

Chapter 3

Learning, Imagery, Tokens, and Types: The Synaptic Matrix

An underlying premise of this book is that the cognitive brain consists of many special-purpose mechanisms synergistically organized in integrated networks. The first multineuronal mechanism that we will consider is a key processing module called a synaptic matrix (Trehub 1967). It is a putative brain mechanism with the capacity to learn and classify complex input patterns, store them in long-term memory, and recall images of them in the absence of external stimulation or when only fragments of them are presented as input. It is a self-organizing neuronal implementation of a parallel, adaptive comb filter that maps large input vectors onto unitary connectively labeled output lines (axons) and selectively maps these output lines onto correlated state vectors that are homologous with learned inputs (Trehub 1975a, 1977, 1987).

In the schematic drawing of a synaptic matrix shown in figure 3.1, there are nine afferent input lines (S_{ij}) and five output lines (Ω_i). In the human brain, a single synaptic matrix may have tens of thousands of axonal inputs and hundreds of thousands of outputs. Two subscripts, i and j , are used in figure 3.1 to index the input vector S because, in this case, it is assumed that the input neurons map from a two-dimensional transducer array. Inputs S_{ij} are in discrete point-to-point synapse with a second set of afferent neurons, called mosaic cells (M). The axon of each mosaic cell is in parallel adaptive synapse with all members of a set of cells in the detection matrix, which are called filter cells (f). Each filter cell is in discrete synapse with an output neuron called a class cell (Ω). The axon of each class cell bifurcates, sending a collateral back in adaptive synapse with the dendrites of all mosaic cells (M) in the imaging matrix. Finally, an inhibitory neuron (marked $-$) receives as its input the axons of all class cells and, in turn, sends its axon in parallel synapse to the dendrites (or cell bodies) of all class cells. This inhibitory cell is called a reset neuron.

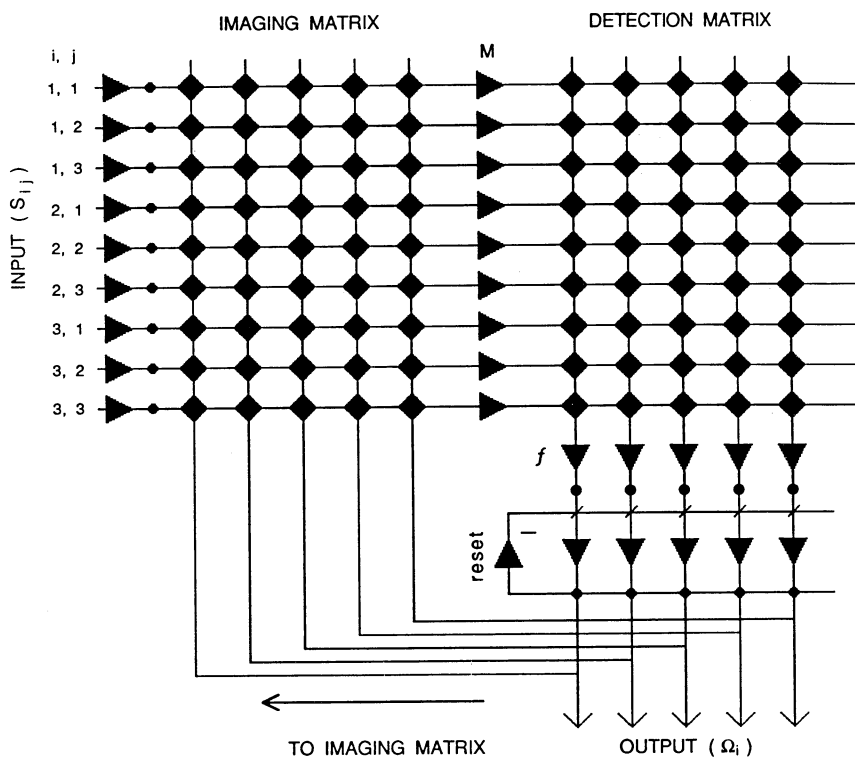


Figure 3.1

Schematic of a synaptic matrix. Afferent inputs are designated S_{ij} , and mosaic cells are designated M . Dots represent fixed excitatory synapses; short oblique slashes represent fixed inhibitory synapses; filled lozenges represent adaptive excitatory synapses. Reset neuron ($-$) generates an inhibitory postsynaptic potential to reset all class cells when discharged. Given any arbitrary input, the class cell coupled with the filter cell having the highest product sum of afferent axon activity and corresponding transfer weights will fire first and inhibit all competing class cells.

