

Adaptive Resonance Theory: How a brain learns to consciously attend, learn, and recognize a changing world[☆]

Stephen Grossberg^{*}

Center for Adaptive Systems, Graduate Program in Cognitive and Neural Systems, Boston University, 677 Beacon Street, Boston, MA 02215, United States
Center for Computational Neuroscience and Neural Technology, and Department of Mathematics, Boston University, 677 Beacon Street, Boston, MA 02215, United States

ARTICLE INFO

Dedicated to Gail Carpenter in appreciation of her many fundamental contributions to Adaptive Resonance Theory.

Keywords:

Adaptive Resonance Theory
Consciousness
Learning
Expectation
Attention
Recognition
Synchrony
Working memory
Reinforcement learning
Adaptive timing
Working memory
Speech perception
Gamma and beta oscillations
Inferotemporal cortex
Parietal cortex
Prefrontal cortex
Entorhinal cortex
Hippocampal cortex
Amygdala
Basal ganglia

ABSTRACT

Adaptive Resonance Theory, or ART, is a cognitive and neural theory of how the brain autonomously learns to categorize, recognize, and predict objects and events in a changing world. This article reviews classical and recent developments of ART, and provides a synthesis of concepts, principles, mechanisms, architectures, and the interdisciplinary data bases that they have helped to explain and predict. The review illustrates that ART is currently the most highly developed cognitive and neural theory available, with the broadest explanatory and predictive range. Central to ART's predictive power is its ability to carry out fast, incremental, and stable unsupervised and supervised learning in response to a changing world. ART specifies mechanistic links between processes of consciousness, learning, expectation, attention, resonance, and synchrony during both unsupervised and supervised learning. ART provides functional and mechanistic explanations of such diverse topics as laminar cortical circuitry; invariant object and scenic gist learning and recognition; prototype, surface, and boundary attention; gamma and beta oscillations; learning of entorhinal grid cells and hippocampal place cells; computation of homologous spatial and temporal mechanisms in the entorhinal–hippocampal system; vigilance breakdowns during autism and medial temporal amnesia; cognitive–emotional interactions that focus attention on valued objects in an adaptively timed way; item–order–rank working memories and learned list chunks for the planning and control of sequences of linguistic, spatial, and motor information; conscious speech percepts that are influenced by future context; auditory streaming in noise during source segregation; and speaker normalization. Brain regions that are functionally described include visual and auditory neocortex; specific and nonspecific thalamic nuclei; inferotemporal, parietal, prefrontal, entorhinal, hippocampal, parahippocampal, perirhinal, and motor cortices; frontal eye fields; supplementary eye fields; amygdala; basal ganglia; cerebellum; and superior colliculus. Due to the complementary organization of the brain, ART does not describe many spatial and motor behaviors whose matching and learning laws differ from those of ART. ART algorithms for engineering and technology are listed, as are comparisons with other types of models.

© 2012 Elsevier Ltd. All rights reserved.

[☆] Invited article submitted to the 25th anniversary special issue of *Neural Networks*.

^{*} Correspondence to: Center for Adaptive Systems, Boston University, 677 Beacon Street, Boston, MA 02215, United States. Tel.: +1 617 353 7858; fax: +1 617 353 7755.

E-mail address: steve@bu.edu.

URL: <http://cns.bu.edu/~steve>.

1. Introduction: the stability–plasticity dilemma and rapid learning throughout life

1.1. Fast recognition learning without catastrophic forgetting

Adaptive Resonance Theory, or ART, is a cognitive and neural theory of how the brain autonomously learns to categorize, recognize, and predict objects and events in a changing world. The problem of learning makes the unity of conscious experience hard to understand, if only because humans are able to rapidly learn enormous amounts of new information, on their own, throughout life. How do humans integrate all this information into unified experiences that cohere into a sense of self? One has only to see an exciting movie just once to marvel at this capacity, since we can then tell our friends many details about it later on, even though the individual scenes flashed by very quickly. More generally, we can quickly learn about new environments, even if no one tells us how the rules of each environment differ. To a remarkable degree, humans can rapidly learn new facts without being forced to just as rapidly forget what they already know. As a result, we can confidently go out into the world without fearing that, in learning to recognize a new friend's face, we will suddenly forget the faces of our family and friends. This is sometimes called the problem of *catastrophic forgetting*.

1.2. Some models that experience catastrophic forgetting

Many contemporary learning algorithms do experience catastrophic forgetting, particularly when they try to learn quickly in response to a changing world. These include the competitive learning, self-organizing map, back propagation, simulated annealing, neocognitron, support vector machine, regularization, and Bayesian models. The brain solves a challenging problem that many current biological and technological learning models have not yet solved: It is a self-organizing system that is capable of rapid, yet stable, autonomous learning in real time of huge amounts of data from a changing environment that can be filled with unexpected events. Discovering the brain's solution to this key problem is as important for understanding ourselves as it is for developing new pattern recognition and prediction applications in technology.

Grossberg (1980) has called the problem whereby the brain learns quickly and stably without catastrophically forgetting its past knowledge the *stability–plasticity dilemma*. The stability–plasticity dilemma must be solved by every brain system that needs to rapidly and adaptively respond to the flood of signals that subserves even the most ordinary experiences. If the brain's design is parsimonious, then similar design principles should operate in all brain systems that can stably learn an accumulating knowledge base in response to changing conditions throughout life. The discovery of such principles should clarify how the brain unifies diverse sources of information into coherent moments of conscious experience. ART has attempted to articulate these principles and the neural mechanisms that realize them. The next sections summarize aspects of how this is proposed to occur.

1.3. Linking consciousness, learning, expectation, attention, resonance, and synchrony

ART clarifies key brain processes from which conscious experiences emerge. It predicts a functional link between processes of Consciousness, Learning, Expectation, Attention, Resonance, and Synchrony (the CLEARs processes). ART predicted that all brain representations that solve the stability–plasticity dilemma use variations of CLEARs mechanisms (Grossberg, 1978a, 1980, 2007a). Synchronous resonances are, in particular, expected to occur between multiple cortical and subcortical areas. Various data support

this prediction; e.g., see Buschman and Miller (2007), Engel, Fries, and Singer (2001), Grossberg (2009b), and Pollen (1999).

Through these CLEARs connections, ART clarifies why many animals are intentional beings who pay attention to salient objects, why “all conscious states are resonant states”, and how brains can learn both *many-to-one maps* (representations whereby many object views, positions, and sizes all activate the same invariant object category) and *one-to-many maps* (representations that enable us to expertly know many things about individual objects and events).

ART accomplishes these properties by proposing how top-down expectations focus attention on salient combinations of cues, and characterizes how attention may operate via a form of self-normalizing “biased competition” (Desimone, 1998). ART explains how such top-down attentive matching may help to solve the stability–plasticity dilemma. In particular, when a good enough match occurs, a synchronous resonant state emerges that embodies an attentional focus and is capable of driving fast learning of bottom-up recognition categories and top-down expectations; hence the name *adaptive resonance*.

All of the main ART predictions have received increasing support from psychological and neurobiological data since ART was introduced in Grossberg (1976a, 1976b). ART has undergone continual development to explain and predict increasingly large behavioral and neurobiological data bases, ranging from normal and abnormal aspects of human and animal perception and cognition, to the spiking and oscillatory dynamics of hierarchically-organized laminar thalamocortical networks in multiple modalities. Indeed, some ART models explain and predict behavioral, anatomical, neurophysiological, biophysical, and even biochemical data. In this sense, they provide a growing set of examples capable of partially solving the classical mind/body problem. All the author's major articles, including those that develop ART, may be downloaded from <http://cns.bu.edu/~steve>.

