

Normal and amnesic learning, recognition and memory by a neural model of cortico-hippocampal interactions

Gail A. Carpenter and Stephen Grossberg

The processes by which humans and other primates learn to recognize objects have been the subject of many models. Processes such as learning, categorization, attention, memory search, expectation and novelty detection work together at different stages to realize object recognition. In this article, Gail Carpenter and Stephen Grossberg describe one such class of model (Adaptive Resonance Theory, ART) and discuss how its structure and function might relate to known neurological learning and memory processes, such as how inferotemporal cortex can recognize both specialized and abstract information, and how medial temporal amnesia might be caused by lesions in the hippocampal formation. This model also suggests how hippocampal and inferotemporal processing might be linked during recognition learning.

A central problem in cognitive neuroscience concerns the processes whereby normal humans and other primates learn to recognize objects, and how these processes break down in different types of amnesic patients. The complexity of these processes has led to the development of neural models that might shed light on these issues. This article focuses on how one particular class of neural models, called Adaptive Resonance Theory (ART) models, can be applied to this task. ART models have been used to help explain and predict a large body of cognitive and neural data about recognition learning, attention and memory search¹⁻⁴. ART systems accomplish this synthesis by developing a solution to a fundamental problem about learning and memory that is called the stability-plasticity dilemma. An adequate self-organizing recognition system must

be capable of plasticity in order to rapidly learn about significant new events, yet its memory must also remain stable in response to irrelevant or often repeated events. For example, how do we learn to recognize new faces without risking unselectively forgetting the faces of our parents? In order to prevent the unselective forgetting of its learned codes by the 'blooming, buzzing confusion' of irrelevant experience, an ART system is sensitive to novelty. It is capable of distinguishing between familiar and unfamiliar events, as well as between expected and unexpected events.

The importance of expectancy and novelty-related processes in conditioning and cognitive processes has been extensively documented since the pioneering work of Tolman⁵, Sokolov^{6,7} and Vinogradova⁸. In ART, interactions between an attentional subsystem and an orienting subsystem, or novelty detector, self-stabilize learning, without an external teacher, as the network familiarizes itself with an environment by categorizing the information within it in a way that leads to behavioral success⁹. This learning system combines several types of processes that have been demonstrated in cognitive and neurobiological experiments, but not synthesized into a model system.

Competitive learning and self-organizing feature maps

All learning takes place in the attentional subsystem. Its processes include activation of short-

Gail A. Carpenter and Stephen Grossberg are at the Center for Adaptive Systems and the Dept of Cognitive and Neural Systems, Boston University, 111 Cummington St, Boston, MA 02215, USA.




Fig. 1. Interactions between the attentional and orienting subsystems of an adaptive resonance theory (ART) circuit. Level F_1 encodes a distributed representation of an event by a short-term memory (STM) activation pattern across a network of feature detectors. Level F_2 encodes the event using a compressed STM representation of the F_1 pattern. Learning of these recognition codes takes place at the long-term memory (LTM) traces within the bottom-up and top-down pathways between levels F_1 and F_2 . The top-down pathways read out learned expectations whose prototypes are matched against bottom-up input patterns at F_1 . The size of mismatches in response to novel events are evaluated relative to the vigilance parameter ρ of the orienting subsystem A. A large enough mismatch resets the recognition code that is active in STM at F_2 and initiates a memory search for a more appropriate recognition code. Output from subsystem A can also trigger an orienting response. (A) Block diagram of circuit. (B) Individual pathways of circuit, including the input level F_0 that generates inputs to level F_1 . The gain control input g_1 to level F_1 helps to instantiate the 2/3 Rule (see text). Gain control g_2 to level F_2 is needed to instate a category in STM.

term memory (STM) traces, incorporation through learning of momentary STM information into longer lasting long-term memory (LTM) traces, and interactions between pathways that carry specific information with nonspecific pathways that modulate the specific pathways. These interactions between specific STM and LTM processes and nonspecific modulatory processes regulate the stability-plasticity balance during normal learning.

The attentional subsystem undergoes both bottom-up learning and top-down learning between the processing levels denoted by F_1 and F_2 in Fig. 1. Level F_1 contains a network of nodes, or cell populations, each of which represents a particular combination of sensory features. Level F_2 contains a network of nodes that represent recognition codes, or categories, that are selectively activated by the activation patterns across F_1 . Each F_1 node sends output signals to a subset of F_2 nodes. Each F_2 node thus receives inputs from many F_1 nodes. The thick arrow from F_1 to F_2 in Fig. 1A represents in a concise way the array of diverging and converging pathways shown in Fig. 1B. Learning takes place at the synapses denoted by semicircular endings in the $F_1 \rightarrow F_2$ pathways. This bottom-up learning enables F_2 nodes to become selectively tuned to particular combinations of activation patterns across F_1 by changing their LTM traces.