Learning and Pattern Recognition

For the purposes of illustration, let us assume a minimal adaptive system in the visual modality. There is a retina with nine discrete receptor cells arranged in a two-dimensional 3×3 array. Each receptor cell connects to an afferent neuron, which is indexed by the position of its associated receptor in the retinal array. We can represent this biological index by using retinotopic two-dimensional coordinate subscripts as index labels. Thus, each afferent neuron in figure 3.1 is designated S_{ij} where ij represents the retinotopic coordinate of its associated retinal cell. Let us also assume that center-surround inhibition (Kuffler 1953) and threshold properties at the retinal layer and lower visual centers result in the extraction of a binary-valued contour transform of the light intensity distribution caused by an image falling on the retina. Each afferent neuron discharges (activity = 1) if an edge is detected at its associated retinal locus; it remains silent (activity = 0) if an edge is not detected at its retinal locus. This array of edge detection signals composes the afferent excitation pattern (input vector) on the mosaic cell array (M). Taking the contour transform of the light-intensity distribution on the retina as input to the synaptic matrix is consistent with the recent experimental finding of Biederman and Ju (1988) showing that initial access to a mental representation of an object can be modeled as a matching of an edge-based representation.

At some early stage of maturation, the synaptic matrix becomes an effective adaptive brain mechanism. ATF and DTF are produced in its mosaic cells and filter cells, and ATF alone is produced in its class cells. With the production of the substances essential for synaptic plasticity, transfer weights (ϕ_i) can be modified by afferent excitation and sustained at their new levels. In this process of learning, the value of ϕ at each synapse will be set in accordance with equation 2.3,

$$\phi_{im} = b + S_{im}(c + kN^{-1}),$$

$b \rightarrow Lim$

where the parameters b , c , and k are subject to the loose constraint $b < c \ll k$.

Before any learning has occurred, all ϕ will be small; taking visual learning as the example, an image falling on the retina would normally produce relatively feeble, infrequent, and random firing of filter cells. If, however, cells in the synaptic matrix are positively biased by arousal (excitation from the reticular activating system), then added visual stimulation can induce a level of activity in some filter cells that exceeds the threshold Θ for ATF-DTF interaction. In this case, any filter cell (f_1) receiving sufficient excitation from a retinal

input pattern (pattern A) via the mosaic cell array will undergo an effectively stable increase in ϕ at those of its synapses that intersect the afferent firing pattern (A) in the detection matrix. After this has happened, since the products of afferent stimulus units (S_{ij}) and their corresponding synaptic transfer weights (ϕ_i) are integrated in each filter cell (equation 2.1), whenever pattern A is presented, f_1 will fire at a higher frequency than other filter cells in the detection matrix because its rate of increase of EPSP, given A, will be higher than that for any other filter cell that has not been conditioned by the pattern A (expression 2.2). If uniform rates of integration and thresholds of discharge in the following class cells are assumed, the class cell connected to f_1 (Ω_1) will fire before any other Ω_i in the matrix output array because its spike input frequency will be higher than any other. In effect, this class cell (Ω_1) represents the biological name for the learned pattern A. Thus, relative spike frequency on labeled lines (class cells) is the effective code in this neuronal system.

The first class cell to fire (in this case, Ω_1) will discharge the inhibitory reset cell ($-$), which will reset the integration level of each neuron in the entire class cell array before EPSP in the other class cells can reach the firing threshold. The same process of Ω_1 discharge and cell reset will recycle continuously as long as the same stimulus (in this case, A) is present (or until the more active cells fatigue). In this fashion, any given visual pattern is identified at the neuronal level by the particular class cell it discharges. If there were no reset mechanism, all class cells in the detection matrix would continuously integrate EPSP and would reach their individual firing thresholds at random moments independently of the input pattern present at the time.

Imagery

Notice in figure 3.1 the bundle of axon collaterals from the set of class cells (Ω_i) that courses back to form adaptive synapses with the dendrites of all mosaic cells in the imaging matrix. When the visual stimulus A is presented, mosaic cells responding to the afferent input of A will discharge, thus firing the appropriate filter cell (f_1) at the highest frequency in the detection matrix and causing its paired class cell (Ω_1) to fire. The transfer weights (ϕ_i) of the adaptive synaptic junctions between the active Ω_1 collateral and mosaic cells (M) that are not discharging in the presence of stimulus A will remain unchanged, but those adaptive synaptic junctions on the mosaic cells that are firing (as afferents responding to the input pattern A) will be modified in accordance with equation 2.3. Synaptic transfer weights

between the class cell Ω_1 , and the mosaic cells will therefore be selectively increased for the M pattern A in the imaging matrix. After this synaptic change has occurred, if Ω_1 is fired (addressed), by either its associated filter cell (f_1) or any other input, it will evoke the afferent firing pattern on the mosaic cell array that is normally evoked by A, even when the stimulus A is not present. In this way, by means of automatic, selective synaptic weighting in the imaging matrix, the neuronal condition is established for retrieving, in the absence of retinal stimulation, an entire afferent pattern previously coded by a particular class cell. The capacity to recreate a specific discharge pattern on the mosaic cell array in the absence of a corresponding retinal stimulus provides one of the neuronal mechanisms for human imagination (Kosslyn 1980, Sheikh 1983, Shepard 1978, Trehub 1977).