1.4. Equations for short-term memory, medium-term memory, and long-term memory

How does ART sit within the corpus of all neural models? In particular, is the brain just a bag of tricks, as some authors have proposed (e.g., Ramachandran (1990))? This article illustrates a contrary view based on the author's view after developing models of mind and brain for 55 years (Grossberg, 1988, <http://cns.bu.edu/Profiles/Grossberg/GrossbergNEditorial2010.pdf>). During this period, I led the discovery and development of a small number of equations (e.g., equations for short-term memory, or STM; medium-term memory, or MTM; and long-term memory, or LTM) and a somewhat larger number of modules or microcircuits (e.g., shunting on-center off-surround networks, gated dipole opponent processing networks, associative learning networks, reinforcement learning networks, spectral timing networks, and the like), which have been specialized and assembled into modal architectures, where the term “modal” stands for modality (e.g., architectures for vision, audition, cognition, cognitive–emotional interactions, sensory–motor control, and the like). Modal architectures are less general than a Turing or von Neumann architecture for general computing, but far more general than a traditional AI algorithm. They are designed to be capable of general-purpose self-organizing processing of a particular modality of biological intelligence and their particular specializations of the basic equations and modules have been selected over the millennia by evolutionary pressures.

ART networks form part of such modal architectures. Modal architectures, in turn, embody new paradigms for brain computing that I have called Complementary Computing (Grossberg, 2000b)

and Laminar Computing (Grossberg, 1999). Complementary Computing describes how the global brain is organized into complementary parallel processing streams whose interactions generate biologically intelligent behaviors. Laminar Computing describes how the cerebral cortex is organized into layered circuits whose specializations can support all forms of higher-order biological intelligence. ART networks exhibit many complementary properties, and are embodied in laminar cortical circuitry for vision and visual object recognition; audition, speech, and language; and cognitive information processing.

Grossberg (1968a, 1968b, 1968c, 1969a, 1969b, 1969c, 1972b) introduced the laws of STM, MTM, and LTM that are used in many contemporary neural models, including ART architectures, and mathematically proved various of their basic properties. These laws are specialized to cope with the evolutionary pressures that are embodied in different modal architectures. One variant of these equations is:

STM: fast activation dynamics

$$\frac{dx_i}{dt} = -Ax_i + (B - Cx_i) \left[I_i + \sum_{k=1}^n f_k(x_k) y_k D_{ki} Z_{ki} \right] - (E + Fx_i) \left[J_i + \sum_{k=1}^n g_k(x_k) Y_k G_{ki} Z_{ki} \right]. \quad (1)$$

Eq. (1) describes the activity, x_i , of the i th cell (population) in a network of n interacting neurons. It includes both the Additive and Shunting models. In the shunting model, the parameters $C \neq 0$ and $F \neq 0$, thereby enabling the automatic gain control and normalization properties that shunting networks exhibit. The parameter $E = 0$ when there is “silent” shunting inhibition, whereas $E \neq 0$ describes the case of hyperpolarizing shunting inhibition. In the Additive model, parameters $C = F = 0$. The excitatory interaction term $[I_i + \sum_{k=1}^n f_k(x_k) y_k D_{ki} Z_{ki}]$ describes an external input I_i plus the total excitatory feedback signal $[\sum_{k=1}^n f_k(x_k) y_k D_{ki} Z_{ki}]$ that is a sum of signals from other populations via their output signals $f_k(x_k)$. The term D_{ki} is a constant connection strength between cell populations k and i , whereas terms y_k and Z_{ki} describe MTM and LTM variables, respectively. The inhibitory interaction term $[J_i + \sum_{k=1}^n g_k(x_k) Y_k G_{ki} Z_{ki}]$ has a similar interpretation. Eq. (1) assumes “fast inhibition”; that is, inhibitory interneurons respond instantaneously to their inputs. Slower finite-rate inhibition with activities X_i uses an equation like (1) to describe the temporal evolution of the inhibitory activities. The output signals from these inhibitory interneurons provide the feedback signals to the excitatory activities. With slow inhibition, the inhibitory feedback signals would be $g_k(X_k)$ instead of $g_k(x_k)$.

MTM: habituated transmitter gates and depressing synapses

$$\frac{dy_i}{dt} = H(K - y_i) - Lf_k(x_k) y_k. \quad (2)$$

Eq. (2) describes how the strength y_i of the habituated transmitter gate, or depressing synapse, in the excitatory feedback term of (2) accumulates at a fixed rate H to its maximum value K via term $H(K - y_i)$ and is inactivated, habituated, or depressed via a mass action interaction between the feedback signal $f_k(x_k)$ and the gate concentration y_k . The mass action term may be more complex than this in some situations; e.g., Gaudiano and Grossberg (1991, 1992). The habituated transmitter gate Y_k in the inhibitory feedback term of (1) obeys a similar equation. It should be noted that not all signals are habituated.

LTM: gated steepest descent learning

$$\frac{dz_{ij}}{dt} = Mf_j(x_j) [h_i(x_i) - z_{ij}] \quad (3)$$

and

$$\frac{dz_{ij}}{dt} = Mf_j(x_j) [h_i(x_i) - z_{ij}]. \quad (4)$$

Eq. (3) describes the *outstar* learning equation, by which the i th source cell can sample and learn a distributed spatial pattern of activation across a network of sampled cells. When the gating signal $f_i(x_i)$ is positive, the adaptive weights z_{ij} can sample the activity-dependent signals $h_j(x_j)$ across the sampled network of cells. Eq. (4) describes the *instar* learning equation, by which the j th target cell can sample and learn the distributed pattern of signals that activated it. This is the learning equation that was used in the competitive learning and self-organizing map models in Grossberg (1976a, 1978a), and later applied by Kohonen (1984). The learning instabilities of competitive learning and self-organizing maps led to the introduction of ART.

Instars and outstars were both used in Grossberg (1976a) to show how to learn arbitrary maps from m -dimensional to n -dimensional vectors, with instars first compressing the m -dimensional input vectors at category cells, and the category cells then sampling and learning the n -dimensional output vectors. By showing how to dynamically self-stabilize instar learning, ART also showed how to learn self-stabilizing maps. This type of instar-outstar map learning scheme was called counterpropagation by Hecht-Nielsen (1987). There are many variations of these gated steepest descent equations (e.g., doubly-gated learning, spike-timing dependent learning, etc.). It should also be noted that not all connections are adaptive.

1.5. Applications of ART to engineering and technology

Although ART models are realized mathematically by nonlinear neural networks, this review focuses on heuristically summarizing ART as a cognitive and neural theory. One part of the mathematical development of ART has been to develop algorithms that computationally embody specific combinations of useful ART design principles. An algorithmic form is used to enable the model to run faster on the computer. These algorithms typically embody a singular approximation to the full dynamics of an ART system. For example, fast dynamics, such as STM activations (see Eq. (1)), are often solved at equilibrium in these algorithms. By simplifying computations, these algorithms have contributed to the mathematical development of the cognitive and neural theory, and are widely used in large-scale engineering and technological applications, such as medical data base prediction, remote sensing, airplane design, and the control of autonomous adaptive robots.

A standard ART algorithm for applications is called Default ARTMAP (Amis & Carpenter, 2007; Carpenter, 2003). Early important ART algorithms for applications include ART 1, ART 2, ARTMAP, fuzzy ART, and fuzzy ARTMAP (Carpenter & Grossberg, 1987; Carpenter, Grossberg, Markuzon, Reynolds, & Rosen, 1992; Carpenter, Grossberg, & Reynolds, 1991; Carpenter, Grossberg, & Rosen, 1991). More recent algorithms from Gail Carpenter and her students include distributed ARTMAP, which combines distributed coding with fast, stable, incremental learning (Carpenter, 1997; Carpenter, Milenova, & Noeske, 1998); ARTMAP Information Fusion, which can incrementally learn a cognitive hierarchy of rules in response to probabilistic, incomplete, and even contradictory data that are collected by multiple observers (Carpenter, Martens, & Ogas, 2005; Carpenter & Ravindran, 2008; Parsons & Carpenter, 2003); Self-supervised ART, which shows how some supervised learning “in school” can lead to effective knowledge acquisition later on by unsupervised learning “in the real world” (Amis & Carpenter, 2009); and Biased ART, which shows how attention can be selectively diverted from features that

cause predictive errors (Carpenter & Gaddam, 2010). Computer code for running various ART algorithms and related neural models that were discovered and developed at Boston University can be found at <http://techlab.bu.edu/resources/software/C51>.

