# Chapter 8
# Composing Behavior: Registers for Plans and Actions

In the conduct of everyday human affairs, the preference of one course of action over another often depends on the perception that the course chosen will tend to maximize or minimize some significant aspect of experience. For example, the decision to wear a coat instead of a light jacket on a fall day might depend on the fact that it is unseasonably cold and wearing a jacket will result in greater discomfort than wearing a coat. The brain must be able to encode and register as neuronal tokens the relative magnitudes of a great variety of experiential parameters, and it must have the capability to recall such tokens as criteria for the selection of episodic adaptive actions—integrated sequences of behavior that can be characterized as the execution of a plan (Miller, Galanter, and Pribram 1960).

## *Encoding Magnitude*

Figure 8.1 shows a schematic representation of a simple neuronal mechanism that receives a pulse of axonal discharge as its input and systematically converts the spike frequency of the input pulse into the discharge of a particular autaptic cell in an ordered string of such cells. This mechanism operates on the same general principles as do the retinoids. In this case, however, the input pulse is of a standard and fixed duration, although its constituent spike frequency will vary according to the intensity of the event that it represents.

The autaptic cell marked 0 is constantly active. Because of the dynamic properties of this kind of neuronal circuit, the extent of excitational translation along the autaptic chain will be proportional to the frequency of the spikes that constitute the input pulse. At the end of the pulse, the ordinal value of the autaptic cell that remains active (in addition to the origin cell 0) will depend directly on the spike frequency of the input pulse. If the input frequency represents the magnitude of an experiential parameter (time or effort to complete a task, for example), then this neuronal analog will be systematically encoded as a particular labeled line (indexed autaptic cell) in an or-

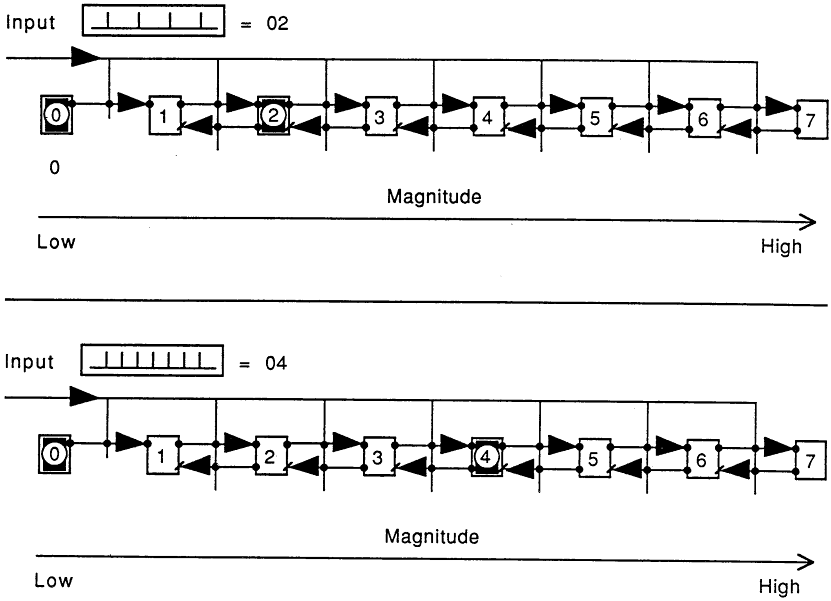


Figure 8.1
A mechanism for encoding magnitude. The degree or intensity of an event that is initially represented by the frequency of neuronal discharge (spike train in inset) is converted to a representation by place in an ordered set of autaptic cells. Each autaptic cell is a labeled line that signals relative magnitude.

dered set. After this transformation, the labeled autaptic line, which is a token (in short-term memory) of the magnitude of its parameter, can provide input to a synaptic matrix to be learned (in long-term memory) as one of the properties of the context in which it was activated.

