BIL 4:2

UPPSALA UNIVERSITY
DEPARTMENT OF COMPUTER SCIENCES
REPORT NR 10
SEPTEMBER, 1967

An environment for an artificial intelligence

by

Erik Sandewall

The research reported here was partly supported by the Swedish Natural Science Research Council (Dnr 2711-2 ans 6826) and by the Advanced Research Projects Agency of the Office of the U.S. Secretary of Defence (SD-183).

0. Introduction

One of the objectives of artificial intelligence research is the creation of a robot which is able to behave intelligently in the real world. The robot must e.g. be able to predict, through logical inference, the consequences of its actions in the environment; and by the use of that ability, it must be able to take the right actions to attain given goals. It should also be able to communicate with a human; to state its beliefs about the environment; and to make use of advice (for how to attain given goals) that the human may give. In the initial stages of work towards such a robot, it may be desirable, or even necessary to let the "intelligence" "live" in an artificial universe, which is simulated inside the computer. The reasons are obvious: first we need not worry about the hardware interface between the computer and the real world, and second, an artificial universe can be made very simple.

In this memo, we shall describe a class of artificial universes, which may be used as play-grounds for primitive robot-type programs. Such artificial universes would be used by programs that consist of two parts: one simulator of the universe and one agent which "lives" in the universe: observes it, and takes actions in it. There must be a two-way communication between these parts: facts about the universe are sent to the agent, and conversely, the simulator is sometimes influenced by the output of the agent.

Our artificial universes are designed to have some features that are characteristic of the physical world, namely, space and time. This would make the work more significant for real-world robots. The simulator operates, therefore, on a two-dimensional picture which is updated in discrete steps of time. Between each updating, the agent would be permitted to do a certain, restricted quantity of reasoning (this quantity should be so small that the agent is not able to figure out the environment completely) and to indicate its actions, if any.

The class of universes is so wide that the complexity, and therefore, the difficulty of the agent's environment can be varied within wide limits.

The system can be fancied up in many ways. For example, we can decide to assign a position in the picture to the agent. The picture would contain some component which "is" the representation of the agent. The agent can then perform actions only in the part of the picture where it (its representation) "is", but this includes an ability to move "itself" around.

This memo is only concerned with the selection of the environment, and we do not even begin to consider the important problems in the construction of an agent that lives in the environment.

1. Proposal for the environment

It is desirable that the environment should operate in discrete steps of space and time, because this makes it easier to manipulate. The development on the board in a game like chess or checkers is an example of such a world. However, the problem in a board game is not very similar to the simplest of the problems a robot encounters. The robot's world is mechanistic: the robot tilts the stool, so that it falls, so that the box on the stool falls to the floor. On the board, there are no such chains of cause and effect: a piece is moved if and only if a player moves it. The robot's world is impartial: the robot may take several units of time to turn around in order to get through a gate, and the gate will still be there. On the board, everything that happens must be understood from the game's character of encounter, and all plans for action must consider the complication that the opponent will try to obstruct the plan.

It is true that artificial intelligences will eventually be required to face an opponent (at least in some applications), but it seems wise to give the agent the favor of a non-hostile environment to begin with. This favor can be withdrawn later, after techniquea developed from work with game-playing programs have been incorporated into the agent.

For these reasons, we select to use a cellular automaton (originally defined by von Neumann, /1/) as a universe. A cellular automaton is an infinite checkerboard, where at each point in time, each square (cell) is in one of a number of states. There is a next-state function which determines the state of a square s at time t+1 from the states of s and its neighbors at time t.

We can obtain a finite subset of the board to work with by "fencing in" a part of the board with a "wall" of cells whose state is constant and independent of the states of its neighbors. Everything that happens inside this wall will then be independent of anything that happens outside. We shall call the area inside the wall, the picture.

A variant of the cellular automaton would have a next-state function which operates in a fashion similar to the Gauss-Seidel algorithm, i.e. which considers the states of some neighbors at time t, and of some neighbors (already updated) at time t+1. With this variant, we can create effects like having a "beam of light" extend across the whole board from one step in time to the next. For its greater richness, we shall prefer to use this variant.

What neighbors are to be used by the next-state function? Obviously, at least for some states s, some neighbors must have an effect, or the history of each square would be independent from that of the others. It might seem natural to select the neighbors immideatily above, below, to the left, and to the right, according to the following spider:

```
 x
xxx
 x
```

Using the Gauss-Seidel variant here would give a next-state function on the form

$$s(i,j,t+1) =$$

$$F(\ s(i,j-1,t+1),\ s(i-1,j,t+1),\ s(i,j,t),\ s(i+1,j,t),\ s(i,j+1,t)\ )$$

where $s(i,j,t)$ is the state of the square s_{ij}(*) at time t. However, this next-state function would destroy the left-right symmetry (for instance, an instantaneous beam could be sent upward or to the right, but not to the left or downward). The alternative is to select neighbors by the spider

```
x x
 x
x x
```

which corresponds to a next-state function like

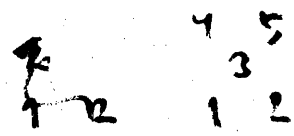

$$s(i,j,t+1) =$$

$$F(\ s(i-1,j-1,t+1),\ s(i+1,j-1,t+1),\ s(i,j,t),\ s(i-1,j+1,t),\ s(i+1,j+1,t))$$

For example, suppose 0 and 1 are states, and F satisfies

$$(x)(y)(z)\ F(x,1,0,y,z) = 1.$$

On a board that consists essentially of 0's one single square in state 1 at time t will cause, at time t+1, a beam of 1's which extends towards "north-west" at least until it reaches a square that was not in state 0 at time t.

Suppose for a moment that the squares on the board are painted black and white checkerboard style. With the next-state function operating on an X-formed spider, anything that happens on black squares is independent of anything that happens on white squares.

(*) with x-coordinate i and y-coordinate j.

Therefore, we can let the program restrict its attention to one of those two sets of squares.

Thus our proposal for a class of universes for the AI program is: the history of the "black" squares of a finite subset of a cellular automaton, "Gauss-Seidel" variant, whose next-state function relies on the states in an X-formed spider.

What actions shall we permit the agent to take? The crudest choice is to give the agent the power of changing the state of at most one square (of its own choice), at each point in time. A small refinement is to restrict the "reach" of the agent to the immideate vicinity of some movable vehicle in the picture. In these cases, the action of the agent completely over-rides the next-state function. It is possible that, for some purposes, we will desire to let the next state be dependent on both the previous and surrounding states, and the output of the agent, so that the output signal of the agent enters F as a sixth argument.