Many variants of ART have been developed and applied to large-scale engineering and technological applications by authors around the world (e.g., Akhbardeh, Junnila, Koivistoinen, and Varri (2007), Anagnostopoulos and Georgiopoulos (2000), Anton-Rodriguez et al. (2009), Brannon, Seiffert, Draelos, and Wunsch (2009), Cai, Wang, Tang, and Yang (2011), Caudell (1992), Caudell, Smith, Johnson, Wunsch, and Escobedo (1991), Cano-Izquierdo, Almonacid, Pinzolas, and Ibarrola (2009), Chao, Hsiao, Su, Hsu, and Wu (2011), Cherng, Fang, Chen, and Chen (2009), Demetgul, Tansel, and Taskin (2009), Dunbar (2012), He, Caudell, Menicucci, and Mammoli (2012); He, Tan, and Tan (2000), Healy, Caudell, and Smith (1993), Ho, Liou, Georgiopoulos, Heileman, and Christodoulou (1994), Hsieh (2008), Hsieh and Yang (2008), Hsu and Chien (2007), Kawamura, Takahashi, and Honda (2008), Kaylani, Georgiopoulos, Mollaghasemi, and Anagnostopoulos (2009), Keskin and Ozkan (2009), Liu, Huang, Lai, and Ma (2009); Liu, Pang, and Lloyd (2008), Lopes, Minussi, and Lotufo (2005), Marchiori, da Silveira, Lotufo, Minussi, and Lopes (2011), Martin-Guerrero, Lisboa, Soria-Olivas, Palomares, and Balaguer (2007), Massey (2009), Mulder and Wunsch (2003), Owega, Khan, Evans, Jervis, and Fila (2006), Prasad and Gupta (2008), Shieh, Yan, and Chen (2008), Sudhakara Pandian and Mahapatra (2009), Takahashi, Murase, Kobayashi, and Honda (2007), Tan (1997), Tan and Teo (1998), Tan, Quek, Ng, and Razvi (2008), Wienke and Buydens (1995), Wunsch, Caudell, Capps, Marks, and Falk (1993), Xu et al. (2009) and Zhang and Kezunovic (2007)). A repository of some applications is found at <http://techlab.bu.edu/resources/articles/C5>.

2. The varieties of learning algorithms

Many current learning algorithms do not emulate the way in which humans and other animals learn. The power of human and animal learning provides high motivation to discover computational principles whereby machines can learn with similar capabilities. Humans and animals experience the world on the fly, and carry out incremental learning of sequences of episodes in real time. Often such learning is unsupervised, with the world itself as the teacher. Learning can also proceed with an unpredictable mixture of unsupervised and supervised learning trials. Such learning goes on successfully in a world that is non-stationary; that is, whose rules can change unpredictably through time. Moreover, humans and animals can learn quickly, and their memories can persist for a long time thereafter, even while new learning takes place. ART's solution of this *stability–plasticity dilemma* (Grossberg, 1980) predicts how brains may learn quickly without forcing catastrophic forgetting of already learned, and still successful, memories.

Thus, ART autonomously carries out fast, yet stable, incremental learning under both unsupervised and supervised learning conditions in response to a complex non-stationary world. In contrast, many current learning algorithms use batch learning in which all the information about the world to be learned is available at a single time, such as support vector machines. Other algorithms are not defined unless all learning trials are supervised, such as back propagation. Most learning algorithms become unstable in a non-stationary world, or cannot learn about important rare cases, or become unstable if learning is fast; that is, if an event can be fully learned on a single learning trial. ART overcomes these problems.

Some machine learning algorithms are feedforward clustering algorithms that undergo catastrophic forgetting in a non-stationary world. These include competitive learning, self-organizing maps, neocognitrons, and back propagation. The ART solution of the stability–plasticity dilemma depends upon feedback, or top-down, learned expectations that are matched against bottom-up data patterns and thereby focus attention upon those combinations of features that are predictive in that context. A good enough match with a critical feature pattern leads to a synchronous resonance and fast learning. A big enough mismatch leads to hypothesis testing, or memory search, that discovers, chooses, and learns a more predictive category. Thus, ART is a self-organizing expert system that avoids the brittleness of traditional expert systems.

The world is filled with uncertainty, so probability concepts seem relevant to understanding how brains learn about uncertain data. This fact has led some machine learning practitioners to assume that brains obey Bayesian laws (e.g., Knill and Pouget (2004) and Doya et al. (2007)). However, the Bayes rule is so general that it can accommodate any system in Nature. This generality makes Bayes a very useful statistical method. However, in order for Bayes concepts to be part of a physical theory, additional computational principles and mechanisms are needed to augment the Bayes rule to distinguish a brain from, say, a hydrogen atom or a hurricane. Because of the generality of the Bayes rule, it does not, in itself, provide heuristics for discovering what these distinguishing physical principles might be.

It is an empirical question whether a Bayesian formalism will be part of the most successful physical theories. This is not true in the physical sciences. It is also not true in the Adaptive Resonance Theory models that the current article reviews. Although these models may exhibit Bayesian *properties*, these properties emerge from organization principles and mechanisms that do not invoke Bayesian ideas.

Probabilistic models such as Bayesian statistics often use non-local computations, that are needed to compute an entire probability distribution, and work best in dealing with stationary data. ART uses only local computations to model how the brain may embody a novel kind of real-time probability theory, hypothesis testing, prediction, and decision-making whose local computations enable fast autonomous adaptation to a non-stationary world whose rules can change through time, and in unexpected ways. These ART principles and mechanisms go beyond Bayesian analysis, and are embodied parsimoniously in the laminar circuits of cerebral cortex. Indeed, cerebral cortex embodies a kind of Laminar Computing that embodies a revolutionary computational synthesis of the best properties of feedforward and feedback processing, digital and analog processing, and data-driven bottom-up processing and hypothesis-driven top-down processing (Grossberg, 2003a, 2007a, see Section 30).

3. The predictive brain: learning and prediction by complementary cortical streams for recognition and action

Biological learning includes both perceptual/cognitive and spatial/motor processes. Accumulating experimental and theoretical evidence show that perceptual/cognitive and spatial/motor processes both need predictive mechanisms to control learning. Thus there is an intimate connection between learning and predictive dynamics in the brain. However, neural models of these processes have proposed, and many experiments have supported, the hypothesis that perceptual/cognitive and spatial/motor processes use different types of predictive mechanisms to regulate the learning that they carry out.

WHAT	WHERE	
Spatially-invariant object learning and recognition	Spatially-variant reaching and movement	
Fast learning without catastrophic forgetting	Continually update sensory-motor maps and gains	
IT	PPC	
	WHAT	WHERE
MATCHING	EXCITATORY	INHIBITORY
LEARNING	MATCH	MISMATCH

Fig. 1. Complementary What and Where cortical processing streams for spatially-invariant object recognition and spatially-variant spatial representation and action, respectively. Perceptual and recognition learning use top-down excitatory matching and match-based learning that achieve fast learning without catastrophic forgetting. Spatial and motor learning use inhibitory matching and mismatch-based learning that enable rapid adaptation to changing bodily parameters. IT = inferotemporal cortex, PPC = posterior parietal cortex. See text for details. Source: Reprinted with permission from Grossberg (2009b).