### *Binding Consequences to Actions*

Chapter 7 gave an example in which two different routes were evaluated for the distance traveled to accomplish the same result. Figure 8.2 illustrates how route and distance can be bound together in a synaptic matrix and represented as a single frame of knowledge that is activated in parallel. The saturation limit (*Lim*) for filter cell transfer weights ($\phi_i$) was arbitrarily set at *Lim* = 2. There are two separate but concurrent sources of afferent excitation to the matrix: a labeled-line token of a particular self-locus excursion (route), either current or recalled from memory, and a token output from the autaptic chain that encodes the overall distance of the route. After route and dis-

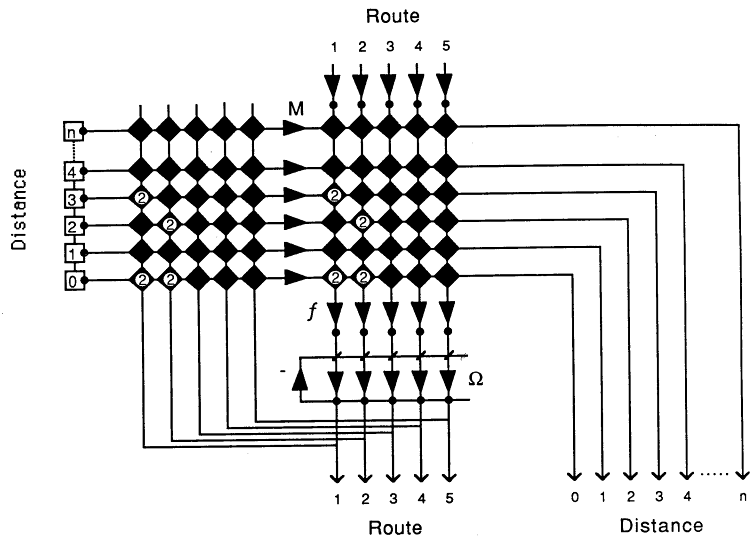


Distance of route 1 = 03

Distance of route 2 = 02

Figure 8.2
A synaptic matrix for binding travel distance to a selected travel route. *Lim* = 2.

tance have been associated during learning, when a filter cell ($f_i$) and class cell ($\Omega_i$) couplet is selectively discharged by the activity of a particular route token, its $\Omega$ collateral to the imaging matrix will discharge a mosaic cell ($M_i$) representing the distance of the route selected.

Recall the example of the two routes for conducting business within the office building. Suppose S were analyzing the distance of each for the first time. By the mechanisms previously described, activation of route 1 on the self-locus retinoid automatically evokes an analog of its overall distance that is represented in the one-dimensional scale of the autaptic cell input to the synaptic matrix shown in figure 8.2. Thus, the matrix receives two parallel token inputs: one representing the route considered, the other representing its overall distance. This results (in accordance with the learning equation) in the strengthening of two specific synaptic junctions on filter cell 1 (in the filter cell–class cell couplet that signals route 1) and the strengthening of one specific synaptic junction on each of two different mosaic cells. In the example, these are at the inputs to $f_1$ from mosaic cells representing distance 0 and distance 3. Synaptic enhancement also occurs at the inputs to mosaic cells distance 0 and distance 3 from an axon collateral of $\Omega_1$. When route 2 is analyzed for the first time, the same kind of process takes place. In this case, however, the synaptic junctions on $f_2$ that are strengthened are from mosaic cells distance 0 and distance 2. At the same time, an axon collateral of $\Omega_2$ strengthens its synapses on the mosaic cells distance 0 and distance 2.

After the synaptic matrix undergoes these changes, discharge of the class cell token of either route will result (via the modified synaptic junctions on the mosaic cells) in the virtually simultaneous discharge of its appropriate token of distance. Thus, recall of route 1 discharges distance token 3, whereas recall of route 2 discharges distance token 2. This provides a neuronal signal that at least one consequence of choosing route 2 is that the distance that must be traveled will be shorter. If minimization of the distance to be traveled (and, indirectly perhaps, time and effort) is a constraint on the selection of behavioral routines, then route 2 will be chosen to guide future behavior in this context. Notice also that discharge of a mosaic cell that signals a particular distance will recall the specific route(s) associated with that distance. On a more general perspective, if there were an opportunity to accomplish one of a variety of necessary errands within a limited period of time, this mechanism enables the selection of an errand that can be completed within the time available.