Why is not just bottom-up learning sufficient? This analysis was carried out in a type of model that is often called a self-organizing feature map, competitive learning or learned vector quantization. Such a model shows how to combine associative learning and lateral inhibition for purposes of learned categorization.

In such a model, as shown in Fig. 2A, an input pattern registers itself as a pattern of activity, or STM, across the feature detectors of level F_1 . Each F_1 output signal is multiplied or gated, by the adaptive weight, or LTM trace, in its respective pathway. All these LTM-gated inputs are added up at their target F_2 nodes. Lateral inhibitory, or competitive, interactions within F_2 contrast-enhance this input pattern. Whereas many F_2 nodes may receive inputs from F_1 , lateral inhibition allows a much smaller set of F_2 nodes to store their activation in STM.

Only the F_2 nodes that win the competition and store their activity in STM can influence the learning process. STM activity opens a learning gate at the LTM traces that abut the winning nodes. These LTM traces can then approach, or track, the input signals in their pathways, a process called 'steepest descent'. This learning law is thus often called 'gated steepest descent' or 'instar learning'. It was introduced into neural-network models in the 1960s¹⁰ and is the learning law that was used to introduce ART (Refs 11, 12). Such an LTM trace can either increase or decrease to track the signals in its pathway. Thus, it is not a Hebbian associative law. It has been used to model neurophysiological data about hippocampal LTP (Refs 13, 14) and adaptive tuning of cortical feature detectors during the visual critical

period^{15,16}, lending support to ART predictions that both systems would employ such a learning law^{11,12}.

Self-organizing feature-map models were introduced and computationally characterized by Grossberg^{12,17,18}, von der Malsburg¹⁹, and Willshaw and von der Malsburg²⁰. These models were subsequently applied and further developed by many authors²¹⁻²⁶. They exhibit many useful properties, especially if not too many input patterns, or clusters of input patterns, perturb level F_1 relative to the number of categorizing nodes in level F_2 . It was proved that under these sparse environmental conditions, category learning is stable, with self-normalizing LTM traces that track the statistics of the environment and oscillate a minimum number of times^{11,12,18}. Also, the category selection rule, like a Bayesian classifier, tends to minimize error. It was also proved, however, that under arbitrary environmental conditions, learning becomes unstable. Such a model could forget the faces of your parents. Although a gradual switching off of plasticity can partially overcome this problem, such a mechanism cannot work in a recognition learning system whose plasticity is maintained throughout adulthood.

This memory instability is due to basic properties of associative learning and lateral inhibition. An analysis of this instability, together with data about categorization, conditioning and attention, led to the introduction of ART models that stabilize the memory of self-organizing feature maps in response to an arbitrary stream of input patterns^{11,12}.

Memory search, feature binding and attentional focusing

In an ART model^{2,27}, learning does not occur when some winning F_2 activities are stored in STM. Instead activation of F_2 nodes may be interpreted as 'making a hypothesis' about an input at F_1 . When F_2 is activated, it quickly generates an output pattern that is transmitted along the top-down adaptive pathways from F_2 to F_1 . These top-down signals are multiplied in their respective pathways by LTM traces at the semicircular synaptic knobs of Fig. 2B. The LTM-gated signals from all the active F_2 nodes are added to generate the total top-down feedback pattern from F_2 to F_1 . This pattern plays the role of a learned expectation. Activation of this expectation may be interpreted as 'testing the hypothesis', or 'reading out the prototype', of the active F_2 category. As shown in Fig. 2B, ART networks are designed to match the 'expected prototype' of the category against the bottom-up input pattern, or exemplar, to F_1 . Nodes that are activated by this exemplar are suppressed if they do not correspond to large LTM traces in the top-down prototype pattern. The resultant F_1 pattern encodes the cluster of input features that the network deems relevant to the hypothesis based upon its past experience. This resultant activity pattern, called X^* in Fig. 2B, encodes the pattern of features to which the network 'pays attention'.