Since only a single class cell collateral is normally active at any given time during the discharge of mosaic cells in the array, the reduction of free DTF in each active mosaic cell is minimal when an image is learned. Thus, a single mosaic cell can participate in the adaptive construction of many different visual representations in the imaging matrix.

An Example of Performance

Following is a simple example of how the synaptic matrix with input from a 3×3 cell retina performs the task of learning, recognizing, and imaging four different stimuli. In this simulation, the saturation limit (Lim) was set at a value of 6 for filter cells in the detection matrix and at a value of 3 for mosaic cells in the imaging matrix. It was assumed that all active afferent lines (S_{ij}) carried uniform excitation. Inputs were designated 1 for an active input line or 0 for an inactive line. Parameters in the learning equation (equation 2.3) were arbitrarily set as follows: $b = 1$, $c = 2$, $k = 10$. For simplicity of illustration, stimuli are designated by numbers corresponding to the order in which they were presented and learned, and each filter cell-class cell couplet is indexed by the pattern that modified it. Figure 3.2 shows the distribution of synaptic transfer weights on each adaptive neuron in the detection matrix and the imaging matrix (rounded to the nearest integer) after each stimulus was presented one time. Filter cell 5 has not been modified and is available for learning another pattern because its store of free DTF has not been depleted and remains at its maximum level.

After the four stimuli (vertical bar, square, horizontal bar, and cross) were learned, the synaptic matrix was tested for recognition of each of the stimuli. Figure 3.3 shows the relative frequency of

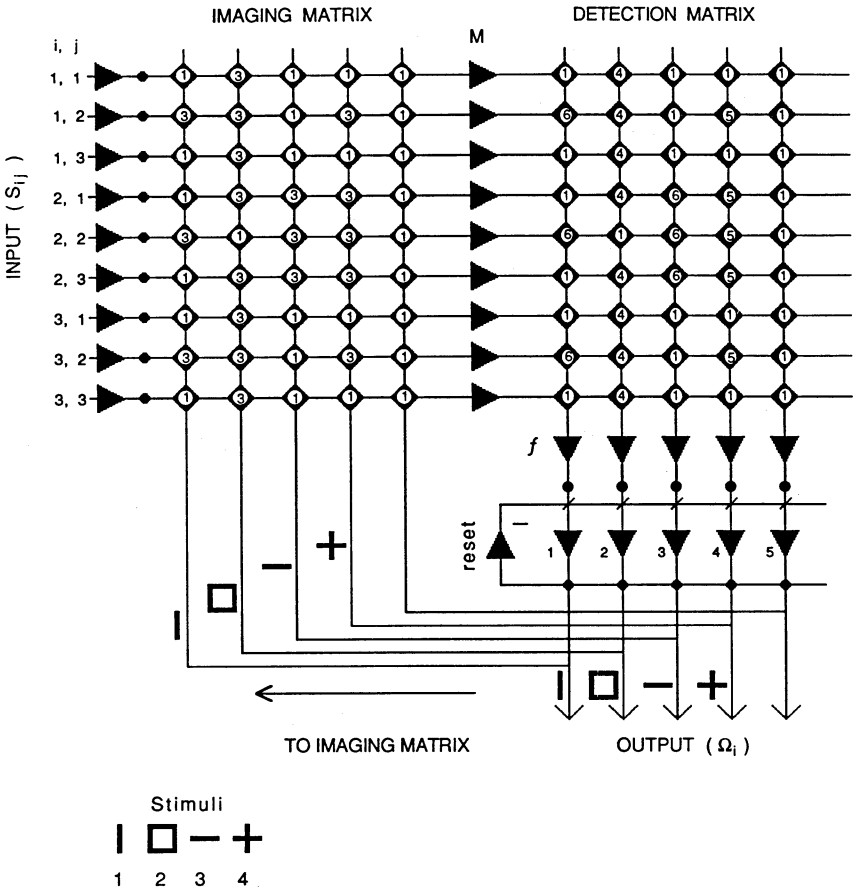


Figure 3.2
Synaptic matrix showing synaptic transfer weights after four different stimuli have been learned. The distribution of transfer weights in the detection matrix reflects a peak tuning of each filter cell to one of the input patterns. If any class cell is discharged, the distribution of transfer weights in the imaging matrix can evoke a pattern of excitation on the mosaic cell array that corresponds to the stimulus originally mapped to that class cell.