The example shows how a neuronal token of a particular consequence (distance traveled) can be bound to a token that represents a specific sequence of actions (travel route). The same kind of mechanism that is illustrated in figure 8.2 provides a general means for appropriately binding tokens of action and consequence over a wide variety of instances.

### *Encoding Plans*

A plan consisting of a sequence of self-directed instructions for gating the behaviors required to travel route 2 (starting at the subway exit) and to transact business in the three offices can be listed as follows:

1. Go (left)
2. Find (building entrance)
3. Go (through)
4. Go (left)
5. Find (door to office-1)
6. Go (through)
7. Do (transaction-1)
8. Find (door to office-1)

9. Go (through)
10. Go (left)
11. Find (door to office-3)
12. Go (through)
13. Do (transaction-3)
14. Find (door to office-3)
15. Go (through)
16. Go (ahead)
17. Find (door to office-4)
18. Go (through)
19. Do (transaction-4)
20. Find (door to office-4)
21. Go (through)
22. Go (right)
23. Go (left)
24. Go (left)
25. Find (building exit)
26. Go (through)
27. Go (left)
28. Find (subway entrance)
29. Branch (to subway routine)

On the assumption that substantial components of complex behavior are hierarchically organized (Rosenbaum 1977; Rosenbaum, Kenny, and Derr 1983) commands such as *Go, Find, Do,* and *Branch* are considered to be superordinate and are printed in this listing without parentheses, while specifiers such as *entrance, left, ahead, through, transaction-1* are considered to be subordinate and are enclosed in parentheses. This kind of hierarchical organization is embodied in the model mechanism illustrated in figure 8.3. It is a putative neuronal structure that can encode and initiate plans and control the particular sequence of actions called for by a plan.

In this model, when a plan of action is committed to memory (learned), the cell marked *Initiate Plan (n)* (*IP(n)*) is fired concurrently with the first step (only) of the plan. Discharge of the *IP(n)* line resets the step ring and fires autaptic cell 1 in the step ring. Thus, in figure 8.3 at step 1, when the mode cell *Go* and the specifier cell *left* were first fired during learning, the *Initiate Plan (n)* line was also discharged, and this evoked input from line 1 in the step ring. As each mode and specifier couplet fires, it evokes a stepping pulse that shifts the locus of activity to the next autaptic cell in the ordered sequence of the step ring. After all the behavioral steps that compose the essential elements of the plan have been learned, because of the pattern