If the expectation is close enough to the input exemplar, then a state of resonance develops as the attentional focus takes hold. The pattern X^* of




Fig. 2. ART search for an F_2 recognition code. (A) The input pattern I generates the specific STM activity pattern X at F_1 , as it nonspecifically activates the orienting subsystem A . X is represented by the tinting across F_1 . Pattern X both inhibits A and generates the output pattern S . Pattern S is transformed by the LTM traces into the input pattern T , which activates the STM pattern Y across F_2 . (B) Pattern Y generates the top-down output pattern U which is transformed into the prototype pattern V . If V mismatches I at F_1 , then a new STM activity pattern X^* is generated at F_1 . X^* is represented by the tinted region. Inactive nodes corresponding to X are untinted. The reduction in total STM activity that occurs when X is transformed into X^* causes a decrease in the total inhibition from F_1 to A . (C) If the vigilance criterion fails to be met, A releases a nonspecific arousal wave to F_2 , which resets the STM pattern Y at F_2 . (D) After Y is inhibited, its top-down prototype signal is eliminated, and X can be reinstated at F_1 . Ending traces of the prior reset lead X to activate a different STM pattern Y^* at F_2 . If the top-down prototype due to Y^* also mismatches I at F_1 , then the search for an appropriate F_2 code continues until a more appropriate F_2 representation is selected. At this point an attentive resonance develops and learning of the attended data is initiated.

attended features reactivates the F_2 category Y which, in turn, reactivates X^* . The network locks into a resonant state through a positive feedback loop that dynamically links, or binds, X^* with Y . Damasio²⁸ has used the term 'convergence zones' to describe such a resonant process. The resonance binds spatially distributed features into either a stable equilibrium or a synchronous oscillation²⁹⁻³¹, much like synchronous feature-binding in visual cortex³²⁻³⁴.

In ART, the resonant state, rather than bottom-up activation, drives the learning process. The resonant state persists long enough, at a high enough activity level, to activate the slower learning process; hence the term adaptive resonance theory. ART systems learn prototypes, rather than exemplars, because the attended feature vector X^* , rather than the

Box 1. ARTMAP benchmark studies

Some machine-learning benchmark studies^{a,b} are given below that compare the performances of supervised ART, or ARTMAP, models with those of alternative models. These benchmarks describe how well these systems predict test sets when they experience equivalent training sets (as in benchmarks 1–4) and the number of epochs, or repetitions of the training set, that are needed to reach the same level of accuracy (benchmark 5).

(1) Medical database

Mortality following coronary bypass grafting surgery
Fuzzy ARTMAP significantly outperforms:

- Logistic regression
- Additive model
- Bayesian assignment
- Cluster analysis
- Classification and regression trees
- Expert panel-derived sickness scores
- Principal component analysis

(2) Mushroom database

Decision trees (90–95% correct)
ARTMAP (100% correct; training set an order of magnitude smaller)

(3) Letter recognition database

Genetic algorithm (82% correct)
Fuzzy ARTMAP (96% correct)

(4) Circle-in-the-square task

Back propagation (90% correct)
Fuzzy ARTMAP (99.5% correct)

(5) Two-spiral task

Back propagation (10 000–20 000 training epochs)
Fuzzy ARTMAP (1–5 training epochs)

References

- a Carpenter, G. A., Grossberg, S. and Reynolds, J. (1991) *Neur. Networks* 4, 565–588
b Carpenter, G. A., Grossberg, S., Markuzon, N., Reynolds, J. H. and Rosen, D. B. (1992) *IEEE Trans. Neur. Networks* 3, 698–713

input exemplar itself, is learned. These prototypes may, however, also be used to encode individual exemplars. How the matching process achieves this is described below. If the mismatch between bottom-up and top-down information is too great, then resonance cannot develop. Instead the F_2 category is quickly reset and a memory search for a better category is initiated. This combination of top-down matching, attention focusing and memory search is what stabilizes ART learning and memory in an arbitrary input environment.

The stabilizing properties of top-down matching may be one reason for the ubiquitous occurrence of reciprocal bottom-up and top-down cortico-cortical and cortico-thalamic processes^{35,36}. Resonant attention has also been suggested to be necessary for conscious experience. The predicted linkage^{9,11} between learning, attention, consciousness and synchronous oscillations has recently attracted much interest³⁷.

Matching, priming and phonemic restoration

The ART attentive matching process is realized by combining bottom-up inputs and top-down expect-

tations with a nonspecific arousal process that is called attentional gain control^{2,27}. An F_1 node can be fully activated only if two of the three input sources that converge on the node send positive signals to the node at a given time. This constraint is called the 2/3 Rule. A bottom-up input pattern turns on the attentional gain control channel in order to instate itself in STM at F_1 (Fig. 2A). A top-down expectation turns off the attentional gain control channel (Fig. 2B). As a result, only those input features that are confirmed by the top-down prototype can be attended at F_1 after an F_2 category is selected.

The 2/3 Rule, first and foremost, enables an ART network to solve the stability–plasticity dilemma. Carpenter and Grossberg²⁷ proved that ART learning and memory are stable in arbitrary environments, but become unstable when 2/3 Rule matching is eliminated. Thus a type of matching that guarantees stable learning also enables the network to pay attention.