In the rest of this memo, we shall:

1. Describe one particular universe, which might be a suitable environment for an agent;
2. Describe a second cellular-automaton universe, which is <u>not</u> a suitable environment for an agent, and discuss the criteria for a good environment.

2. Desciption of one universe

We shall describe a very simple universe with four states (*). The sceneries in this universe can be visualized in terms of "signals" which can slide "down" (i.e. toward south-west or south-east in the picture) along "paths"; which can "meet" and "split" at branchings of and intersections between these paths; and which can vanish when the end of the path is reached. The universe shall have a complete left-right symmetry.

The next-state function g is of this universe is defined in a COMIT-like notation. Let us define a <u>transformation</u> as an expression of the form

$$\begin{matrix} d & & e \\ & c & \\ a & & b \end{matrix} \quad = \quad n$$

Such a transformation shall mean $g(a,b,c,d,e) = n = g(b,a,c,e,d)$. A transformation can be <u>applied</u> to square (i,j) at time t if either of the following two conditions hold:

$$(1) \quad \begin{cases} s(i-1,j-1,t+1) = a, \text{ \underline{and}} \\ s(i+1,j-1,t+1) = b, \text{ \underline{and}} \\ s(i,\ j,\ t) = c, \text{ \underline{and}} \\ s(i-1,j+1,t) = d, \text{ \underline{and}} \\ s(i+1,j+1,t) = e \end{cases}$$

(2) The transformation $\begin{matrix} e & & d \\ & c & \\ b & & a \end{matrix} \quad = \quad n$ can be applied to the square (i,j) at time t according to condition (1).

(*) besides the state that makes up the "wall".

The next-state function g is defined by a <u>sequence</u> of transformations. s(i,j,t+1) shall be the right-hand side of the first transformation of the sequence, which can be applied to the square (i,j) at time t.

We use the symbols +, =, and ¢ as variables in the transformations. If one variable occurs several times in one transformation, it need not be matched against the same state in all positions. Thus the transformation

```
¢ ¢
 o     ==-   o   (where o and s are states) can be
¢ ¢
```

specialized to the case

```
s o
 o     ==-   o .
s s
```

The four states are

o (the "blank state")

x (which forms the paths)

s (which is used as a signal)

e (an "excited" state, used e.g. when a signal changes direction).

The next-state function g is defined by the following sequence of transformations:

$$\begin{array}{ccc} \not c & & \not c \\ & o & \\ \not c & & \not c \end{array} \quad \text{==-} \quad o$$

$$\begin{array}{ccc} s & & s \\ & x & \\ \not c & & \not c \end{array} \quad \text{==-} \quad x$$

$$\begin{array}{ccc} \not c & & s \\ & x & \\ + & & \not c \end{array} \quad \text{==-} \quad s \qquad \text{if + stands for x or s}$$

$$\begin{array}{ccc} \not c & & s \\ & x & \\ o & & \not c \end{array} \quad \text{==-} \quad e$$

$$\begin{array}{ccc} \not c & & \not c \\ & x & \\ \not c & & \not c \end{array} \quad \text{==-} \quad x$$

$$\begin{array}{ccc} \not c & & \not c \\ & s & \\ e & & + \end{array} \quad \text{==-} \quad x \qquad \text{if + stands for x or s}$$

$$\begin{array}{ccc} \not c & & \not c \\ & s & \\ e & & \not c \end{array} \quad \text{==-} \quad e$$

$$\begin{array}{ccc} \not c & & \not c \\ & s & \\ \not c & & \not c \end{array} \quad \text{==-} \quad x$$

$$\begin{array}{ccc} + & & = \\ & e & \\ o & & o \end{array} \quad \text{==-} \quad e \qquad \text{if = and + are both different from o, and at least one of them is x.}$$

$$\begin{array}{ccc} \not c & & \not c \\ & e & \\ \not c & & \not c \end{array} \quad \text{==-} \quad s$$

On the following pages, a number of computer-produced histories in this universe are displayed to show main things that can happen. The program (which was written in LISP) is available from the author. - In the figures, the state o is represented as ".".