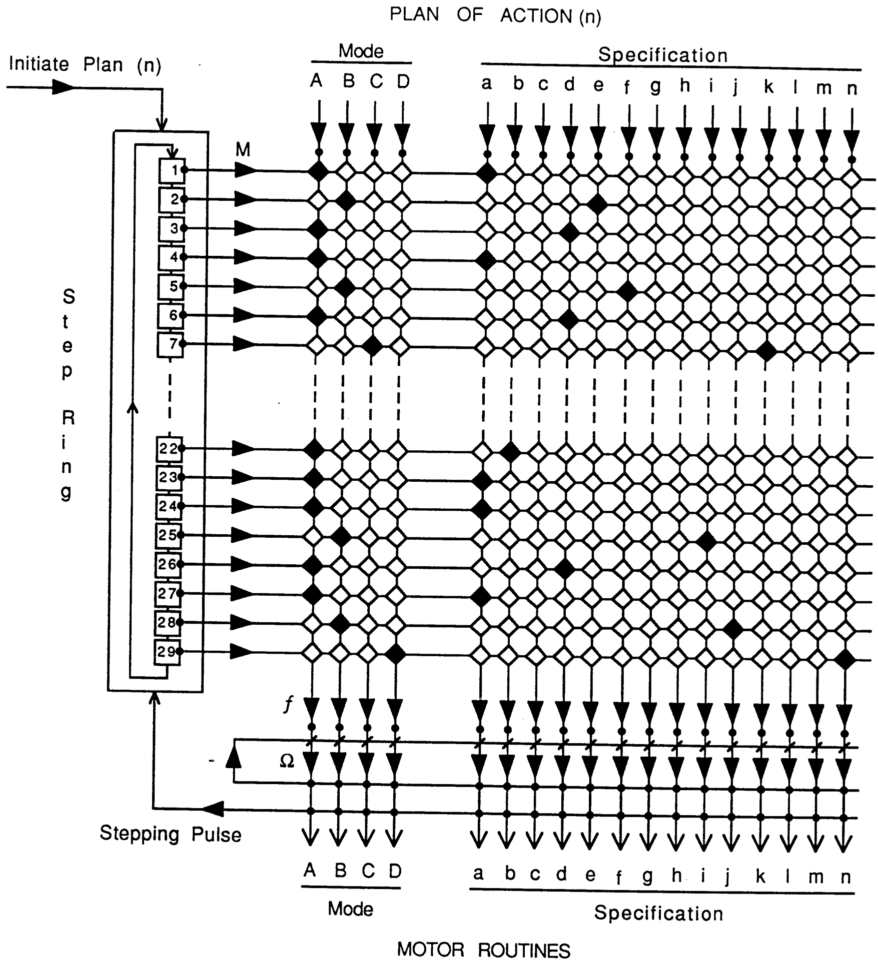


Figure 8.3
A mechanism for registering and recalling a plan of action. Filled lozenges represent synaptic junctions having augmented transfer weights ($\phi_i$).

of synaptic enhancement in the detection matrix, if *IP(n)* alone is fired, the entire action sequence represented by the stored plan will be automatically recalled and can be executed under the appropriate circumstances.

The synaptic changes represented in figure 8.3 are the result of successive inputs from the step ring and the sequence of concurrent discharge of mosaic cells (*M*) and filter cells (*f*) that correspond to the first seven and last eight instructions of the plan listed above for conducting business in a particular building. The elements of the overall plan are represented here as follows.

Mode: A = *Go*; B = *Find*; C = *Do*; D = *Branch*.
Specification: a = *left*; b = *right*; c = *ahead*; d = *through*; e = *building entrance*; f = *door to office-1*; g = *door to office-3*; h = *door to office-4*; i = *building exit*; j = *subway entrance*; k = *transaction-1*; l = *transaction-3*; m = *transaction-4*; n = *subway routine*.

During learning, where ⟨=, . . . , =⟩ denotes parallel excitation, the order of activation that established the synaptic structure shown in figure 8.3 is ⟨1, *IP(n)*, A, a⟩; ⟨2, B, e⟩; ⟨3, A, d⟩; ⟨4, A, a⟩; ⟨5, B, f⟩; ⟨6, A, d⟩; ⟨7,C, k⟩; . . . ⟨22, A, b⟩; ⟨23, A, a⟩; ⟨24, A, a⟩; ⟨25, B, i⟩; ⟨26, A, d⟩; ⟨27, A, a⟩; ⟨28, B, j⟩; ⟨29, D, n⟩. The behavioral subroutines represented by steps 1–7 provide for the appropriate actions from leaving the subway to completion of transaction-1 in office-1. The subroutines represented by steps 22–29 provide for leaving the building after final business is accomplished (transaction-4) and finally for branching to the subway routine for the ride home.

The mechanism shown in figure 8.3 does not directly control motor behavior. What it does is organize and store an addressable neuronal representation of action sequences (a plan of action) that can prime a series of motor subroutines competent to accomplish a selected task. Whether the actions represented within any such register dedicated to a given plan are actually expressed in overt behavior depends on the veridical ecological context and neuronal gating (by increase in general excitatory bias) of the specific motor subroutines primed by the plan. In a strictly cognitive review (covert activation) of a plan of action, there are no necessary constraints on the time interval between successive steps. This is clearly not the case when a plan is overtly expressed. The action called for by each subroutine must be completed in real time before the next step is initiated. This suggests that in the overt execution of a plan (as distinguished from its cognitive construction or review), each stepping pulse to the step ring will be initiated by the terminus of each subroutine in the overall plan rather than by discharge of class cell tokens for mode and specifier.