In the brain, 2/3 Rule matching is illustrated by experiments on phonemic restoration^{38–42}. Suppose that a noise spectrum replaces a letter sound in a word heard in an otherwise unambiguous context. Then subjects hear the correct letter sound, not the noise, to the extent that the noise spectrum includes the letter formants. If silence replaces the noise, then only silence is heard. Top-down expectations thus amplify expected input features while suppressing unexpected features, but do not create activations not already in the input.

Matching by the 2/3 Rule also explains paradoxical reaction time and error data from priming experiments during lexical decision and letter gap detection tasks^{43,44}. Although priming is often thought of as a residual effect of previous bottom-up activation, a combination of bottom-up activation and top-down 2/3 Rule matching was needed to explain the complete data pattern. This analysis combined bottom-up priming with a type of top-down priming; namely, the top-down activation that prepares a network for an expected event that may or may not occur. The 2/3 Rule clarifies why top-down priming, by itself, is subliminal and unconscious, even though it can facilitate supraliminal processing of a subsequent expected event.

Vigilance, memory search and generalization

The criterion of an acceptable 2/3 Rule match is defined by a parameter ρ called 'vigilance'^{2,27}. The vigilance parameter is computed in the orienting subsystem A. Vigilance weighs how similar an input exemplar I must be to a top-down prototype V in order for resonance to occur. Resonance occurs if $\rho|| - |X^*| \leq 0$. This inequality says that the F_1 attentional focus X^* inhibits A more than the input I excites it. If A remains quiet, then an $F_1 \leftrightarrow F_2$ resonance can develop.

Vigilance calibrates how much novelty the system can tolerate before activating A and searching for a different category. If the top-down expectation and the bottom-up input are too different to resonate,

then hypothesis testing, or memory search, is triggered. During search, the orienting subsystem interacts with the attentional subsystem (Figs 2C,D) to rapidly reset mismatched categories and to select better F_2 representations with which to learn about novel events at F_1 , without risking unselective forgetting of previous knowledge. Search may select a familiar category if its prototype is similar enough to the input to satisfy the resonance criterion. The prototype may then be refined by 2/3 Rule attentional focusing. If the input is too different from any previously learned prototype, then an uncommitted population of F_2 cells is selected and learning of a new category is initiated.

Because vigilance can vary across learning trials, recognition categories capable of encoding widely differing degrees of generalization or abstraction can be learned by a single ART system. Low vigilance leads to broad generalization and abstract prototypes. High vigilance leads to narrow generalization and to prototypes that represent fewer input exemplars, even a single exemplar. Thus a single ART system may be used, say, to recognize abstract categories of faces and dogs, as well as individual faces and dogs. A single system can learn both, as the need arises, by increasing vigilance just enough to activate A if a previous categorization leads to a predictive error⁴⁵⁻⁴⁷.

ART systems provide a new answer to the question of whether the brain learns prototypes or exemplars. Various authors have realized that neither alternative is satisfactory, and that a hybrid system is needed⁴⁸. ART systems can perform this hybrid function in a manner that is sensitive to environmental demands. Box 1 summarizes how such a supervised ART system performs relative to other machine learning, genetic algorithm and back propagation networks in benchmark simulations.

Memory consolidation and direct access to familiar categories

As inputs are practiced over learning trials, the search process eventually converges upon stable categories. Familiar inputs directly access the category whose prototype provides the globally best match, while unfamiliar inputs trigger memory searches for better categories, until the memory capacity is fully utilized². The process whereby search is automatically disengaged is a form of memory consolidation that emerges from network interactions. Emergent consolidation does not preclude structural consolidation at individual cells, since persistent resonance may be a trigger for learning-dependent cellular processes.

Face recognition and inferotemporal cortex

Level F_2 properties may be compared with properties of cell activations in inferotemporal cortex (IT) during recognition learning in monkeys. The ability of F_2 nodes to learn categories with different levels of generalization clarifies how some IT cells can exhibit high specificity, such as selectivity to views of particular faces, while other cells respond to broader features of the animal's environment⁴⁹⁻⁵⁷. Moreover,