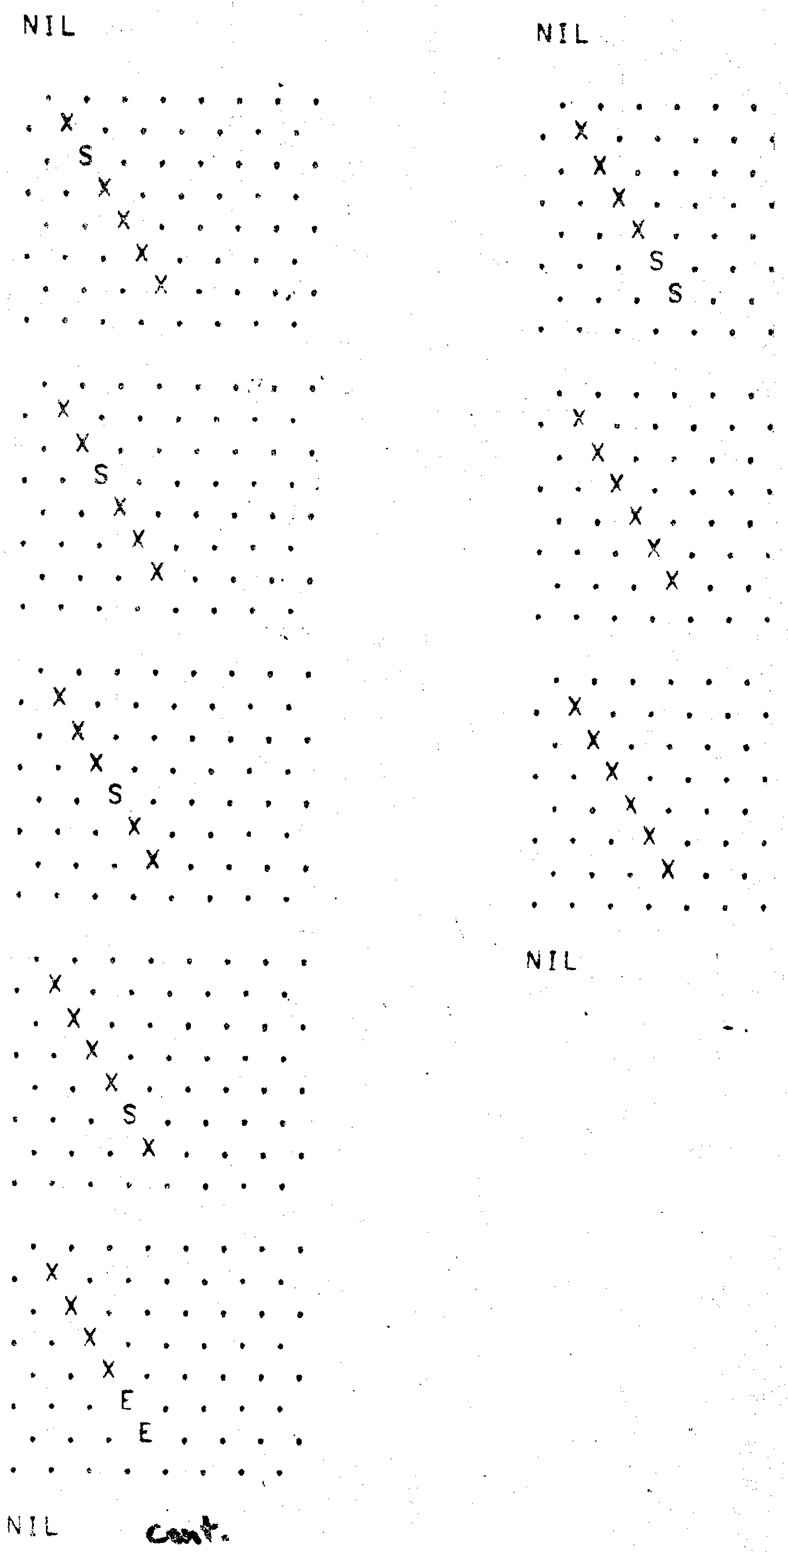

Read first the leftmost column from top to bottom, then the second column.

The signal travels downward along the path, one level for each step in time.

When it reaches the end of the path, it goes into the excited state (it is dangling like a drop) but eventually falls off.

A path with no signal is of course constant in time.

Due to the symmetry of the next-state function, all snapshot sequences can be mirrored along a vertical axis.

```
NIL

 . . . . . . . .
. . X . X . . .
 . . S . S . . .
. . . X . X . .
 . . . X . X . .
. . . X . . X .
 . . X . . . X .
. . X . . . . X

 . . . . . . . .
. . X . X . . .
 . . X . X . . .
. . . S . S . .
 . . . X . X . .
. . . X . . X .
 . . X . . . X .
. . X . . . . X

 . . . . . . . .
. . X . X . . .
 . . X . X . . .
. . . E . X . .
 . . . E . S . .
. . . X . . X .
 . . X . . . X .
. . X . . . . X

 . . . . . . . .
. . X . X . . .
 . . X . X . . .
. . . S . X . .
 . . . S . X . .
. . . X . . S .
 . . X . . . X .
. . X . . . . X

 . . . . . . . .
. . X . X . . .
 . . X . X . . .
. . . X . X . .
 . . . X . X . .
. . . S . . X .
 . . X . . . S .
. . X . . . . X

NIL
```

In a "curve", the signal goes into the "excited" state E; then goes back to S and proceeds. As compared to the signal that goes the straight path, the signal ~~is~~ is delayed one time unit in the curve.

In the ~~xxxxx~~ agent's thinking, the overall view will very often be sufficient. It is important that the agent shall be able to look away from the "local disturbances" of e.g. excitation into E state.

```
NIL

 . . . . . . . .
. . X . X . . .
 . . S . S . . .
. . . X . X . .
 . . . X . X . .
. . . X . . X .
 . . X . . . X .
. . X . . . . X

 . . . . . . . .
. . X . X . . .
 . . X . X . . .
. . . S . S . .
 . . . X . X . .
. . . X . . X .
 . . X . . . X .
. . X . . . . X

 . . . . . . . .
. . X . X . . .
 . . X . X . . .
. . . E . X . .
 . . . E . S . .
. . . X . . X .
 . . X . . . X .
. . X . . . . X

 . . . . . . . .
. . X . X . . .
 . . X . X . . .
. . . S . X . .
 . . . S . X . .
. . . X . . S .
 . . X . . . X .
. . X . . . . X

 . . . . . . . .
. . X . X . . .
 . . X . X . . .
. . . X . X . .
 . . . X . X . .
. . . S . . X .
 . . X . . . S .
. . X . . . . X

NIL
```

In a "curve", the signal goes into the "excited" state E; then goes back to S and proceeds. As compared to the signal that goes the straight path, the signal is delayed one time unit in the curve.

In the agent's thinking, the overall view will very often be sufficient. It is important that the agent shall be able to look away from the "local disturbances" of e.g. excitation into E state.