*Composing a Plan*

Given a neuronal repertoire of $N$ different subroutines of action, what controls the composition and initiation of a particular plan? A very general answer might be that a plan is composed and activated by an individual in the belief that it will achieve a particular salient goal in the context of a perceived environmental situation (either current or anticipated). The concept of belief carries the flavor of a resident homunculus, and in the interest of biological explication I will try to characterize it later in terms of a specific kind of neuronal state. For the moment, however, consider some conditions that might account for such a belief. First, real-life experience might have shown that the achievement of a goal was a reliable consequence of executing a particular plan. In this case, a neuronal token that signals the achievement of a goal (say $G_1$) is bound to the token of a plan (say $P_1$) by synaptic modification (learning) in a synaptic matrix of the kind illustrated in figure 8.2. Second, covert analysis of a new situation might reveal that out of a set of possible action sequences that can, in principle, lead to the achievement of a goal, only certain individual actions can be performed by the actor but that this subset, if properly chained, can attain the goal. Finally, one might accept a plan of action offered by an outside source in the belief that it will be successful because of relevant expertise attributed to the source.

Consider the microworld schematized by the directed graph at the top of figure 8.4. Each geometric shape represents a particular local situation in which the possible actions of an individual are constrained by local affordances so that the actor can accomplish only those transitions to other situations that are joined to the local one by directed paths. There are three goals ($G_1$, $G_2$, $G_3$). Each can be reached only by means of a particular prior situation-action. Goal 1 is reached by action enabled in the situation represented by the horizontally oriented rectangle; goal 2 by action in the situation represented by the vertically oriented rectangle; and goal 3 from the situation represented by the hourglass figure.

Because of the nature of environmental and behavioral constraints in the natural world, the initial situation at the time an actor (S) first aims to reach a specific goal will determine the particular sequence of actions that can attain the goal. At the bottom-left side of figure 8.4 are two columns—one a list of geometric shapes designating S's starting situation and the second an associated list of goals ($G_i$), each to be reached from its paired starting situation. For each goal there is a plan of action—a series of action-mediated transitions from one situation to another in a path that should culminate in the attainment