Figure 3.3

Bar graph of relative discharge frequency for each filter cell in the synaptic matrix shown in figure 3.2 when each of the four stimuli is presented after all have been learned. The cell showing maximum output in response to the stimulus presented has an asterisk above it. Shapes next to bars in graph identify both the stimulus presented on the trial and the filter cell that originally learned the stimulus.

discharge in all filter cells (constituting input to each of their coupled class cells) for each of the stimuli. Activity is maximum for the appropriate filter cell in all presentations, resulting in the discharge of the correct class cells (figure 3.2).

Notice that stimulus 4 (the cross) completely includes two other stimuli that had been learned (vertical bar and horizontal bar), yet the system discriminated the cross and gave the proper recognition response. And the vertical and the horizontal bars, both fully included in the cross stimulus, were also correctly recognized. The problems of normalization discussed in chapter 2 are remedied by the learning mechanism implicit in equation 2.3. The quasi-normalized DTF contribution and the small, fixed ATF contribution have a joint effect on synaptic transfer weights that gives the detection matrix the ability to make appropriate discriminations even when a learned stimulus is a substantial integral part of another learned stimulus.

Turning to figure 3.2 again, examination of the pattern of modified transfer weights in the imaging matrix reveals that discharge of any arbitrarily chosen class cell will evoke a distribution of excitation over the array of mosaic cells that corresponds to the afferent pattern signalled by that class cell. Since each mosaic-cell dendrite receives an active input from only one excited class-cell collateral at any given time, depletion of DTF will be minimal. Thus, each mosaic cell can adaptively change its synaptic weights to accommodate many subsequent image representations (memories).

Operating Characteristics of the Synaptic Matrix

In addition to its primary processing capabilities, other properties of the synaptic matrix deserve particular attention. An early simulation test of the model (Trehub 1975a) presented nine different visual patterns through a 5×5 -cell retina. To satisfy the constraint $b < c \ll k$, the following values were arbitrarily assigned: $b = 1$, $c = 5$, $k = 100$. After all the patterns had been presented once during the learning phase, each was used in repeated stimulus-response trials to test recognition. Each stimulus was either an intact figure, as originally learned, or a degraded version of the original figure formed by randomly eliminating approximately 40 percent of the pixels forming the pattern (figure 3.4).

The upper half of figure 3.4 shows that when the learned patterns were presented intact in the test phase, response was 100 percent correct. The robustness of performance of the synaptic matrix is illustrated in the bottom half of figure 3.4, which shows what happens




Figure 3.4

Results of simulation test. Numbers at top identify the nine stimuli that were learned. Notice the error (indicated by asterisk at bottom) made in response to the randomly degraded cross.

when the matrix must recognize pattern fragments having approximately 40 percent of the originally learned pattern randomly omitted. In this highly degraded condition, only one error of classification is committed—mistaking the fragment of a cross for a vertical bar. Inspection of those fragments of the cross presented as a stimulus reveals how intuitively appropriate the erroneous response was. Repeated tests with randomly generated fragments of the original stimuli (approximately 60 percent of each learned prototype) yielded correct responses in 92 percent of the tests. This early simulation demonstrated that the synaptic matrix performs well in the presence of subtractive noise. (Other simulations to be presented demonstrate robust recognition performance in the presence of both additive and subtractive noise, as well as when stimuli are confounded by the superposition of learned patterns.)

The synaptic matrix operates on the basis of ordinal logic rather than point logic. When a stimulus is presented, the sum of the products for mosaic cells and their corresponding synaptic transfer weights on the filter cells (f_m) does not have to be any particular value (a point criterion) for the mechanism to respond appropriately. All that is required is that the sum of these products for a correct filter cell be greater than the sum for each of the other competing filter cells (an ordinal criterion).

The model has direct implications for our understanding of what a concept might be in terms of brain function. The detection matrix does not compute a single prototypical distribution of synaptic transfer weights for a given class of objects. Thus, there is no representa-

tion of an "average" pattern that might be taken as a person's concept of an object. Rather, in the course of human learning, filter cells may be tuned to different exemplars of a single object class. Since the decision about class (category) membership is made on the basis of competitive, first-order pickoff among the possible class cell outputs, the system exhibits robust intrinsic generalization and can, in principle, assign proper category membership to many objects that it has never learned as exemplars of their class. This is true because the mechanism operates on the principle of ordinal logic, which ensures that even in the absence of a good match between an input pattern and its appropriate filter cell, if there is no better match with an inappropriate filter cell, the correct recognition response will be given. The effect of this operating principle is to make the system appear as if it is responding on the basis of the fit of each stimulus against a set of averaged internal standards (Estes 1986) when, in fact, there are no averaged standards. It will be seen later that when the bare synaptic matrix is augmented by other putative brain mechanisms, the occurrence of a poor match against all filter cells in the detection matrix results in a signal that the current stimulus is novel, and if the ecological and motivational context warrants it, the level of arousal will increase and the stimulus will be automatically learned (Trehub 1977).