NIL

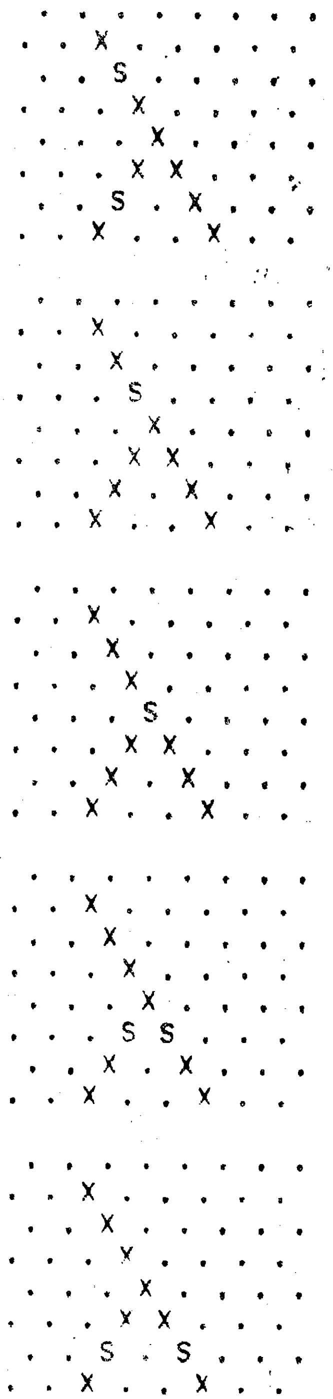

NIL

NIL

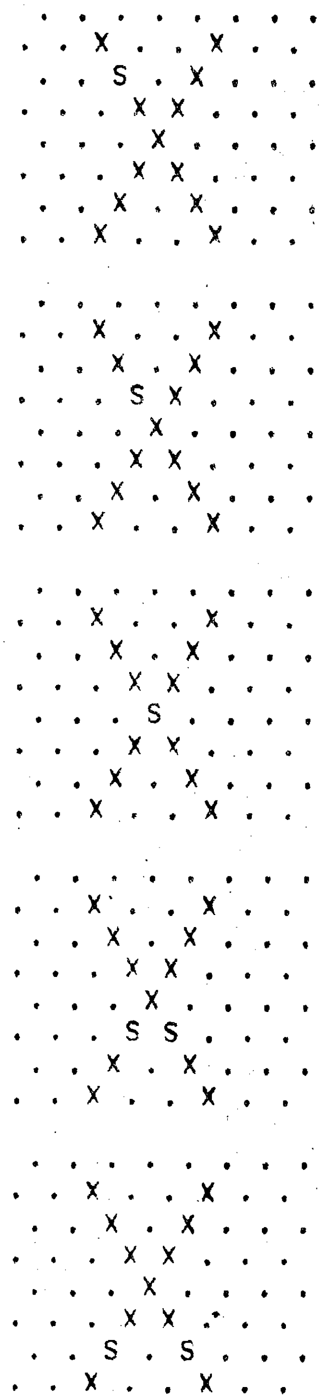

NIL

In simple "intersections" like these, one signal will split into two.

The lower signal in the first picture to the left, vanishes prematurely because it has hit the "wall" around the picture. (In the program that produced these pictures, the boundary conditions are handled somewhat non-strictly).

NIL

```
 . . . . . . . .
. . X . . X . .
 . . S . S . . .
. . . X X . . .
 . . . X . . . .
. . . X X . . .
 . . X . X . . .
. . X . . X . .

 . . . . . . . .
. . X . . X . .
 . . X . X . . .
. . . S S . . .
 . . . X . . . .
. . . X X . . .
 . . X . X . . .
. . X . . X . .

 . . . . . . . .
. . X . . X . .
 . . X . X . . .
. . . X X . . .
 . . . X . . . .
. . . X X . . .
 . . X . X . . .
. . X . . X . .

 . . . . . . . .
. . X . . X . .
 . . X . X . . .
. . . X X . . .
 . . . X . . . .
. . . X X . . .
 . . X . X . . .
. . X . . X . .
```

NIL

Two signals that arrive simultaneously, annihilate.

NIL

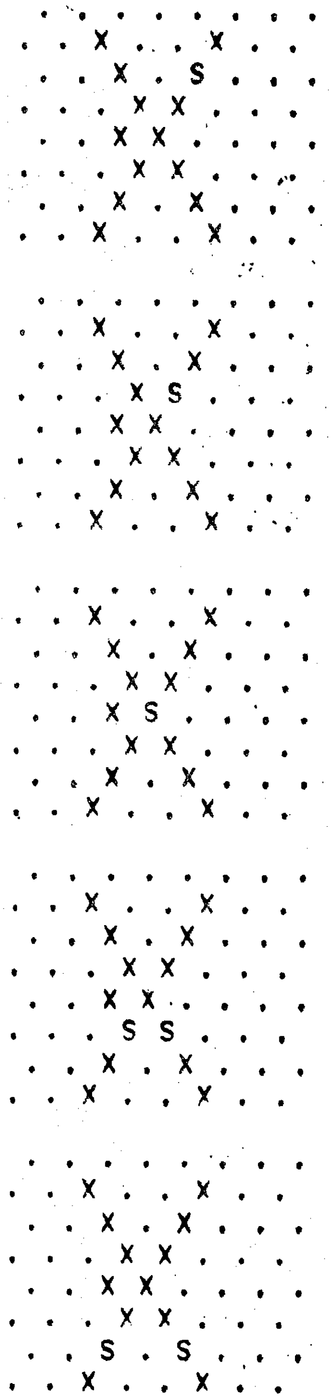

NIL

NIL

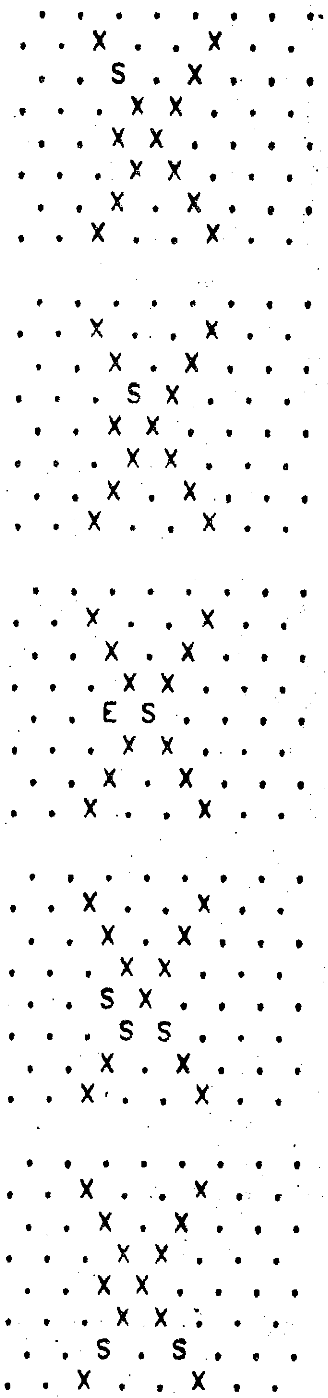

NIL

This slightly more complicated intersection lets through signals (from either direction!) in grossly the same way as the simple intersectic

There are some "local disturbances"in the right-hand column, but there is no delay when the signal passes through the inter-section.