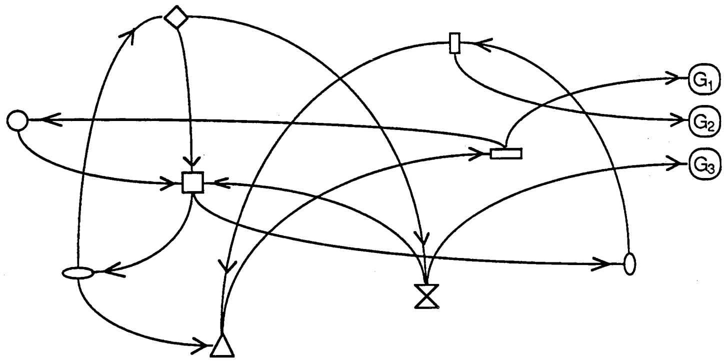


| | Start | Goal | | Plan of Action |
|---|---|---|---|---|
| 1 | △ | $G_1$ | 1 | [ △ → ▭ → $G_1$ ] |
| 2 | ○ | $G_1$ | 2 | [ ○ → □ → ⬭ → △ → ▭ → $G_1$ ] |
| 3 | △ | $G_2$ | 3 | [ △ → ▭ → ○ → □ → ⬯ → ▯ → $G_2$ ] |
| 4 | ⧖ | $G_2$ | 4 | [ ⧖ → □ → ⬯ → ▯ → $G_2$ ] |
| 5 | ◇ | $G_1$ | 5 | [ ◇ → □ → ⬭ → △ → ▭ → $G_1$ ] |
| 6 | ◇ | $G_2$ | 6 | [ ◇ → □ → ⬯ → ▯ → $G_2$ ] |
| 7 | ◇ | $G_3$ | 7 | [ ◇ → ⧖ → $G_3$ ] |
| 8 | □ | $G_1$ | 8 | [ □ → ⬭ → △ → ▭ → $G_1$ ] |
| 9 | □ | $G_2$ | 9 | [ □ → ⬯ → ▯ → $G_2$ ] |
| 10 | □ | $G_3$ | 10 | [ □ → ⬯ → ▯ → △ → ▭ → ○ → ⊡ ] |
| | | | 11 | [ □ → ⬭ → △ → ▭ → ○ → ⊡ ] |
| | | | 12 | [ □ → ⬭ → ◇ → ⧖ → $G_3$ ] |

Figure 8.4
*Top:* Schematic representation of transition affordances among a number of situation actions and finally to one of three separate goals. *Bottom:* Ten states, each characterized by a starting situation action (indicated by a distinctive geometric shape) and a goal to be achieved ($G_i$). States 1–9 are each followed by a plan of action that leads directly to the goal. State 10 is followed by two ineffective plans (plans 10 and 11, leading back to the starting state), and plan 12, which leads directly to the desired $G_3$.

of the goal. Plans 1–9 will effectively reach the goals that have been set. Plans 10 and 11 will not; in these two cases, S ends up back in the same situation that he started from, without reaching the goal. Plan 12, however, although starting at the same point as plans 10 and 11, is effective in achieving the goal.

Let us assume that initially S has knowledge of only the local transition affordances of each situation (states from which a target condition can be directly achieved and states that can be directly achieved from a given target condition), including those from which a one-step transition reaches a goal. Given a starting situation from which more than two transitions are required, and in absence of relevant prior global knowledge, there is no way that S can ensure on the basis of local affordances that any particular local transition will be part of a transition sequence (plan of action) that attains a desired goal. Thus, S must engage in either overt or covert trials, testing different chains of situation-action transitions to determine which, if any, provide a continuous path between a current situation (or one known to be attainable) and the desired goal. In a situation-action domain like that depicted in figure 8.4, there is a basic asymmetry in the information available with respect to the attainment of a goal. That is, if one focuses on the local affordances of the starting situation, there is no guarantee that the next step will lie on a path that reaches a desired goal, but if one focuses on the local affordances of the goal, an effective prior state (situation-action on the path to the goal) is revealed with certainty. This suggests that the appropriate strategy for covert testing of possible action sequences in developing a plan is to recall the local affordances at the goal and then, in successive steps, test backward to an acceptable starting condition.

To see how such an analysis might be performed in the neuronal mechanisms of the brain, consider plan 12 shown at the bottom of figure 8.4. The fact that the action represented by the hourglass figure (say $A_3$) has the immediate consequence of achieving $G_3$ can be learned and represented by neuronal tokens in a synaptic matrix like the one that serves to organize a semantic network. The same kind of mechanism can represent the local transition consequences of any action depicted in the situation-action domain of figure 8.4. Moreover, because discharge of the class cell token of a consequence will evoke a token of the action that produces the consequence and because each action can be considered as having been enabled by an immediately preceding action, it is possible to evoke a chain of neuronal tokens that represents a backward sequence of transition affordances from any goal state to earlier states on the path to the goal.

All of the situation-action affordances illustrated at the top of figure 8.4 can be captured and represented in a distribution of synaptic transfer weights in a matrix similar to the one shown in figure 6.2. The following sentences presented to the simulated synaptic matrix provided all of the information required for planning in the environment of figure 8.4. The lexical labels for the environmental tokens are self-explanatory except for *hoval* (horizontally oriented oval), *voval* (vertically oriented oval), *hrect* (horizontally oriented rectangle), and *vrect* (vertically oriented rectangle). The word *gets* indicates a transition affordance from a given situation-action to another or else to a goal:

1. Circle gets-square .
2. Hoval gets-diamond .
3. Hoval gets-triangle .
4. Diamond gets-square .
5. Diamond gets-hourglass .
6. Square gets-hoval .
7. Square gets-voval .
8. Triangle gets-hrect .
9. Hourglass gets-square .
10. Hourglass gets-goal3 .
11. Vrect gets-triangle .
12. Vrect gets-goal2 .
13. Hrect gets-goal1 .
14. Hrect gets-circle .
15. Voval gets-vrect .