With respect to the distinction between concept and category, one might say (from the biophysical standpoint) that a concept consists of that subset of exemplar-tuned filter cells together with their coupled class cells that has at least one common output effect distinct from the output effects of any other subset in the detection matrix. For example, the presentation of the stimulus object A will maximally stimulate one of the cells in the subset of filter cells that defines the concept of A-ness. The discharge of its coupled class cell will cause the output signal (A), which then defines the category of the stimulus instance. A proper category signal can be any arbitrary but distinct and stable correlate of the class cells that evoke the signal.

The synaptic matrix does not employ feature extraction or feature list checkoff in the learning and subsequent recognition process. For instance, in the examples presented, learning the cross did not require that its constituent features of a vertical and horizontal bar be isolated and learned as separate representations. Nor did the recognition of the cross require that there be discrete representations of the bar features in the detection matrix. The excitation pattern of a stimulus on the mosaic cell array was learned holistically and represented holistically in the distribution of synaptic transfer weights on the filter cells.

The competence of this theoretical model suggests that natural pattern recognition can proceed very well without prior feature extraction, as an example at the level of insect behavior shows. In a carefully controlled study that explored how bees remember the shape of flowers, it was found that bees apparently store flower patterns as "low-resolution eidetic images" (Gould 1985). But this does not imply that organisms are necessarily incapable of feature extraction. The issue of what constitutes a feature and how it might be extracted from a complex visual stimulus is still quite murky (Pinker 1984). For example, study of the visual system in lower vertebrates suggests that they respond best to simple stimuli, such as small, horizontally elongated blobs in motion (Ingle 1968). In such cases, however, it seems more appropriate to ascribe the feature selectivity to an absolute constraint on the visuo-motor system imposed by primitive, relatively simple, pretuned filtering mechanisms that are suited only to the basic survival needs (prey catching, avoidance of predators). It is assumed that in the higher mammals or at least in humans, mosaic cell excitation can be limited to salient parts of a complex object, and these pattern parts can then be learned separately as constituent features of the larger object. In this way, filter cells can be constructed that detect and classify not only whole objects and scenes but their salient components as well.

In perceptual situations where immediate discrimination of a particular object on the basis of holistic filters is difficult, analytic detection of the previously learned parts of that object can assist the recognition process—for example, recognizing a semirigid object, such as the human body. In this case, isolation of a relatively rigid part, such as the head, enables the viewer to infer the whole figure from a recognized component.

To the extent that features may be regarded as important aspects of shapes that deserve special weighting, this approach provides intrinsic weighting of especially significant object parts on the following bases:

1. Parts that are invariant aspects of objects over a variety of perturbations will, on average, be represented in the transfer weight distributions of more filter cells than will noninvariant parts.
2. Parts of objects that have particular ecological relevance or utility (the handle of a cup) will naturally tend to be the focus of perceptual orientation, and more exemplars of such parts (features) will tend to be represented in the detection matrix.
3. If experience teaches that an ambiguous stimulus can be re-

solved by determining the presence or absence of a particular part, the part that solves the problem will tend to be isolated and learned as a useful feature.

Thus, in this model, recognition by feature detection is a special case within the general pattern recognition process.

Another significant property of the synaptic matrix is that its reaction time depends on the relationship between the current stimulus and the learning experience of the matrix. Because class cells integrate filter cell output over time and fire when their integrated EPSP reaches threshold, the reaction time for output from the detection matrix is a monotonic function of the correlation between the current input pattern and the best-fitting ϕ representation among the filter cells that have been modified by previous learning. Thus, the more familiar a stimulus is, the faster will the system respond.