NIL

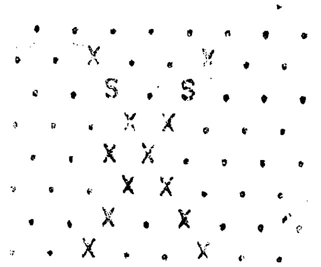

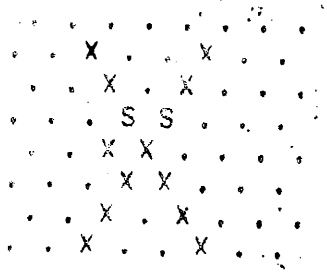

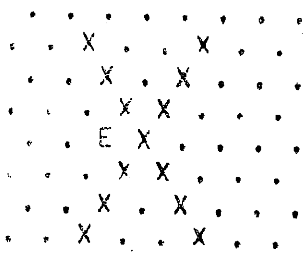

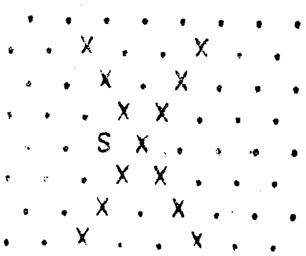

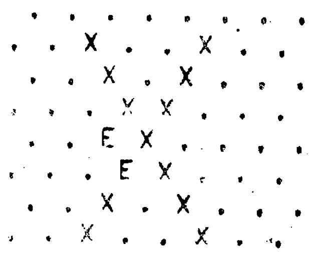

NIL

Cont.

NIL

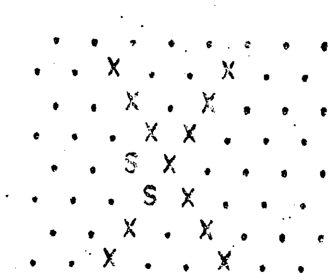

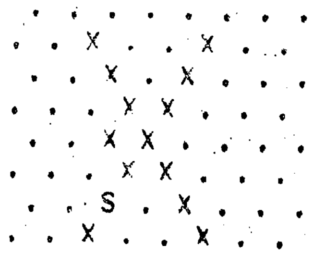

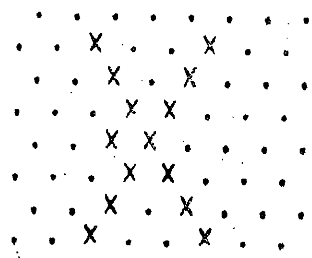

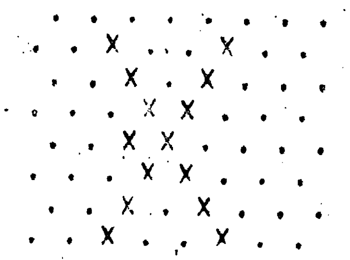

NIL

The same intersection does function differently when two signals meet in the intersection: In the simple intersection, they would have annihilated.

NIL

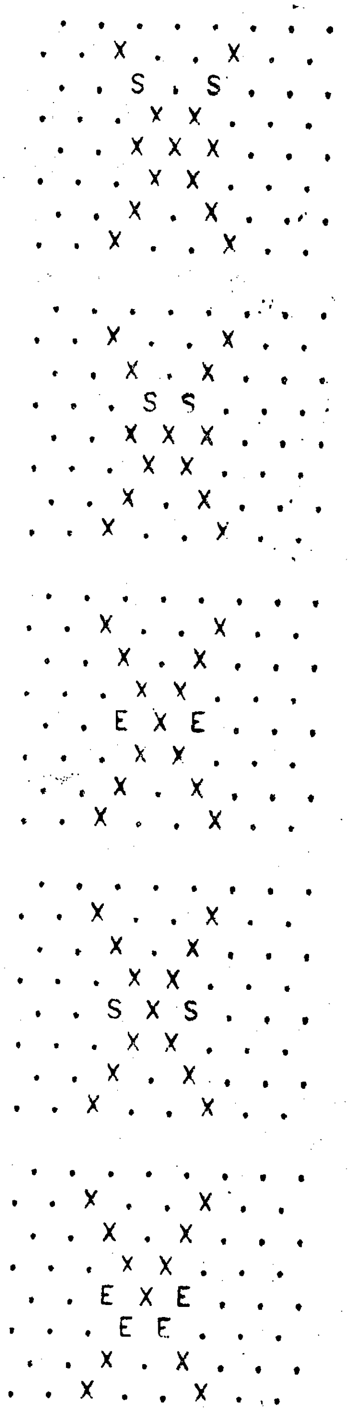

NIL

Cont.

NIL

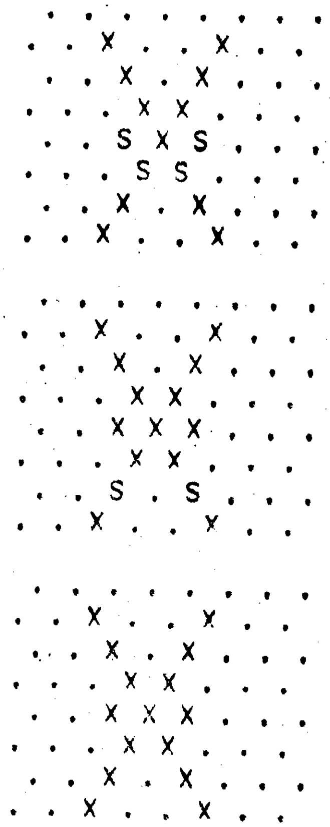

NIL

Putting <u>two</u> such side-paths in the same intersection, we obtain the same process on both sides, as we saw on one side on the previous sheet.

NIL

```
 . . . . . . . .
. S . . . . X .
 . X . . . X . .
. . X . . X . .
 . . X . X . . .
. . . X X . . .
 . . . X . . . .
. . . . . . . .
```

```
 . . . . . . . .
. X . . . . X .
 . S . . . X . .
. . X . . X . .
 . . X . X . . .
. . . X X . . .
 . . . X . . . .
. . . . . . . .
```

```
 . . . . . . . .
. X . . . . X .
 . X . . . X . .
. . S . . X . .
 . . X . X . . .
. . . X X . . .
 . . . X . . . .
. . . . . . . .
```

```
 . . . . . . . .
. X . . . . X .
 . X . . . X . .
. . X . . X . .
 . . S . X . . .
. . . X X . . .
 . . . X . . . .
. . . . . . . .
```

```
 . . . . . . . .
. X . . . . X .
 . X . . . X . .
. . X . . X . .
 . . X . X . . .
. . . S X . . .
 . . . X . . . .
. . . . . . . .
```

NIL cont.