After this information was learned, the model could formulate a plan of action for reaching a desired goal from any arbitrary starting situation. Consider the problem illustrated at the bottom of figure 8.4: S must compose an effective plan that will achieve goal $G_3$ from the starting situation designated by the square. S uses the following process of self-query within the matrix module that is automatically initiated when he is first motivated to reach $G_3$. At the biological level, a self-query corresponds to an endogenously evoked discharge of the class cell token of a predicate *gets-⟨situation-action⟩*. The answer corresponds to the mosaic cell token that is then discharged by the axon collateral of the activated class cell:

*Question:* What gets-goal3?
*Answer:* Hourglass gets-goal3.

*Question:* What gets-hourglass?
*Answer:* Diamond gets-hourglass.

*Question:* What gets-diamond?
*Answer:* Hoval gets-diamond.

*Question:* What gets-hoval?
*Answer:* Square gets-hoval.

Discharge of the neuronal token for *square* brings S to the internal state that corresponds to the starting situation, and self-query is terminated. At this point, a series of neuronal tokens has been discharged, beginning with the goal to be reached ($G_3$) and ending with the situation at which the plan is to start (*square*). It is assumed that this backward chain of situation-action tokens is captured in short-term memory by a string of autaptic cells and then registered in the proper sequence (square → $G_3$) as a plan of action in the mechanism illustrated in figure 8.3. The plan—square → hoval → diamond → hourglass → $G_3$—corresponds to plan 12 shown at the bottom of figure 8.4.

Similar simulation tests of the model were run for a variety of planning problems, all designed to determine the sequences of action that will ensure that each selected goal will be reached from arbitrary starting situations. In each case, a satisfactory plan was produced. For example, starting at the situation represented by *triangle*, with $G_2$ as the goal, the process of self-query was initiated as follows and proceeded until the token for *triangle* was discharged:

*Question:* What gets-goal2?
*Answer:* Vrect gets-goal2.

*Question:* What gets-vrect?
*Answer:* Voval gets-vrect.

*Question:* What gets-voval?
*Answer:* Square gets-voval.

*Question:* What gets-square?
*Answer:* Circle gets-square.
Diamond gets-square.
Hourglass gets-square.

*Question:* What gets-circle?
*Answer:* Hrect gets-circle.

*Question:* What gets-hrect?
*Answer:* Triangle gets-hrect.

On completion of self-query, the plan of action was: triangle → hrect → circle → square → voval → vrect → $G_2$. A check of this action sequence against plan 3 in figure 8.4 shows that it is correct.

Notice, however, that in response to the query *What gets-square?* there was more than one answer. In addition to *circle, diamond* and *hourglass* provided transition affordances to the situation represented by *square*. As it turned out, querying *circle* led to a direct path back to the appropriate start at *triangle*. But, if instead of *circle*, either *diamond* or *hourglass* were queried, then subsequent transitions would have led to the previously established link in the sequence—*square*. For example,

*Question:* What gets-square?
*Answer:* Circle gets-square.
Diamond gets-square.
Hourglass gets-square.

*Question:* What gets-diamond?
*Answer:* Hoval gets-diamond.

*Question:* What gets-hoval?
*Answer:* Square gets-hoval.

If the same query were repetitively evoked, the system would be locked in an ineffective recurrent loop. It is assumed that in processes of this kind, when a previously established link is repetitively fired before a token of the selected starting situation is evoked, a different candidate for transition to the repeated state is discharged for the next query (*circle*). A neurophysiological basis for such a shift in query might be the fatigue (and associated rise in relative threshold) of repetitively discharged cells.