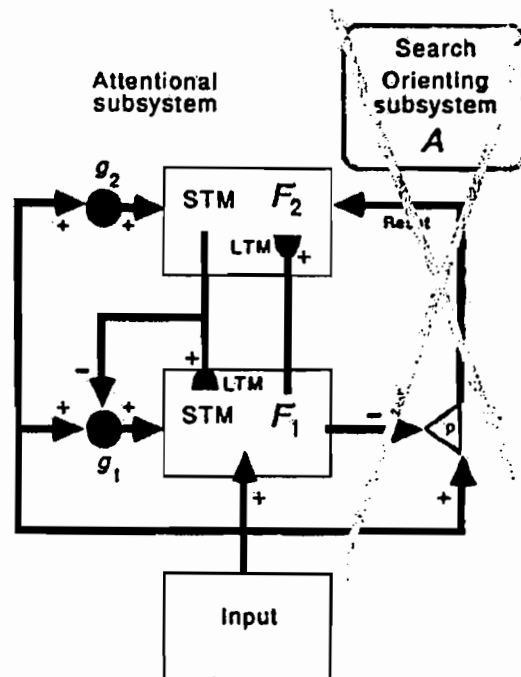


Fig. 3. A memory disturbance with formal symptoms similar to those of medial temporal amnesia is caused by a lesion of the model's orienting subsystem. The symptoms are emergent properties due to interactions among the non-lesioned network components. The formal amnesic syndrome is strikingly similar to the one caused in humans and monkeys by lesioning the hippocampal formation.

when monkeys are exposed to easy and difficult discriminations, 'in the difficult condition the animals adopted a stricter internal criterion for discriminating matching from nonmatching stimuli . . . the animals' internal representations of the stimuli were better separated, independent of the criterion used to discriminate them . . . Increased effort appears to cause enhancement of the responses and sharpened selectivity for attended stimuli' (Ref. 58; pp. 339-340). These are also properties of model cells in F_2 . Prototypes represent smaller sets of exemplars at higher vigilance levels, so a stricter matching criterion is learned. These exemplars match their finer prototypes better than do exemplars which match a coarser prototype. This better match more strongly activates the corresponding F_2 nodes.

Data from IT support the hypothesis that unfamiliar or unexpected stimuli nonspecifically activate level F_2 via the orienting subsystem. As Desimone has noted, 'the fact that IT cortex has a reduced level of activation for familiar or expected stimuli suggests that a high level of cortical activation may itself serve as a trigger for attentional and orienting systems, causing the subject to orient to the stimulus causing the activation. This link between the mnemonic and attentional systems would "close the loop" between the two systems, resulting in orienting behavior that is influenced by both current stimuli and prior memories. Such a mechanism has a number of similarities to the adaptive

resonance theory' (Ref. 59; p. 359). IT cells during working memory tasks are reset after each trial^{60,61}. Reset also occurs in ART. Some data suggest that the pulvinar may mediate attentional gain^{59,62}.

Orienting, the hippocampus and amnesia

The hypothesis that the ART orienting system has a neural analog in the hippocampal formation has considerable experimental support. A lesion of the ART orienting subsystem (Fig. 3) creates formal symptoms like those of humans with medial temporal amnesia, including: unlimited anterograde amnesia; limited retrograde amnesia; failure of consolidation; tendency to learn the first event in a series; abnormal reactions to novelty, including perseverative reactions; normal priming; and normal information processing of familiar events⁶³⁻⁷¹. Unlimited anterograde amnesia occurs because the network cannot carry out the memory search to learn a new recognition code. Limited retrograde amnesia occurs because familiar events can directly access correct recognition codes². Before events become familiar, memory consolidation occurs that utilizes the orienting subsystem (Fig. 2C). This failure of consolidation does not necessarily prevent learning *per se*. Instead, learning influences the first recognition category activated by bottom-up processing, much as 'amnesics are particularly strongly wedded to the first response they learn' (Ref. 72; p. 253). Perseverative reactions can occur because the orienting subsystem cannot reset sensory representations or top-down expectations that may be persistently mismatched by bottom-up cues. The inability to search memory prevents ART from discovering more appropriate stimulus combinations to attend to. Normal priming occurs because it is mediated by the attentional subsystem.

Similar behavioral problems have been identified in hippocampectomized monkeys. Gaffan noted that fornix transection 'impairs ability to change an established habit... in a different set of circumstances that is similar to the first and therefore liable to be confused with it' (Ref. 73; p. 94). In ART, a defective orienting subsystem prevents the memory search whereby different representations could be learned for similar events. Pribram called such a process a 'competence for recombinant context-sensitive processing' (Ref. 74; p. 362). These ART mechanisms illustrate how memory consolidation and novelty detection may be mediated by the same neural structures⁷¹, why hippocampectomized rats have difficulty orienting to novel cues⁷⁵, and why there is a progressive reduction in novelty-related hippocampal potentials as learning proceeds in normal rats^{76,77}. In ART, the orienting system is automatically disengaged as events become familiar during the memory consolidation process.