NIL

```
 . . . . . . . .
. X . . . . X .
 . X . . . X . .
. . X . . X . .
 . . X . X . . .
. . . E X . . .
 . . . E . . . .
. . . . . . . .
```

```
 . . . . . . . .
. X . . . . X .
 . X . . . X . .
. . X . . X . .
 . . X . X . . .
. . . S X . . .
 . . . E . . . .
. . . . . . . .
```

```
 . . . . . . . .
. X . . . . X .
 . X . . . X . .
. . X . . X . .
 . . X . X . . .
. . . E X . . .
 . . . E . . . .
. . . . . . . .
```

```
 . . . . . . . .
. X . . . . X .
 . X . . . X . .
. . X . . X . .
 . . X . X . . .
. . . S X . . .
 . . . E . . . .
. . . . . . . .
```

NIL

A signal falling into a "ditch" creates an oscillating structure: a constant E in the bottom, and an oscillating S-E above it.

NIL

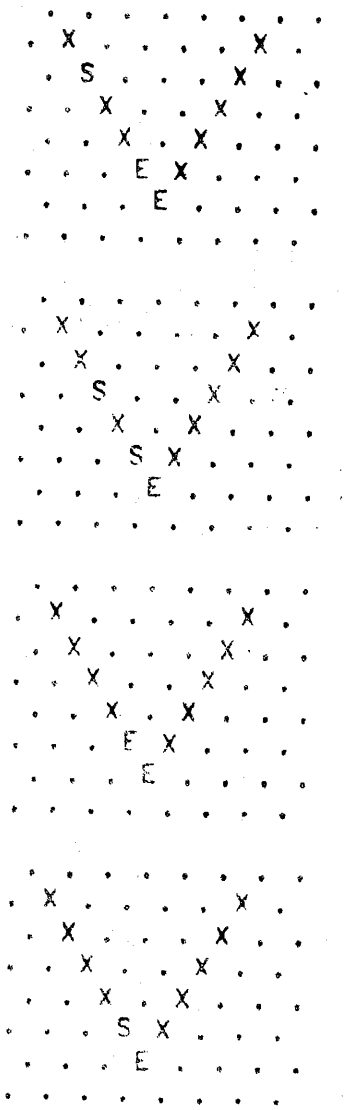

NIL

NIL

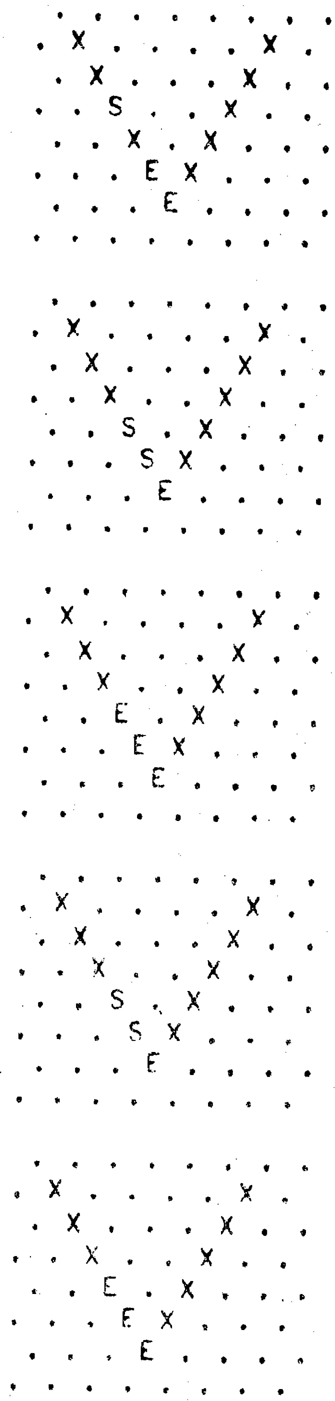

NIL

The S-E oscillator can be extended by another signal, which comes in in the right phase. In the left-hand column, the signal comes wrong and is lost.

