.: A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. Kybernetik **13**, 55—80 (1973)
- Wilson, K.G.: The renormalization group: Critical phenomena and the Kondo problem. Reviews Mod. Phys. **47**, 773—840 (1975)

Received: November 28, 1977

Dr. T. Poggio
MPI für biolog. Kybernetik
Spemannstr. 38
D-7400 Tübingen
Federal Republic of Germany