3.1. Excitatory matching and match learning vs. inhibitory matching and mismatch learning

The need for different predictive mechanisms is clarified by accumulating theoretical and empirical evidence that brain specialization is governed by computationally complementary cortical processing streams that embody different predictive and learning mechanisms (Grossberg, 2000b). As summarized in Fig. 1, perceptual/cognitive processes in the What ventral cortical processing stream often use *excitatory matching* and *match-based learning* to create predictive representations of objects and events in the world. Match-based learning solves the stability–plasticity dilemma and is the kind of learning used in ART. This sort of learning can occur quickly without causing catastrophic forgetting, much as we quickly learn new faces without forcing rapid and unselective forgetting of familiar faces. However, match learning, and by extension ART, does not describe the only kind of learning that the brain needs to accomplish autonomous adaptation to a changing world. If only for this reason, ART is not a theory about “everything”.

3.2. Learning to be an expert in a changing body

There are just as essential, but complementary, spatial/motor processes in the Where dorsal cortical processing stream that often use *inhibitory matching* and *mismatch-based learning* to continually update spatial maps and sensory–motor gains as our bodily parameters change through time (Bullock, Cisek, & Grossberg, 1998; Bullock & Grossberg, 1988; Gaudiano & Grossberg, 1991; Georgopoulos, Kalaska, Caminiti, & Massey, 1982; Georgopoulos, Schwartz, & Kettner, 1986). Indeed, we would be ill-served by spatial and motor learning processes that solve the stability–plasticity dilemma, since we do not want the spatial representations and motor gains that were suitable for controlling our infant bodies to be remembered as we grow up and used to control our adult bodies. In this sense, catastrophic forgetting is a good property during spatial and motor learning.

As an example of inhibitory spatial matching, consider how we make an arm movement. To make such a movement, a representation of where the arm is now (its present position vector) is subtracted from a representation of where we want the arm to move (its target position vector), thereby computing a difference vector that represents the direction and distance of movement needed to attain the target. After moving to the target,

the target and present positions agree, so the difference vector is zero. In other words, this sort of matching is inhibitory (Bullock & Grossberg, 1988).

Neither type of matching and learning is sufficient to design an adaptive autonomous agent, but each is necessary. By combining these two types of processes together, our brains can incrementally learn and stably remember perceptual and cognitive representations of a changing world, leading to a self-stabilizing front end that solves the stability–plasticity dilemma and enables us to become increasingly expert in understanding the world and predicting outcomes in the world. At the same time, our brains can adaptively update their representations of where objects are and how to act upon them using bodies whose parameters change continuously through time due to development, exercise, illness, and aging.

3.3. Why procedural memories are not conscious

Brain systems that use inhibitory matching and mismatch learning cannot generate excitatory resonances. Hence, if “all conscious states are resonant states”, then spatial and motor representations are not conscious. This way of thinking provides a mechanistic reason why declarative memories (or “learning that”), which are the sort of memories learned by ART, may be conscious, whereas procedural memories (or “learning how”), which are the sort of memories that control spatial orienting and action, are not conscious (Cohen & Squire, 1980).

3.4. Spatially-invariant recognition vs. spatially localized action

There is another basic reason why these complementary What and Where processes need to work together. The What stream attempts to learn spatially-invariant object categories, so that a combinatorial explosion does not occur wherein every view of every object at every position and distance needs to be represented by a different category.

Indeed, learning in the What cortical stream leads to recognition categories that tend to be increasingly independent of object size and position at higher cortical levels. The anterior inferotemporal cortex exhibits such invariance (Bar et al., 2001; Sigala & Logothetis, 2002; Tanaka, Saito, Fukada, & Moriya, 1991). Although how this occurs needs careful discussion (e.g., Zoccolan, Kouh, Poggio, and DiCarlo (2007)), such object invariance prevents a combinatorial explosion in memory of object representations that could otherwise occur at every perceived size and position. Cao, Grossberg, and Markowitz (2011) and Grossberg, Markowitz, and Cao (2011) have used ART to simulate recent neurophysiological data about neurophysiologically observed properties of invariant category learning and recognition in inferotemporal cortex.

In becoming spatially invariant, recognition categories lose information about how to direct action towards the locations in space where desired objects may be found. In contrast, the Where stream learns spatial maps that do enable us to locate such desired objects, as well as the movement gains that enable us to accurately act with respect to them. On the other hand, Where stream spatial processing gives up information about which objects are in those spatial locations. Interactions between the What and Where stream (“What–Where fusion”) overcome these complementary deficiencies to enable spatially-invariant object representations to control actions towards desired goals in space (e.g., Brown, Bullock, and Grossberg (2004), Fazl, Grossberg, and Mingolla (2009), Grossberg (2009b) and Grossberg and Vladusich (2011)).

In summary, because of their different types of matching and learning, perceptual and cognitive learning provide a self-stabilizing front end to control the more labile spatial and motor learning that enables changing bodies to effectively act upon recognized objects in the world. The present article summarizes how this may happen.

4. Learning, expectation, attention, and intention

Humans are *intentional* beings who learn expectations about the world and make predictions about what is about to happen. Humans are also *attentional* beings who focus processing resources upon a restricted amount of incoming information at any time. Why are we both intentional and attentional beings, and are these two types of processes related? The stability–plasticity dilemma and its solution using resonant states provides a unifying framework for understanding these issues.

4.1. Top-down attentional priming

To clarify the role of sensory or cognitive expectations, and of how a resonant state is activated, suppose you were asked to “find the yellow ball as quickly as possible, and you will win a \$100,000 prize”. Activating an expectation of a “yellow ball” enables its more rapid detection, and with a more energetic neural response. Sensory and cognitive top-down expectations hereby lead to *excitatory matching* with consistent bottom-up data. A mismatch between top-down expectations and bottom-up data can suppress the mismatched part of the bottom-up data, while attention is focused upon the matched, or expected, part of the bottom-up data.

4.2. Learning of attended critical feature patterns

Excitatory matching and attentional focusing on bottom-up data using top-down expectations generates resonant brain states: When there is a good enough match between bottom-up and top-down signal patterns between two or more levels of processing, their positive feedback signals amplify, synchronize, and prolong their mutual activation, leading to a resonant state that focuses attention on a combination of features (the “critical feature pattern”) that are needed to correctly classify the input pattern at the next processing level and beyond. Amplification, synchronization, and prolongation of activity triggers learning in the more slowly varying adaptive weights that control the signal flow along pathways between the attended features and the recognition category with which they resonate. Resonance hereby provides a global context-sensitive indicator that the system is processing data worthy of learning, hence the name *Adaptive Resonance Theory*.

In summary, ART predicts a link between the mechanisms which enable us to learn quickly and stably about a changing world, and the mechanisms that enable us to learn expectations about such a world, test hypotheses about it, and focus attention upon information that may predict desired consequences. ART clarifies this link by asserting that, in order to solve the stability–plasticity dilemma, only resonant states can drive fast new learning.

5. Linking brain to behavior: all conscious states are resonant states

It is just a step from here to propose that those experiences which can attract our attention and guide our future lives after being learned are also among the ones that are conscious. Support for the predicted link between resonance and consciousness comes many modeling studies wherein the parametric properties of brain resonances map onto parametric properties of conscious behavioral experiences in the simulated experiments. Indeed, without such a linking hypothesis between brain mechanisms and behavioral functions, no theory of consciousness can be fully tested.

Although it is predicted that “all conscious states are resonant states”, it is not predicted that “all resonant states are conscious

states”. Indeed, some resonant states, such as the storage of a sequence of events in working memory before rehearsal occurs (see Sections 50–52), or the entorhinal–hippocampal resonances that may dynamically stabilize the learning of entorhinal grid cells and hippocampal place cells (see Section 40), are not accessible to consciousness.

6. ART Matching Rule and biased competition: modulatory on-center, off-surround network

6.1. Attention obeys the ART Matching Rule

How are What stream top-down expectations computed? How do they focus attention on expected combinations of features? Carpenter and Grossberg (1987) mathematically proved that the simplest attentional circuit that solves the stability–plasticity dilemma is a *top-down, modulatory on-center, off-surround network*, which provides excitatory priming of critical features in the on-center, and driving inhibition of irrelevant features in the off-surround. Eq. (1) describes an on-center off-surround network when the excitatory connections are spatially more localized than the inhibitory connections. The modulatory on-center emerges from a balance between top-down excitation and inhibition. The neurons in the network obey the membrane equations of neurophysiology. The entire attentional circuit is said to satisfy the ART Matching Rule.