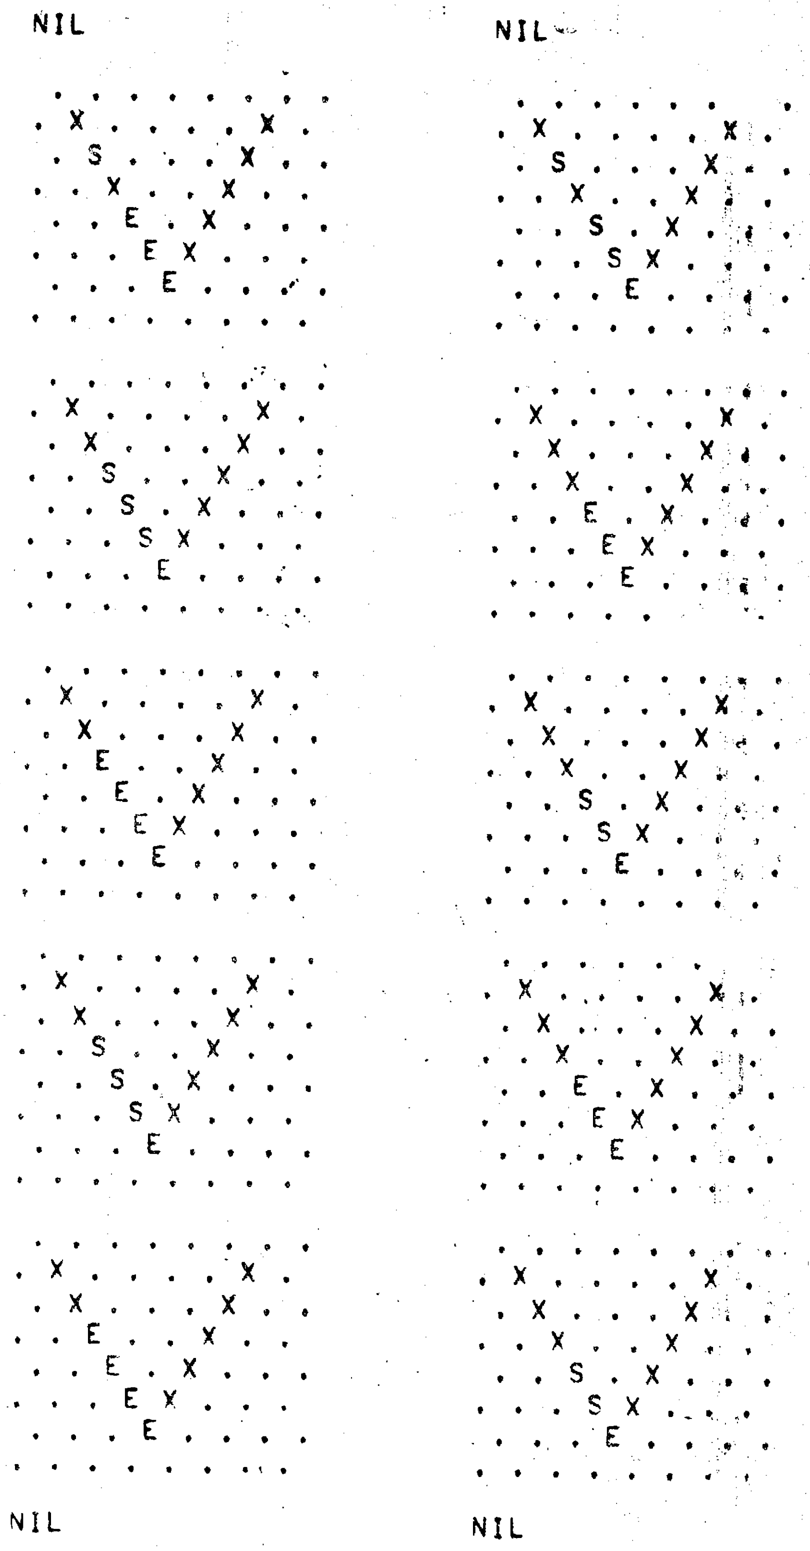

The same phenomenon appears again. Oscillating stacks of indefinite length can be built.

NIL

```
 . . . . . . . .
. X . . . . S .
 . X . . . X . .
. . S . . X . .
 . . S . X . . .
. . . S X . . .
 . . . E . . . .
. . . . . . . .
```

```
 . . . . . . . .
. X . . . . X .
 . X . . . S . .
. . E . . X . .
 . . E . X . . .
. . . E X . . .
 . . . E . . . .
. . . . . . . .
```

```
 . . . . . . . .
. X . . . . X .
 . X . . . X . .
. . S . . S . .
 . . S . X . . .
. . . S X . . .
 . . . E . . . .
. . . . . . . .
```

```
 . . . . . . . .
. X . . . . X .
 . X . . . X . .
. . E . . X . .
 . . E . S . . .
. . . E X . . .
 . . . E . . . .
. . . . . . . .
```

```
 . . . . . . . .
. X . . . . X .
 . X . . . X . .
. . S . . X . .
 . . S . X . . .
. . . S X . . .
 . . . E . . . .
. . . . . . . .
```

NIL

NIL

```
 . . . . . . . .
. X . . . . S .
 . X . . . X . .
. . E . . X . .
 . . E . X . . .
. . . E X . . .
 . . . E . . . .
. . . . . . . .
```

```
 . . . . . . . .
. X . . . . X .
 . X . . . S . .
. . S . . X . .
 . . S . X . . .
. . . S X . . .
 . . . E . . . .
. . . . . . . .
```

```
 . . . . . . . .
. X . . . . X .
 . X . . . X . .
. . E . . S . .
 . . E . X . . .
. . . E X . . .
 . . . E . . . .
. . . . . . . .
```

```
 . . . . . . . .
. X . . . . X .
 . X . . . X . .
. . S . . X . .
 . . S . S . . .
. . . S X . . .
 . . . E . . . .
. . . . . . . .
```

```
 . . . . . . . .
. X . . . . X .
 . X . . . X . .
. . E . . X . .
 . . E . X . . .
. . . E X . . .
 . . . E . . . .
. . . . . . . .
```

Irrespective of phase, a signal coming from the other direction has no influence.

3. A second example, and the restriction to entropic universes.

Let us consider the following simple cellular automaton of the "ordinary" kind (not the Gauss-Seidel variant):

States: 0,1

Next-state function: h(a,b,c,d,e,) =

if c = 1 or a+b+d+e is odd then 1 else 0

If a square has once attained the state 1, it will stay there. In this world, therefore, any initial figure of 1's will grow larger and larger, and develop intricate patterns (see illustrations following page).

This universe (h) is different from the universe (g) of the previous section (with next-state function g) in several ways:

(1) Everything interrelated. In the g universe, it is possible to have in the same picture, several signals with some distance between them; and these signals will then mainly be unrelated, so that they can be studied one at a time. (This is especially possible if we use a larger picture than in the above examples). In the h universe, on the other hand, changes of state take place along the whole circumference of the growing corpus, and it does not seem possible to isolate semi-independent subsystems in the picture.