In summary, the hypothesis that the hippocampal formation is linked to orienting subsystem functions⁷⁸ helps to explain amnesic symptoms as manifestations of a breakdown in the orienting and memory search mechanisms that normally solve the stability-plasticity dilemma. This interpretation does not contradict other data that suggest additional

functions for the hippocampal formation^{75,79}. A hippocampal role in adaptive timing, conditioned reinforcement, spatial approach and avoidance, and attentional blocking has been mechanistically outlined within the larger model-neural system that includes ART recognition networks^{1-4,78,80}. Such a hybrid function is consistent with data about hippocampal cells with place fields in a radial-arm maze and conditioned responses in classical conditioning tasks⁸¹. These results clarify how the hippocampus may subserve LTP-based learning, without suggesting that it temporarily stores recognition codes of many types of sensory events until these memories can consolidate in their respective sensory cortices. The disengagement of the orienting subsystem during memory consolidation does not imply that the orienting subsystem ever learns a sensory recognition code.

This larger model system also includes spatial and motor learning circuits^{23,82,83} whose properties shed new light on the popular distinctions between knowing that and knowing how⁸⁴, memory with record and memory without record⁸⁵, taxon and locale⁷⁵, memory and habit⁸⁶, and declarative memory and procedural memory⁸⁷ by clarifying aspects of how these distinct processes work and interact.

Concluding remarks

Many properties of data about the inferotemporal cortex and the hippocampal formation are rationalized by the ART circuits that solve the stability-plasticity dilemma. These model circuits also suggest predictions that may be tested by novel neurobiological experiments. For example, varying the vigilance parameter of the orienting subsystem alters the specificity of recognition codes that are learned by the attentional subsystem by calibrating how different an input needs to be from a prototype before the orienting subsystem triggers search. This property suggests that operations which make the novelty-related potentials of the hippocampus more sensitive to input changes may trigger the formation of more selective inferotemporal recognition categories. Can such a correlation be recorded, say, when monkeys learn easy and hard discriminations? Conversely, operations that progressively block the expression of hippocampal novelty potentials may lead to the learning of coarser recognition categories, with amnesic symptoms as a limiting case.

Selected references

- 1 Grossberg, S. (ed.) (1987) *The Adaptive Brain, Volumes I and II*, Elsevier
- 2 Carpenter, G. A. and Grossberg, S. (eds) (1992) *Pattern Recognition by Self-Organizing Neural Networks*, MIT Press
- 3 Grossberg, S. (1988) *Neural Networks and Natural Intelligence*, MIT Press
- 4 Grossberg, S. (1982) *Studies of Mind and Brain: Neural Principles of Learning, Perception, Development, Cognition, and Motor Control*, Reidel Press
- 5 Tolman, E. C. (1932) *Purposive Behavior in Animals and Men*, Appleton-Century-Crofts
- 6 Sokolov, E. N. (1958) *Perception and the Conditioned Reflex*, Moscow University Press
- 7 Sokolov, E. N. (1968) *Mechanisms of Memory*, Moscow University Press
- 8 Vinogradova, O. S. (1975) in *The Hippocampus*, Vol. 2