### *Storing and Recalling Plans*

Effective plans, once composed, are stored as long-term resources in the brain's neuronal networks to serve the efficient pursuit of goals and the solution of diverse ecological problems. A plan is selectively initiated when a current situation is recognized as similar to an earlier situation in which the particular plan (in fact or in imagination) was successful.

Figure 8.5 illustrates a model neuronal register in which plans (discrete action sequences) are selectively associated with the tokens of those situations where they were successfully applied. The states that are mapped to the input tokens each consist of a particular starting situation and a goal to be achieved from that situation. Thus, in figure 8.5 if situations 1, 2, 3, . . . , *n* are identified with the corresponding pairs of *Start* and *Goal* in the lower-left columns of figure 8.4, the

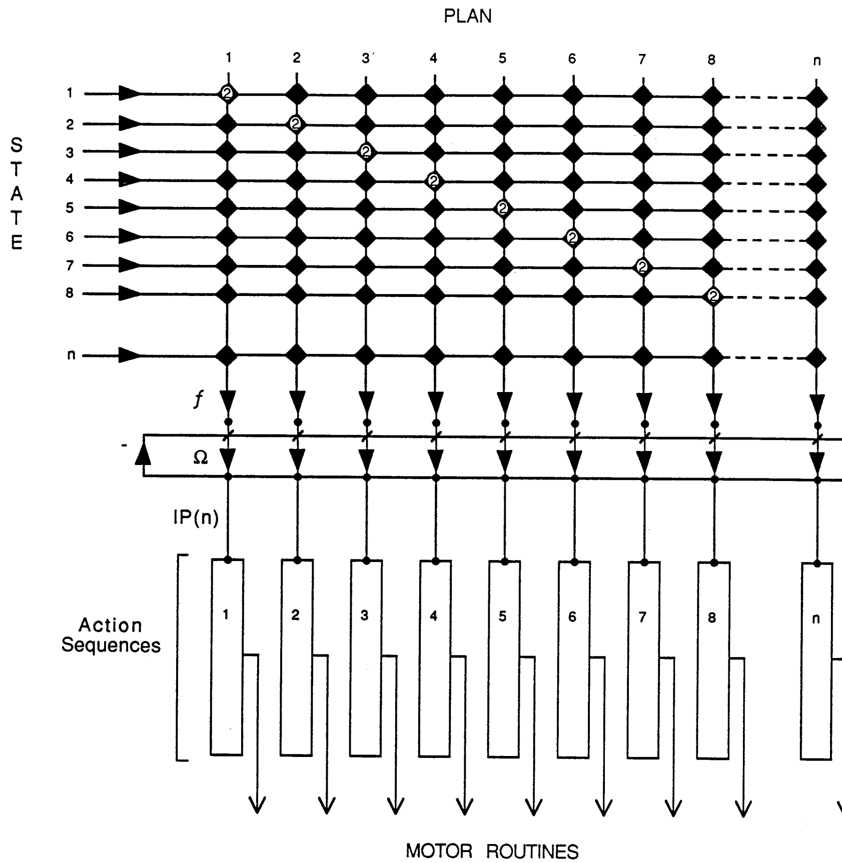


Figure 8.5
A neuronal register for selectively binding plans of action to the states in which they represent effective solutions. Evocation of a particular state recalls the appropriate plan. *Lim* = 2.

action sequences that are selectively initiated by this matrix correspond to the effective plans of action shown in figure 8.4. In figure 8.5, each indexed box triggered by an output *IP(n)* is assumed to contain a stepping ring and a matrix to control the modes and specifications for the proper sequence of actions constituting a plan, as detailed in figure 8.3. Taking state 4 (figure 8.4) as an example, all of the behavioral commands required to move from *hourglass* to *square* to *voval* to *vrect* to $G_2$ would be represented by the distribution of synaptic weights in box 4 of figure 8.5.