6.2. Solving the noise–saturation dilemma with shunting on-center off-surround networks

It was first proved by Grossberg (1973) that the shunting, or gain control, properties of membrane equation neurons in an on-center off-surround network (Eq. (1)) enable them to self-normalize their activities, and thereby solve a design problem that is just as basic as the stability–plasticity dilemma. This design problem is called the *noise–saturation dilemma*. Without suitable interactions between neurons, their inputs can be lost in cellular noise if they are too small, or can saturate cell activities at their maximum values if they are too large. Moreover, input amplitudes can vary greatly through time. What sort of network interactions enable neurons to retain their sensitivities to the *relative sizes* of their inputs across the network, even while these inputs may vary in size through time over several orders of magnitude? The answer is: an on-center off-surround network whose cells exhibit shunting properties.

Modeling studies clarify how a top-down, modulatory on-center, off-surround network can participate in many different types of brain processes (e.g., Dranias, Grossberg, and Bullock (2008), Gove, Grossberg, and Mingolla (1995), Grossberg, Govindarajan, Wyse, and Cohen (2004), Grossberg and Kazerounian (2011)). Models of how cerebral cortex embodies attention within its layered circuits have discovered that identified cell types and connections occur with the necessary properties to realize the ART Matching Rule (Grossberg, 1999; Grossberg & Pearson, 2008; Grossberg & Versace, 2008; Raizada & Grossberg, 2003).

6.3. Data support for the ART Matching Rule

Many anatomical and neurophysiological experiments have provided support for the ART prediction of how attention works, including data about modulatory on-center, off-surround interactions; excitatory priming of features in the on-center; suppression of features in the off-surround; and gain amplification of matched data (e.g., Bullier, Hupé, James, and Girard (1996), Caputo and Guerra (1998), Downing (1988), Hupé, James, Girard, and Bullier (1997), Mounts (2000), Reynolds, Chelazzi, and

Desimone (1999), Sillito, Jones, Gerstein, and West (1994), Somers, Dale, Seiffert, and Tootell (1999), Steinman, Steinman, and Lehmkuhle (1995) and Vanduffel, Tootell, and Orban (2000)). The ART Matching Rule is often called the “biased competition” model of attention by experimental neurophysiologists (Desimone, 1998; Kastner & Ungerleider, 2001). The property of the ART Matching Rule that bottom-up sensory activity may be enhanced when matched by top-down signals is in accord with an extensive neurophysiological literature showing the facilitatory effect of attentional feedback (Luck, Chelazzi, Hillyard, & Desimone, 1997; Roelfsema, Lamme, & Spekreijse, 1998; Sillito et al., 1994), but not with models, such as Bayesian “explaining away” models, in which matches with top-down feedback cause only suppression (Mumford, 1992; Rao & Ballard, 1999).

The ART Matching Rule helps to explain the existence of top-down modulatory connections at multiple stages of cortical processing. For example, Zeki and Shipp (1988, p. 316) wrote that “backward connections seem not to excite cells in lower areas, but instead influence the way they respond to stimuli”; that is, they are modulatory. Likewise, the data of Sillito et al. (1994, pp. 479–482) on attentional feedback from cortical area V1 to the Lateral Geniculate Nucleus (LGN) support an early prediction that the ART Matching Rule should exist in this pathway as well (Grossberg, 1976b). In this regard, Sillito et al. (1994) concluded that “the cortico-thalamic input is only strong enough to exert an effect on those dLGN cells that are additionally polarized by their retinal input. . . the feedback circuit searches for correlations that support the ‘hypothesis’ represented by a particular pattern of cortical activity”. Their experiments demonstrated all of the properties of the ART Matching Rule, since they also found that “cortically induced correlation of relay cell activity produces coherent firing in those groups of relay cells with receptive-field alignments appropriate to signal the particular orientation of the moving contour to the cortex. . . this increases the gain of the input for feature-linked events detected by the cortex”. In other words, top-down priming, by itself, cannot fully activate LGN cells; it needs matched bottom-up retinal inputs to do so; and those LGN cells whose bottom-up signals support cortical activity get synchronized and amplified by this feedback. In addition, anatomical studies have shown that the V1-to-LGN pathway realizes a top-down on-center off-surround network (Dubin & Cleland, 1977; Sillito et al., 1994; Weber, Kalil, & Behan, 1989); see Fig. 5(d) below. Nobuo Suga and colleagues have shown that feedback from auditory cortex to the medial geniculate nucleus (MGN) and the inferior colliculus (IC) also has an on-center off-surround form (Zhang, Suga, and Yan (1997), and Temereanca and Simons (2001) have produced evidence for a similar feedback architecture in the rodent barrel system.

6.4. Mathematical form of the ART Matching Rule

There is also a convergence across models of how to mathematically instantiate the ART Matching Rule attentional circuit. For example, the “normalization model of attention” (Reynolds & Heeger, 2009) simulates several types of experiments on attention using the same equation for self-normalizing attention as the distributed ARTEXture (dARTEX) model (Bhatt, Carpenter, & Grossberg, 2007, Eq. (A5)) used to simulate human psychophysical data about Orientation-Based Texture Segmentation (OBTS, (Ben-Shahar & Zucker, 2004)).

7. Imagining, planning, and hallucinations: prediction without action

A top-down expectation is not always modulatory. The excitatory/inhibitory balance in the modulatory on-center of a top-down expectation can be modified by volitional control from the

basal ganglia. If, for example, volitional signals inhibit inhibitory interneurons in the on-center, then read-out of a top-down expectation from a recognition category can fire cells in the on-center prototype, not merely modulate them. Such volitional control has been predicted to control mental imagery and the ability to think and plan ahead without external action, a crucial type of predictive competence in humans and other mammals. If these volitional signals become tonically hyperactive, then top-down expectations can fire without overt intention, leading to properties like schizophrenic hallucinations (Grossberg, 2000a). In summary, our ability to learn quickly without catastrophic forgetting led to circuits that can be volitionally modulated to enable imagination, internal thought, and planning. This modulation, which brings a huge evolutionary advantage to those who have it, also carries with it the risk of causing hallucinations.

A similar modulatory circuit, again modulated by the basal ganglia, is predicted to control when sequences of events are stored in short-term working memory in the prefrontal cortex (Grossberg & Pearson, 2008, see Sections 50–52) and the span of spatial attention (“useful-field-of-view”) in the parietal and prefrontal cortex (Foley, Grossberg, & Mingolla, 2012, see Section 17). All of these properties build upon the fundamental ability to learn quickly throughout life without catastrophic forgetting by using top-down expectations to stabilize learned memories.

8. Complementary attentional and orienting systems: expected/unexpected, resonance/reset

8.1. The cycle of resonance and reset

As noted above, the type of learning within the sensory and cognitive domain that ART mechanizes is *match learning*: Match learning occurs only if a good enough match occurs between bottom-up information and a learned top-down expectation that is read out by an active recognition category, or code. When such an approximate match occurs, a resonance can be triggered, whereupon previous knowledge can be refined through learning. It has been mathematically proved that match learning within an ART model leads to stable memories of arbitrary events presented in any order (e.g., Carpenter and Grossberg (1987, 1991)).

However, match learning also has a serious potential weakness: If you can only learn when there is a good enough match between bottom-up data and learned top-down expectations, then how do you ever learn anything that is really novel? ART proposes that this problem is solved by the brain by using an interaction between complementary processes of *resonance* and *reset* that are predicted to control properties of attention and memory search, respectively. These complementary processes help our brains to balance between the complementary demands of processing the familiar and the unfamiliar, the expected and the unexpected.