Propositional and Analogical Representation

If we examine the physical state of a synaptic matrix after a stimulus pattern has been learned, we can characterize the modified distribution of synaptic weights as a latent representation, a deep structure that determines the subsequent surface behavior of the system. In the detection matrix, this synaptic structure results in the selection of output lines (class cells) that are selectively contingent on particular stimulus properties. In effect, the discharge of a class cell that is evoked by a pattern of activity (S) in the mosaic array is the physical equivalent of a propositional representation in the simple form S (noun) *is* p (predicate) (see Jackendoff 1987). It is comparable, for example, to the assertion that a stimulus object A belongs to the specific category Ω_1 or that B belongs to the category Ω_{394} . Discharge of any particular class cell can also be considered as a physical symbol (Derthick and Plaut 1986, Newell 1980) or token of the object (more accurately, of the object's proximal neuronal analog in the imaging matrix) to which its coupled filter cell has been synaptically tuned.

On the other hand, when the imaging matrix is excited by the input of a class cell collateral, the deep distribution of synaptic weights causes a surface response on the mosaic array that is an analog of some particular stimulus object. In effect, activation of a specific physical predicate evokes an analog representation of its proper subject—somewhat like depositing a monetary token in a slot and getting back its face value in currency. The ability of the synaptic matrix to operate interactively and with facility in the domains of both analogical and propositional representation provides an efficient biological mechanism for higher-order cognitive processing.

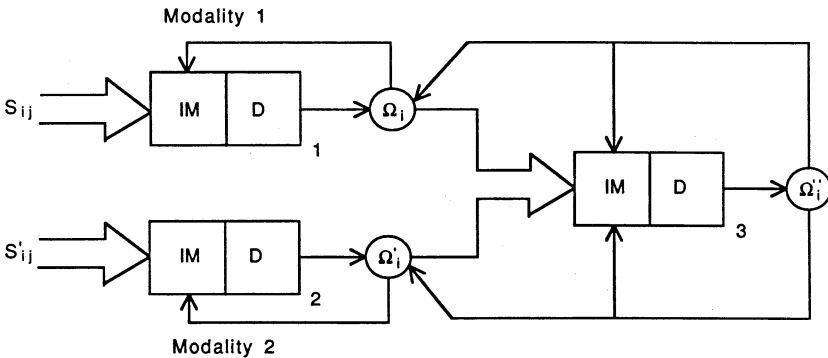


Figure 3.5

Flow diagram of three synaptic matrices in a feedforward-feedback network. Parallel inputs from two different sensory modalities. IM = imaging matrix; D = detection matrix.

Networks of Matrices

Up to this point, we have considered the basic capabilities of the synaptic matrix as a single neurocognitive module. Other important properties become evident when multiple synaptic matrices are organized in complex networks with parallel, feedforward, and feedback connections.

Figure 3.5 is a flow diagram illustrating three interconnected synaptic matrices. In this example, two separate modules (1 and 2) receive parallel volleys of stimulation from their respective sensory input modalities S_{ij} and S'_{ij} (vision and audition). The class cell outputs (Ω_i and Ω'_i) of these matrices, in turn, constitute an input vector to a third synaptic matrix that lacks direct access to the sensory arrays. After learning has occurred, each token Ω_i and Ω'_i can directly evoke an afferent image of its associated sensory pattern, but Ω''_i can directly only evoke a representation of its associated token inputs in its mosaic cell array. However, if associative excitatory collateral connections from class cells Ω''_i to Ω_i and Ω'_i are learned, output from the synaptic matrix that does not have access to sensory input (module 3) can indirectly evoke the sensory images originally associated with the class cells in sensory modules 1 and 2. In networks of this kind, the dendrites of class cells would have adaptive synapses just as mosaic cells do and would receive axon collaterals as feedback from the synaptic matrix at the next higher level (figure 3.6). Thus, a backward chain of selective excitation from higher-level class cells to lower-level class cells can elicit neuronal analogs of the appropriate




Figure 3.6

Schematic of a synaptic matrix in a network with feedback from a higher level to a lower level through adaptive class cells.

environmental stimuli. That analogs of endosomatic stimulus states (kinesthetic, somasthetic, vestibular) can also be recovered in the same fashion is implicit in this account.

The type-token distinction (Jackendoff 1985) can be roughly captured in a natural way in synaptic matrix networks of this kind. Class cells that signal particular exemplars of an object type (as in the sensory modules) can be taken as tokens of those objects, whereas class cells that fire in response to the input of any of a particular subset of class cell tokens are physical signals of the subset, and in effect, their activity announces a type name in the biological system. As synaptic matrices are cascaded, class cells can signal types of types (figure 3.7). Thus, neuronal representations at increasingly higher levels of abstraction can be instantiated in these putative brain networks. For example, in a feed-forward chain of synaptic matrices, my pet cat named Duffy can be represented by a series of distinct class cell signals as follows: Duffy (sitting on the floor in front of me) $\rightarrow \Omega_1$ (Duffy) $\rightarrow \Omega_2$ (cat) $\rightarrow \Omega_3$ (mammal) $\rightarrow \Omega_4$ (animal) $\rightarrow \Omega_5$ (animate). Moreover, given the backward-chaining mechanism illustrated in figure 3.6, excitation of any class cell in the chain can evoke appropriate subordinate representations as well as those that are superordinate.