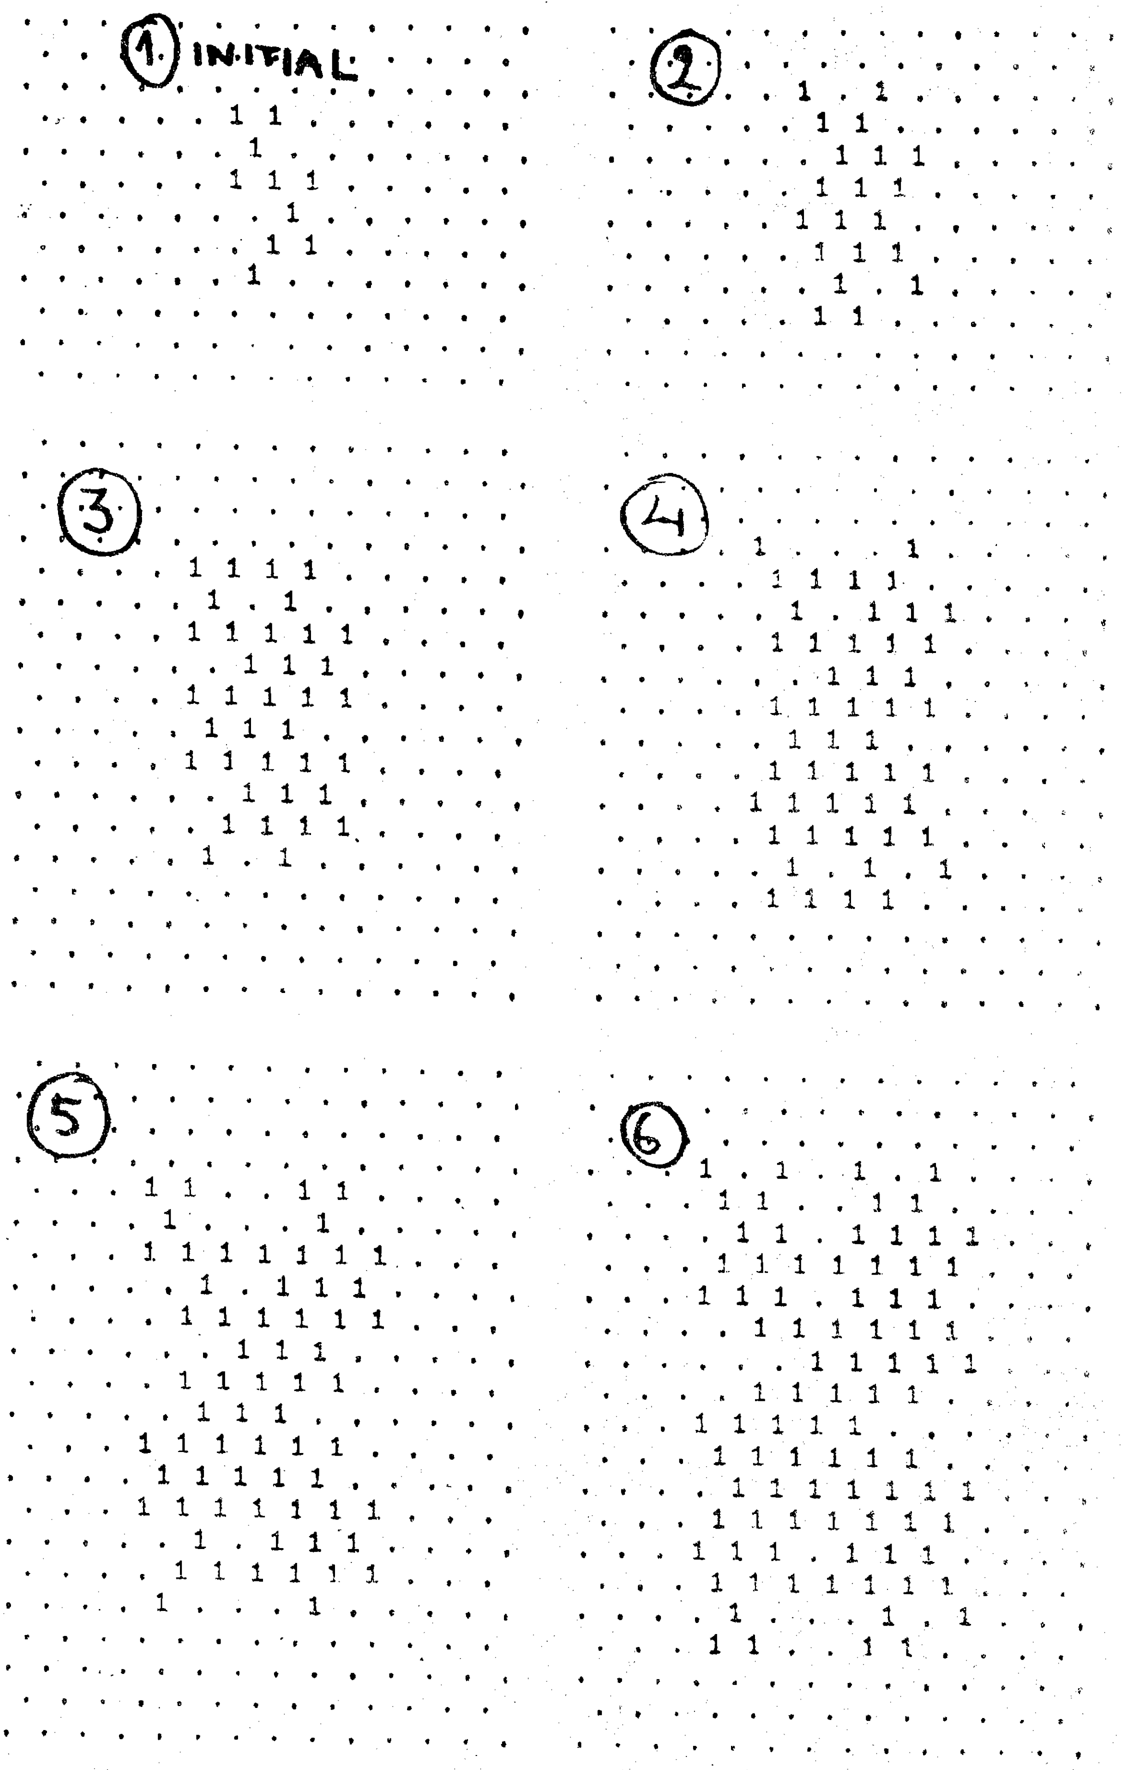
1. INITIAL
2
3
4
5
6

(2) <u>No entropy</u>. In the g universe, new "independent subsystems of change of state" can occasionally be created (e.g. when a signal branches in an intersection), but such subsystems will also die off (e.g. when a signal reaches the bottom of a path). Therefore, there is no risk that the whole picture shall eventually be covered by activities. In the h universe, there is no such "entropy". One disturbance in the area where growth takes place (which is the interesting area) may be able to perpetuate itself indefinitely.

(3) <u>Low predictability</u>. The simple, characteristic facts about the g universe (signals follow paths; signals go downwards with constant speed) can be used for crude guesses about the future effects of possible actions. We must of course have more facts for a good description (e.g. "signals are delayed one time unit in curves"), but the simplest facts can be used, e.g. to eliminate all those suggested actions for attaining a given goal, which are clearly irrelevant to the goal. In h, there do not seem to be such broad rules: one small disturbance in the growth area may have farreaching effects, but the exact character of the effects depends on the circumstances.

In all these three aspects, universe h seems to offer more difficulties to the agent than does universe g. Furthermore, as g type universes are easier to conceptualize and to reason about, it can well be claimed that they are more relevant to our work.

Let us introduce the adjective entropic for a universe which houses reasonably predictable subsystems of change of state; subsystems which do not interrelate too heavily; which do not breed too fast, and which die off fast enough to keep the picture clean and easy to understand. Clearly, entropic universes are the ones that are the most suitable as playgrounds for artificial intelligences.

References

/1/ von Neumann, J.

Theory of Automata: Construction, Reproduction, Homogeneity

(Part II of

Burks, A.W. (ed)

The Theory of Self-Reproducing Automata

University of Illinois Press)