This problem raises the following basic questions: How does a brain learn to balance between expected and unexpected events? How does a brain learn to incorporate unexpected and unfamiliar events within the corpus of previously learned events, and do so without causing catastrophic forgetting? ART proposes that, when novel inputs cannot match a known recognition category, a memory search, or hypothesis testing, process is activated that enables our brains to discover and learn new recognition categories that best match novel objects or events.

Organization of the brain into complementary processes is predicted to be a general principle of brain design that is not just found in ART (Grossberg, 2000b). A complementary process can individually compute some properties well, but cannot, by itself, process other complementary properties. In thinking intuitively about complementary properties, one can imagine puzzle pieces fitting together. Both pieces are needed to finish the puzzle.

Complementary brain processes are, however, more dynamic than any such analogy: Pairs of complementary processes interact to form emergent properties which overcome their complementary deficiencies to compute complete information with which to represent or control some aspect of intelligent behavior.

The resonance process in the complementary pair of resonance and reset is predicted to take place in the What cortical stream, notably in the sensory, temporal, and prefrontal cortices. Here top-down expectations are matched against bottom-up inputs. When a top-down expectation achieves a good enough match with bottom-up data, this match process focuses attention upon those feature clusters in the bottom-up input that are expected. If the expectation is close enough to the input pattern, then a state of resonance develops as the attentional focus takes hold, which is often realized by oscillatory dynamics that synchronize the firing properties of the resonant neurons. Such a resonance opens the learning gates in the gated steepest descent learning laws (Eqs. (3) and (4)).

However, as noted above, a sufficiently bad mismatch between an active top-down expectation and a bottom-up input, say because the input represents an unfamiliar type of experience, can drive a memory search. Such a mismatch within the attentional system is proposed to activate a complementary *orienting system*, which is sensitive to unexpected and unfamiliar events. ART suggests that this orienting system includes the nonspecific thalamus and the hippocampal system. See Carpenter and Grossberg (1993) and Grossberg and Versace (2008) for a summary of data supporting this prediction. Output signals from the orienting system rapidly reset the recognition category that has been reading out the poorly matching top-down expectation. The cause of the mismatch is hereby removed, thereby freeing the system to activate a different recognition category. In this way, a reset event triggers memory search, or hypothesis testing, which automatically leads to the selection of a recognition category that can better match the input.

If no such recognition category exists, say because the bottom-up input represents a truly novel experience, then the search process automatically activates an as yet uncommitted population of cells, with which to learn about the novel information. In order for a top-down expectation to match the features that activated a new recognition category, its top-down adaptive weights initially have large values, which are pruned by the learning of a particular expectation; see Section 39.

8.2. ART search cycle

Fig. 2 illustrates these ART ideas in a two-level network. Here, a bottom-up input pattern, or vector, I activates a pattern X of activity across the feature detectors of the first level F_1 . For example, a visual scene may be represented by the features comprising its boundary and surface representations (see Section 20). This feature pattern represents the relative importance of different features in the inputs pattern I . In Fig. 2(a), the pattern peaks represent more activated feature detector cells, the troughs less activated feature detectors. This feature pattern sends signals S through an adaptive filter to the second level F_2 at which a compressed representation Y (also called a recognition category, or a symbol) is activated in response to the distributed input T . Input T is computed by multiplying the signal vector S by a matrix of adaptive weights, or long-term memory traces, that can be altered through learning. The representation Y is compressed by competitive interactions – in particular, shunting recurrent lateral inhibition – across F_2 that allow only a small subset of its most strongly activated cells to remain active in response to T . These active cells are the recognition category that represents the pattern of distributed features across level F_1 . The pattern Y in the figure




Fig. 2. Search for a recognition code within an ART learning circuit: (a) Input pattern I is instated across feature detectors at level F_1 as an activity pattern X , while it nonspecifically activates the orienting system A with gain ρ , which is called the *vigilance* parameter. Output signals from activity pattern X inhibits A and generates output pattern S . S is multiplied by learned adaptive weights to form the input pattern T . T activates category cells Y at level F_2 . (b) Y generates the top-down signals U which are multiplied by adaptive weights and added at F_1 cells to form a *prototype* V that encodes the learned expectation of active F_2 categories. If V mismatches I at F_1 , then a new STM activity pattern X^* (the hatched pattern) is selected at F_1 . X^* is active at I features that are confirmed by V . Mismatched features (white area) are inhibited. When X changes to X^* , total inhibition decreases from F_1 to A . (c) If inhibition decreases sufficiently so that the total inhibition due to X^* is less than the total excitation due to I multiplied by the vigilance parameter ρ , then A is activated and releases a nonspecific arousal burst to F_2 ; that is, “novel events are arousing”. Arousal resets F_2 by inhibiting Y . (d) After Y is inhibited, X is reinstated and Y stays inhibited as X activates a different activity pattern Y^* . Search for better F_2 category continues until a better matching or novel category is selected. When search ends, an attentive resonance triggers learning of the attended data. Source: Adapted with permission from Carpenter and Grossberg (1993).

indicates that a small number of category cells may be activated to different degrees.

These category cells, in turn, send top-down signals U to F_1 . The vector U is converted into the top-down expectation V by being multiplied by another matrix of adaptive weights. When V is received by F_1 , a matching process takes place between the input vector I and V which selects that subset X^* of F_1 features that were “expected” by the active F_2 category Y . The set of these selected features is the emerging “attentional focus” that is gain amplified by the top-down match.

9. Synchronous binding of feature patterns and categories during conscious resonances

If the top-down expectation is close enough to the bottom-up input pattern, then the pattern X^* of attended features reactivates the category Y which, in turn, reactivates X^* . The network hereby locks into a resonant state through a positive feedback loop that dynamically links, or binds, the attended features across X^* with their category, or symbol, Y .

9.1. Resonant synthesis of complementary categories and distributed feature patterns

The resonance process itself embodies another type of complementary processing. Indeed, there seem to be complementary processes both within and between cortical processing streams (Grossberg, 2000b). This particular complementary relation occurs between distributed feature patterns and the compressed categories, or symbols, that selectively code them:

Individual features at F_1 have no meaning on their own, just like the pixels in a picture are meaningless one-by-one. The category, or symbol, in F_2 is sensitive to the global patterning of these features, and can selectively fire in response to this pattern. But it cannot represent the “contents” of the experience, including their conscious qualia, due to the very fact that a category is a compressed, or “symbolic” representation. Practitioners of Artificial Intelligence have claimed that neural models can process distributed features, but not symbolic representations. This is not, of course, true in the brain. Nor is it true in ART.

Resonance between these two types of information converts the *pattern* of attended features into a coherent context-sensitive state that is linked to its category through feedback. Coherent binding of the attended features to the category give them a meaning as a context-sensitive “event” rather than as just isolated pixels. Such coherent states between distributed features and symbolic categories are often expressed dynamically as synchronously oscillating activations across the bound cells, and can enter consciousness.

9.2. Order-preserving limit cycles and synchronous oscillations

The original ART article (Grossberg, 1976b) predicted the existence of such synchronous oscillations, which were there described in terms of their mathematical properties as “order-preserving limit cycles”. The property of “order-preservation” means that the relative sizes, and thus importance, of the feature activations should not reverse during the oscillation, which could occur, for example, during a traveling wave. Many neurophysiological experiments have been done confirming the existence of synchronous oscillations since the original confirmatory experimental reports of Eckhorn et al. (1988) and Gray and Singer (1989). See Raizada and Grossberg (2003) and Grossberg and Versace (2008) for reviews of confirmed ART predictions, including predictions about synchronous oscillations.

10. Resonance links intentional and attentional information processing to learning

In ART, the resonant state, rather than bottom-up activation alone, is predicted to drive fast learning. The synchronous resonant state persists long enough, and at a high enough activity level, to activate the slower learning processes in the adaptive weights that guide the flow of signals between bottom-up adaptive filter and top-down expectation pathways between levels F_1 and F_2 in Fig. 2. This viewpoint helps to explain how adaptive weights that were changed through previous learning can regulate the brain’s present information processing, without necessarily learning about the signals that they are currently processing unless the network as a whole can initiate a resonant state. Through resonance as a mediating event, one can understand from a deeper mechanistic view why humans are intentional beings who are continually predicting what may next occur, and why we tend to learn about the events to which we pay attention.