Selective Association of Input and Output

So far, we have considered only models in which synaptic modification (learning) occurs as a result of an increase in arousal (excitatory bias) on all adaptive cells in the synaptic matrix. Since this kind of neuronal priming is diffuse, which particular filter cells will happen to fire in the presence of a stimulus and thus undergo a change in their distribution of synaptic transfer weights is a matter of chance (though the specific input-output mapping, once established, becomes a systematic property of the matrix). However, in order to achieve adaptive goals, it is necessary that particular stimuli be mapped onto just those output lines (class cells) that can evoke (either directly or indirectly) the appropriate behavioral responses.

Let us say that I already have vocalization routines in my motor repertoire that allow me to utter the words *Duffy* and *cat*. The problem is to forge privileged excitation paths from the retinal excitation pattern of Duffy (object) through its class cell token (Ω_1) to a cell that gates the utterance *Duffy* and from the first token (Ω_1) through another class cell (Ω_2 , at a higher level of abstraction) to a cell that gates the utterance *cat*. The desired selection of associations will occur if, instead of initiating diffuse priming (general excitatory bias) of filter cells (f_i), only those f_i that are coupled to the appropriate class cells

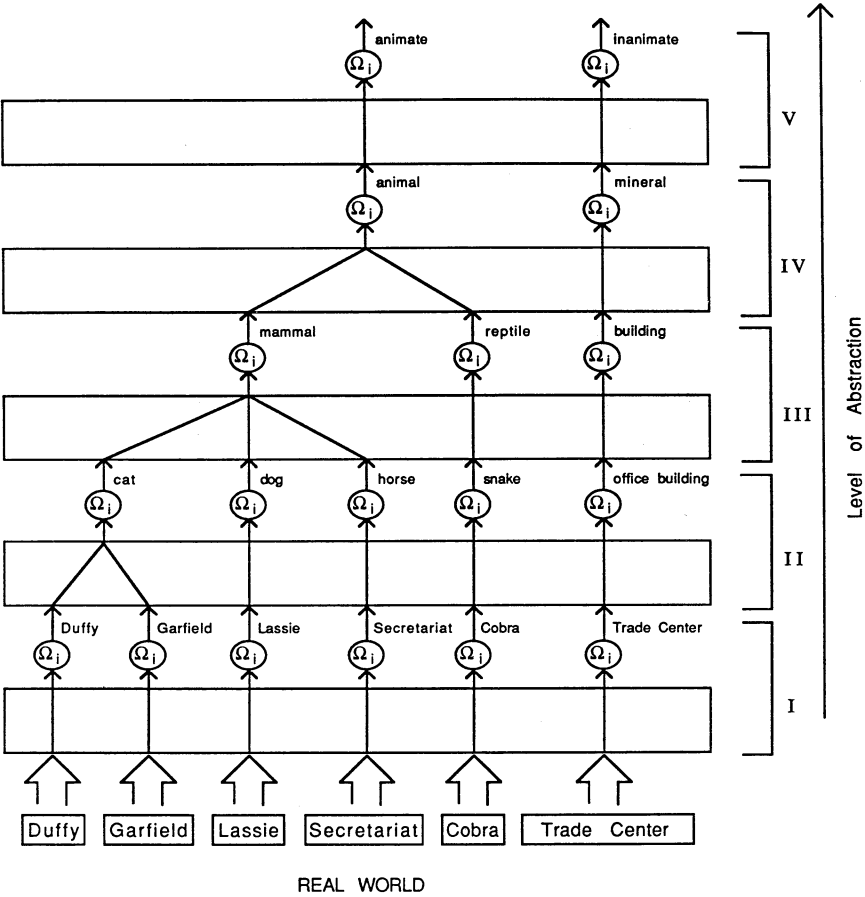


Figure 3.7
Flow diagram through five synaptic matrices. Token mappings from sensory input to progressively higher levels of abstraction.