- (Isaacson, R. L. and Pribram, K. H., eds), pp. 3–69, Plenum Press
- 9 Grossberg, S. (1980) *Psych. Rev.* 87, 1–51
 - 10 Grossberg, S. (1969) *J. Stat. Physics* 1, 319–350
 - 11 Grossberg, S. (1976) *Biol. Cybern.* 23, 121–134
 - 12 Grossberg, S. (1976) *Biol. Cybern.* 23, 187–202
 - 13 Levy, W. B. (1985) in *Synaptic Modification, Neuron Selectivity, and Nervous System Organization* (Levy, W. B., Anderson, J. and Lehmkuhle, S., eds), pp. 5–33, Lawrence Erlbaum Associates
 - 14 Levy, W. B. and Desmond, N. L. (1985) in *Synaptic Modification, Neuron Selectivity, and Nervous System Organization* (Levy, W. B., Anderson, J. and Lehmkuhle, S., eds), pp. 105–121, Lawrence Erlbaum Associates
 - 15 Rauschecker, J. P. and Singer, W. (1979) *Nature* 280, 58–60
 - 16 Singer, W. (1983) in *Synergetics of the Brain* (Basar, E., Flohr, H., Haken, H. and Mandell, A. J., eds), pp. 89–101, Springer-Verlag
 - 17 Grossberg, S. (1972) *Kybernetik* 10, 49–57
 - 18 Grossberg, S. (1978) in *Progress in Theoretical Biology, Vol. 5* (Rosen, R. and Snell, F., eds), pp. 233–374, Academic Press
 - 19 von der Malsburg, C. (1973) *Kybernetik* 14, 85–100
 - 20 Willshaw, D. J. and von der Malsburg, C. (1976) *Proc. R. Soc. London Ser. B* 194, 431–445
 - 21 Amani, S. and Takeuchi, A. (1978) *Biol. Cybern.* 29, 127–136
 - 22 Bienenstock, E. L., Cooper, L. N. and Munro, P. W. (1982) *J. Neurosci.* 2, 32–48
 - 23 Grossberg, S. and Kuperstein, M. (1986) *Neural Dynamics of Adaptive Sensory-Motor Control*, North-Holland; expanded edition (1989), Pergamon Press
 - 24 Kohonen, T. (1984) *Self-Organization and Associative Memory*, Springer-Verlag
 - 25 Linsker, R. (1986) *Proc. Natl Acad. Sci. USA* 83, 8779–8783
 - 26 Rumelhart, D. E. and Zipser, D. (1985) *Cogn. Sci.* 9, 75–112
 - 27 Carpenter, G. A. and Grossberg, S. (1987) *Comput. Vis. Graph. Image Process.* 37, 54–115
 - 28 Damasio, A. R. (1989) *Neur. Comp.* 1, 123–132
 - 29 Eckhorn, R. and Schanze, T. (1991) in *Self-Organization, Emerging Properties, and Learning* (Babloyantz, A., ed.), Plenum Press
 - 30 Grossberg, S. and Somers, D. (1991) *Neur. Networks* 4, 453–466
 - 31 Grossberg, S. and Somers, D. (1992) in *Neural Networks for Vision and Image Processing* (Carpenter, G. A. and Grossberg, S., eds), pp. 385–405, MIT Press
 - 32 Eckhorn, R. et al. (1988) *Biol. Cybern.* 60, 121–130
 - 33 Gray, C. M. and Singer, W. (1989) *Proc. Natl Acad. Sci. USA* 86, 1698–1702
 - 34 Gray, C. M., Konig, P., Engel, A. K. and Singer, W. (1989) *Nature* 338, 334–337
 - 35 Macchi, G. and Rinvik, E. (1976) *Handbook of Electroencephalography and Clinical Neurophysiology, Vol. 2, Pt 4* (Rémond, A., ed.), Elsevier
 - 36 Tsumoto, T., Creutzfeldt, O. D. and Legendy, C. R. (1978) *Exp. Brain Res.* 25, 291–306
 - 37 Crick, F. and Koch, C. (1990) *Cold Spring Harbor Symp. Quant. Biol.* 55, 953–962
 - 38 Repp, B. H. (1991) *Haskins Laboratories Status Report on Speech Research SR-107/108*, 147–170
 - 39 Samuel, A. G. (1981) *J. Exp. Psychol. General* 110, 474–494
 - 40 Samuel, A. G. (1981) *J. Exp. Psychol. Hum. Percept. Perform.* 7, 1124–1131
 - 41 Warren, R. M. (1984) *Psych. Bull.* 96, 371–383
 - 42 Warren, R. M. and Sherman, G. L. (1974) *Percept. Psychophys.* 16, 150–156
 - 43 Grossberg, S. and Stone, G. O. (1986) *Psych. Rev.* 93, 46–74
 - 44 Schvaneveldt, R. W. and MacDonald, J. E. (1981) *J. Exp. Psychol. Hum. Percept. Perform.* 7, 673–687
 - 45 Carpenter, G. A. and Grossberg, S. (1992) *IEEE Commun. Mag.* 30, 38–49
 - 46 Carpenter, G. A., Grossberg, S. and Reynolds, J. (1991) *Neur. Networks* 4, 565–588
 - 47 Carpenter, G. A., Grossberg, S., Markuzon, N., Reynolds, J. H. and Rosen, D. B. (1992) *IEEE Trans. Neur. Networks* 3, 698–713
 - 48 Smith, E. E. (1990) in *An Invitation to Cognitive Science* (Osherson, D. O. and Smith, E. E., eds), pp. 33–53, MIT Press
 - 49 Desimone, R. (1991) *J. Cogn. Neurosci.* 3, 1–8
 - 50 Desimone, R. and Ungerleider, L. G. (1989) in *Handbook of Neuropsychology, Vol. 2* (Boller, F. and Grafman, J., eds), pp. 267–299, Elsevier