This *match-based learning* process stabilizes learned memories both in the bottom-up adaptive filters that activate recognition categories and in the top-down expectations that are matched

against feature patterns. It embodies a fundamental form of prediction that can be activated either bottom-up by input data, or top-down by an expectation that predictively primes a class of events whose future occurrence is sought. Match-based learning allows memories to change only when input from the external world is close enough to internal expectations, or when something completely new occurs.

11. Resonance vs. reset implies gamma vs. beta oscillations

The Synchronous Matching ART (SMART) model (Grossberg & Versace, 2008) clarifies how ART processes may be embodied in laminar thalamocortical circuits that experience spiking dynamics. SMART demonstrates how a top-down attentive match may lead to fast gamma oscillations that facilitate spike-timing dependent plasticity (STDP), whereas mismatch and reset can lead to slower beta oscillations that help to prevent mismatched events from being learned. This match–mismatch gamma–beta story seems to occur in quite a few brain systems, with examples of data supporting the Grossberg–Versace prediction having recently been reported in cortical area V1, hippocampus, and frontal eye fields (see Section 38).

12. Mixing unsupervised with supervised learning

The ART category learning process works well under both unsupervised and supervised conditions. Variants of the ARTMAP architecture can carry out both types of learning (e.g., Carpenter et al. (1992)). Unsupervised learning means that the system can learn how to categorize novel input patterns without any external feedback. Supervised learning uses predictive errors to let the system know whether it has categorized the information correctly or not.

Supervision can force a search for new categories that may be culturally determined, and are not based on feature similarity alone. For example, separating the featurally similar letters *E* and *F* into separate recognition categories is culturally determined. Such error-based feedback enables variants of *E* and *F* to learn their own category and top-down expectation, or prototype. The complementary, but interacting, processes of attentive-learning and orienting-search together realize a type of error correction through hypothesis testing that can build an ever-growing, self-refining internal model of a changing world.

13. Mismatch-activated nonspecific arousal regulates reset and search

13.1. Complementary attentional and orienting systems

The attentional and orienting systems in an ART network (Fig. 2) also experience complementary informational deficiencies. At the moment when a predictive error occurs, the system does not know why the currently active category was insufficient to predict the correct outcome. In particular, when the orienting system gets activated by a mismatch in the attentional system, the orienting system has no way of knowing what went wrong in the attentional system. Thus, the attentional system has information about how inputs are categorized, but not whether the categorization is correct, whereas the orienting system has information about whether the categorization is correct, but not what is being categorized. How, then, does the orienting system cope with the daunting challenge of resetting and driving a memory search within the attentional system in a way that leads to a better outcome after the search ends.

13.2. Novelty-sensitive nonspecific arousal: novel events are arousing!

Because the orienting system does not know what cells in the attentional system caused the predictive error, its activation needs to influence all potential sources of the error equally. Thus, mismatch triggers a burst of *nonspecific arousal* that activates all cells in the attentional system equally. In other words, novel events are arousing! Said in a more philosophical way, a novelty-sensitive burst of nonspecific arousal implements the principle of sufficient reason. As illustrated in Fig. 2, the current state of activation of the attentional system interacts with such an arousal burst to selectively reset cells that caused the mismatch, and to thereby drive a search leading to a better predictive outcome.

13.3. Medium-term memory: habituated transmitter gates in non-stationary hypothesis testing

The search process exploits a type of medium-term memory acting at category cells that is realized by activity-dependent habituated transmitter gates (Carpenter & Grossberg, 1990; Grossberg, 1968c, 1972b, 1976b). See Eq. (2). Such habituated transmitter gates are sometimes called *depressing synapses* after the work of Abbott, Varela, Sen, and Nelson (1997), which experimentally confirmed predicted properties of such gates in the visual cortex. Habituated transmitter gates have been used to help explain a wide range of data about processes other than category learning, including the dynamics of visual perception, cognitive–emotional interactions, and sensory–motor control (Francis & Grossberg, 1996; Francis, Grossberg, & Mingolla, 1994; Gaudiano & Grossberg, 1991, 1992; Grossberg, 1972b, 1980, 1984a, 1984b).

Due to habituated gating, recently active cells are more habituated than inactive cells. Activity-dependent habituation interacts with self-normalizing competition among the category cells to help suppress cells that are most active when the arousal burst is received. Once the maximally activated cells are suppressed by this combination of habituation and competition during the search cycle, the self-normalizing network activity is available to enable other cells, which got smaller inputs than the original winning cells, to become active in the next time interval. This cycle of mismatch–arousal–reset continues until resonance can again occur.

The ability of the category cell network to self-normalize its total activity enables the activities of these categories to be interpreted as a kind of real-time probability distribution, and the ART search cycle to be interpreted as a kind of probabilistic hypothesis testing and decision making that works in response to non-stationary time series of input patterns.

14. Vigilance regulates the content of conscious experiences: exemplars and prototypes

14.1. Vigilance controls whether concrete or general categories are learned

What combinations of features or other information are bound together into conscious object or event representations? One popular view in cognitive psychology is that exemplars, or individual experiences, are learned, because humans can have very specific memories. For example, we can all recognize the faces of our friends. On the other hand, storing every remembered experience as an exemplar could lead to a combinatorial explosion of memory, as well as to unmanageable problems of memory retrieval. A possible way out is suggested by the fact that humans can learn prototypes which represent general properties

of the environment (Posner & Keele, 1968). For example, we can recognize that everyone has a face. But then how do we learn specific episodic memories? ART provides an answer to this question that overcomes problems faced by earlier models.

ART prototypes are not merely averages of the exemplars that are classified by a category, as is often assumed in classical prototype models. Rather, they are the actively selected *critical feature patterns* upon which the top-down expectations of the category focus attention. The generality of the information that is coded by these critical feature patterns is controlled by a gain control process, called *vigilance* control, which can be influenced by environmental feedback or internal volition (Carpenter & Grossberg, 1987). Low vigilance permits the learning of general categories with abstract prototypes. High vigilance forces a memory search to occur for a new category when even small mismatches exist between an exemplar and the category that it activates. As a result, in the limit of high vigilance, the category prototype may encode an individual exemplar.

14.2. Vigilance is computed in the orienting system

Vigilance is computed within the orienting system of an ART model (Fig. 2(b)–(d)). It is here that bottom-up excitation from all the active features in an input pattern I are compared with inhibition from all the active features in a distributed feature representation across F_1 . If the ratio of the total activity across the active features in F_1 (that is, the “matched” features) to the total activity due to all the features in I is less than a *vigilance* parameter ρ (Fig. 2(b)), then a nonspecific reset wave is activated (Fig. 2(c)), which can drive the search for another category with which to classify the exemplar. This can be accomplished by letting ρ multiply the bottom-up inputs I to the orienting system; that is, ρ is the gain of the inputs to the orienting system. The orienting system is then activated when the total excitatory input ρI is greater than the total inhibition from the features X^* across F_1 that survive top-down matching; that is, when $\rho |I| - |X^*| > 0$, where $|\cdot|$ denotes the number of positive inputs or matched features. This inequality can be rewritten as $\rho > |X^*| |I|^{-1}$ to show that the orienting system is activated whenever ρ is chosen higher than the ratio of active X^* matched features in F_1 to total features in I . In other words, the vigilance parameter controls how bad a match can be before search for a new category is initiated. If the vigilance parameter is low, then many exemplars can all influence the learning of a shared prototype, by chipping away at the features that are not shared with all the exemplars. If the vigilance parameter is high, then even a small difference between a new exemplar and a known prototype (e.g., F vs. E) can drive the search for a new category with which to represent F .