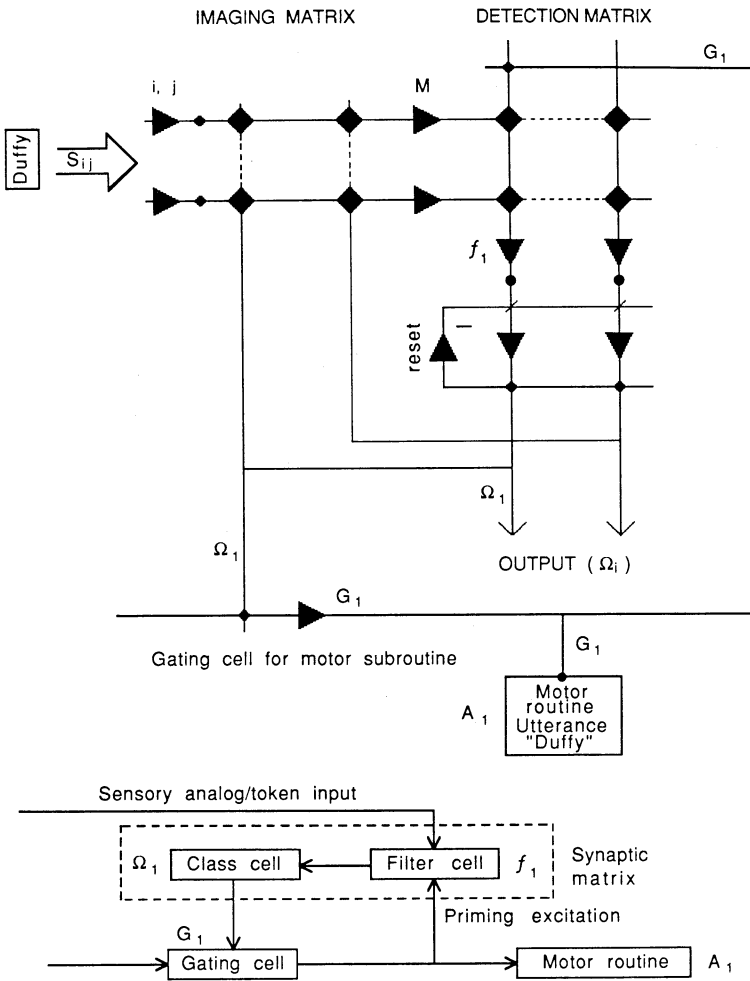


Figure 3.8

Top: Gating cell G_1 selectively primes a filter cell-class cell channel dedicated to a particular motor routine (A_1). Bottom: Simplified flow diagram of sensory motor loop. Analog (S_{ij}) or propositional (Ω) input learned by filter cell (f_1) becomes a potential trigger for its associated behavior.

are primed (in this case, f_1 and f_2). This can be accomplished if there is a collateral axonal-input line leading from each cell that gates a specific motor routine to a discrete filter cell in a synaptic matrix that is an adaptive processor of the retinal input. Figure 3.8 shows how the basic synaptic matrix is augmented to learn this kind of stimulus-response selectivity. In these networks, the initiation of a particular behavior in the presence of a particular stimulus (or a token of it) will induce a privileged excitation path from a class cell token of the stimulus to its coactivated behavior. Thus, if I say *Duffy* (vocally or subvocally) when I see the cat that is to be my household pet, this animal is labeled in my brain in two different ways: by a biological "name" (token Ω_1) and by a specific efferent routine that can generate the signal utterance *Duffy*. It is a specific claim of the model that all learned sensory/token-motor behaviors are mediated by action-dedicated neuronal loops of the form $G_1 \rightarrow f_1 \rightarrow \Omega_1 \rightarrow G_1$. After the filter cell in such a loop has been tuned to a particular input, any subsequent instance of this input can evoke the appropriate action.

Putting Motivation into the Network

In the previous examples, we considered two principal kinds of inputs to the synaptic matrix: exteroceptive sensory stimuli and class cell tokens of such stimuli. But motivational states produce specific interoceptive stimuli, and these can be processed in the same way as exteroceptive inputs. Thus, for example, an interoceptive vector that corresponds to the physiological state of hunger can be coded (mapped) to a class cell in a detection matrix, and activation of this cell can serve as a token in subsequent adaptive associations (Trehub 1970). For example, say that a hungry person wants to eat. Because there must be food available to eat, the joint tokens of hunger (Ω_1) and available food (Ω_2) must be activated to initiate the motor subroutines that constitute the act of eating. The double $\Omega_1\Omega_2$ tokens could constitute a pattern that is mapped to Ω_3 in a sensory motor loop that gates the action of eating. In this case, Ω_1 alone would be in a sensory motor loop that gates the look-for-food subroutines (go to the refrigerator, open the door, take an apple). After completion of the search behavior induced by Ω_1 , the sensory input evoked by the apple would activate Ω_2 , the condition for discharge of Ω_3 would be established, and the apple would be eaten.

This is a glimpse of how motivational and sensory tokens can be processed as joint patterns. As the structure and dynamics of the model are elaborated, we will see that high-level, goal-directed behavior can be initiated and adaptively controlled in networks of synaptic matrices.