 - 51 Gochin, P. M., Miller, E. K., Gross, C. G. and Gerstein, G. L. (1991) *Exp. Brain Res.* 84, 505–516
 - 52 Harries, M. H. and Perrett, D. I. (1991) *J. Cogn. Neurosci.* 3, 9–24
 - 53 Mishkin, M. (1982) *Philos. Trans. R. Soc. London Ser. B* 298, 85–95
 - 54 Mishkin, M. and Appenzeller, T. (1987) *Sci. Am.* 256, 80–89
 - 55 Perrett, D. I., Mistlin, A. J. and Chitty, A. J. (1987) *Trends Neurosci.* 10, 358–364
 - 56 Schwartz, E. L., Desimone, R., Albright, T. and Gross, C. G. (1983) *Proc. Natl Acad. Sci. USA* 80, 5776–5778
 - 57 Seibert, M. and Waxman, A. M. (1991) in *Neural Networks for Perception, Vol. 1* (Wechsler, H., ed.), pp. 426–444, Academic Press
 - 58 Spitzer, H., Desimone, R. and Moran, J. (1988) *Science* 240, 338–340
 - 59 Desimone, R. (1992) in *Neural Networks for Vision and Image Processing* (Carpenter, G. A. and Grossberg, S., eds), pp. 343–364, MIT Press
 - 60 Miller, E. K., Li, L. and Desimone, R. (1991) *Science* 254, 1377–1379
 - 61 Riches, I. P., Wilson, F. A. W. and Brown, M. W. (1991) *J. Neurosci.* 11, 1763–1779
 - 62 Robinson, D. L. and Peterson, S. E. (1992) *Trends Neurosci.* 15, 127–132
 - 63 Cohen, N. J. (1984) in *The Neuropsychology of Memory* (Squire, L. and Butters, N., eds), pp. 83–103, Guilford Press
 - 64 Graf, P., Squire, L. R. and Mandler, G. (1984) *J. Exp. Psychol. Learn. Memory Cognition* 10, 164–178
 - 65 Lynch, G., McGaugh, J. L. and Weinberger, N. M. (eds) (1984) *Neurobiology of Learning and Memory*, Guilford Press
 - 66 Mattis, S. and Kovner, R. (1984) in *Neuropsychology of Memory* (Squire, L. and Butters, N., eds), pp. 115–121, Guilford Press
 - 67 Squire, L. R. and Butters, N. (eds) (1984) *Neuropsychology of Memory*, Guilford Press
 - 68 Squire, L. R. and Cohen, N. J. (1984) in *Neurobiology of Learning and Memory* (Lynch, G., McGaugh, J. and Weinberger, N. M., eds), pp. 3–64, Guilford Press
 - 69 Warrington, E. K. and Weiskrantz, L. (1970) *Nature* 228, 628–630
 - 70 Warrington, E. K. and Weiskrantz, L. (1974) *Neuropsychology* 12, 419–428
 - 71 Zola-Morgan, S. M. and Squire, L. R. (1990) *Science* 250, 288–290
 - 72 Gray, J. A. (1982) *The Neuropsychology of Anxiety: An Enquiry into the Functions of the Septo-Hippocampal System*, Oxford University Press
 - 73 Gaffan, D. (1985) *Philos. Trans. R. Soc. London Ser. B* 308, 87–99
 - 74 Pribram, K. H. (1986) in *The Hippocampus, Vol. 4* (Isaacson, R. L. and Pribram, K. H., eds), pp. 329–370, Plenum Press
 - 75 O'Keefe, J. and Nadel, L. (1978) *The Hippocampus as a Cognitive Map*, Oxford University Press
 - 76 Deadwyler, S. A., West, M. O. and Lynch, G. (1979) *Brain Res.* 169, 29–43
 - 77 Deadwyler, S. A., West, M. O. and Robinson, J. H. (1981) *Science* 211, 1181–1183
 - 78 Grossberg, S. (1975) *Int. Rev. Neurobiol.* 18, 263–327
 - 79 Isaacson, R. L. and Pribram, K. H. (eds) (1986) *The Hippocampus: Vol. 4*, Plenum Press
 - 80 Grossberg, S. and Merrill, J. W. M. (1992) *Cogn. Brain Res.* 1, 3–37
 - 81 Eichenbaum, H. and Cohen, N. J. (1988) *Trends Neurosci.* 11, 244–248
 - 82 Bullock, D. B. and Grossberg, S. (1992) *Hum. Mov. Sci.* 10, 3–53
 - 83 Bullock, D. B., Grossberg, S. and Guenther, F. J. *Cogn. Neurosci.* (in press)
 - 84 Ryle, G. (1949) *The Concept of Mind*, Hutchinson
 - 85 Bruner, J. S. (1969) *The Pathology of Memory* (Talland, G. A. and Waugh, N. C., eds), Academic Press
 - 86 Mishkin, M. (1982) *Philos. Trans. R. Soc. London Ser. B* 298, 85–95
 - 87 Squire, L. R. and Cohen, N. J. (1984) *Neurobiology of Learning and Memory* (Lynch, G., McGaugh, J. and Weinberger, N. M., eds), pp. 3–64, Guilford Press

Acknowledgements

This research was supported in part by the Air Force Office of Scientific Research (AFOSR 90-0175), British Petroleum (BP 89A-1204), DARPA (AFOSR 90-0083), the National Science Foundation (NSF IRI-90-00530) and the Office of Naval Research (ONR N00014-91-J-4100). The authors wish to thank Cynthia E. Bradford, Diana J. Meyers and Robin L. Locke for their valuable assistance in the preparation of the manuscript.