14.3. Minimax learning via match tracking: learning the most general predictive categories

One way to control vigilance is by a process of *match tracking* (Carpenter & Grossberg, 1991; Carpenter et al., 1992). Here, in response to a predictive error (e.g., D is predicted in response to F), the vigilance parameter ρ increases just enough to trigger reset and search for a better-matching category. Match tracking gives up the minimum amount of generalization in the learned categories to search for a better-matching category. In other words, vigilance “tracks” the degree of match between input exemplar and matched prototype. Because match tracking increases vigilance by the minimum amount to trigger a reset and search for a new category, it realizes a Minimax Learning Rule that conjointly *maximizes* category generality while it *minimizes* predictive error. Match tracking thus uses the least memory resources that can correct errors in classification.

Because the baseline level of vigilance is initially set at the lowest level that has led to predictive success in the past, ART models try to learn the most general category that is consistent with the data. This tendency can, for example, lead to the type of overgeneralization that is seen in young children until further learning leads to category refinement. However, because vigilance can vary during match tracking in a manner that reflects current predictive success, 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 learn abstract prototypes with which to recognize abstract categories of faces and dogs, as well as “exemplar prototypes” with which to recognize individual views of faces and dogs, depending on task requirements.

15. Memory consolidation and the emergence of rules: direct access to globally best match

As sequences of inputs are practiced over learning trials, the search process eventually converges upon stable categories. It has been mathematically proved (e.g., Carpenter and Grossberg (1987)) that familiar inputs directly access the category whose prototype provides the globally best match, without undergoing any search, while unfamiliar inputs engage the orienting subsystem to trigger memory searches for better categories until they become familiar. In other words, ART provides a solution of the local minimum problem that various other algorithms, such as back propagation (Baldi & Hornik, 1989; Gori & Tesi, 1992), do not solve. This process of search and category learning continues until the memory capacity, which can be chosen arbitrarily large, is fully utilized.

15.1. Memory consolidation and medial temporal amnesia

The process whereby search is automatically disengaged is a form of *memory consolidation* that emerges from network interactions. The first example of memory consolidation that was described by ART concerns cortico-hippocampal interactions, and proposed how a hippocampal ablation may cause symptoms of medial temporal amnesia (Carpenter & Grossberg, 1993). Emergent consolidation does not preclude structural consolidation at individual cells, since the amplified and prolonged activities that subserve a resonance may be a trigger for learning-dependent cellular processes, such as protein synthesis, synapse formation, and transmitter production.

15.2. Learning of fuzzy IF-THEN rules by a self-organizing production system

It has been proved that the adaptive weights which are learned by some ART models can, at any stage of learning, be translated into fuzzy IF-THEN rules (Carpenter et al., 1992). Thus the ART model is a self-organizing rule-discovery production system as well as a neural network. These examples show that the claims of some cognitive scientists and AI practitioners that neural network models cannot learn rule-based behaviors are as incorrect as the claims that neural models cannot learn symbols.

16. Where’s Waldo? Positionally-invariant recognition codes and positionally precise actions

What kind of categories can an ART system learn? How can such learning be incorporated into the kinds of percep-




Fig. 3. Where’s Waldo: the position in a scene of a valued target can be located by linking What stream recognition to Where stream action. Interactions between cortical areas ITp, ITa, amygdala, orbitofrontal cortex (ORB), and posterior parietal cortex (PPC) can bridge between positionally-invariant ITa object categories that are selected by motivated attention, and parietal cortical representations of object positions. The numbers indicate the order of pathway activations. If there are two numbers, the larger one represents the stage when feedback activates that pathway. See text for details.

Source: Reprinted with permission from Grossberg (2009b).

tion–cognition–emotion–action cycles that are characteristic of many mammalian behaviors? Actions directed towards valued goal objects cannot be made until goal objects are recognized, their currently perceived value used to help direct attention towards them, and their spatial locations specified. As noted in Fig. 1, the What cortical stream learns object representations that are increasingly independent of object position and size, whereas the Where cortical stream represents object positions and how to move. Interactions between the What and Where streams overcome these complementary informational deficiencies to generate actions towards recognized objects that are currently valued.

Whereas object representations in posterior inferotemporal cortex (ITp) combine feature and positional information, object representations in anterior inferotemporal cortex (ITa) are more positionally invariant. These two types of representations are linked by reciprocal learned connections, as described by ART. ITp representations also project to the posterior parietal cortex (PPC) as target locations of an object. Given this background, consider what happens when multiple objects in a scene all try to activate their corresponding ITp and ITa representations. Suppose that a particular ITa category represents a valued goal object in that situation. As explained in Section 42, the ITa representation can get amplified by an inferotemporal–amygdala–orbitofrontal resonance. When this happens, the amplified ITa representation can better compete for object attention, and can send larger top-down priming signals to its ITp representations. The ITp representation that corresponds to the valued object is thereby selectively amplified, and sends an amplified signal to the parietal cortex, where its target location can win the competition for where the next movement will go. See Fig. 3. This scheme can help to solve the Where’s Waldo problem, or the rapid discovery of a desired goal object in a cluttered scene.

17. ARTSCAN: learning invariant object categories using attentional shrouds

17.1. Solving the View-to-Object Binding Problem while scanning a scene

One crucial part of this behavioral cycle is the learning of view-invariant and positionally-invariant categories by the brain. To understand how this happens, several basic questions need to be answered: What is an object? How can multiple views of an object that is seen in different positions and distances with respect to an observer all activate an invariant object category at a sufficiently high processing level? How does the brain learn invariant object

properties under both unsupervised and supervised learning conditions? How does the brain learn to bind multiple views of an object into a view-invariant and positionally-invariant object category while freely scanning a scene with eye movements?

To answer these questions, one also needs to solve the following basic View-to-Object Binding Problem: As eyes scan a scene, two successive eye movements may focus on different parts of the same object or on different objects. How does the brain avoid erroneously classifying views of different objects together, even before the brain knows what the object is? One cannot say that the brain does this by knowing that some views belong together whereas others do not, because this can happen even before the brain has a concept of what the object is. Indeed, such scanning eye movements may be used to learn the object concept in the first place.

17.2. Coordinating spatial and prototype attention during invariant category learning

The ARTSCAN model (Fig. 4) clarifies how the brain uses scanning saccadic eye movements to learn view-invariant object categories (Cao et al., 2011; Fazl et al., 2009; Foley et al., 2012; Grossberg, 2007b, 2009b; Grossberg et al., 2011). The discussion about ART above considered only one form of object attention (Posner, 1980) in the What cortical stream, the kind that focuses attention upon the critical feature pattern of a category prototype. ARTSCAN explains how this kind of object attention, called *prototype attention*, interacts with *spatial attention* (Duncan, 1984) in the Where cortical stream to direct eye movements that explore object surfaces. ARTSCAN makes a major new prediction about how spatial and object attention are related; namely, spatial attention coordinates the learning of invariant object categories during free viewing conditions. The ART dynamics schematized in Fig. 2 learn the view-specific categories that are bound together through such coordination into view-invariant and positionally-invariant object categories.

The process begins when a view-specific category of a novel object is learned, and activates cells at a higher cortical level that will become a view-invariant object category as multiple view categories are associated with it. Indeed, as the eyes move around an object surface, multiple view-specific categories are learned of the object (e.g., in ITp; see Fig. 4) and are associated with the emerging invariant object category (e.g., in ITa; see Fig. 4). How does the brain know how to prevent the invariant object category from being reset while it is getting associated with multiple view-specific categories of a single object, each of which must be reset to enable the next view-specific category to be activated and learned?

17.3. Attentional shroud inhibits reset of an invariant object category during object search

ARTSCAN predicts that a *pre-attentively* formed surface representation activates an *attentional shroud* (Tyler & Kontsevich, 1995), or form-fitting distribution of spatial attention, even before the brain can recognize the surface as representing a particular object. This shroud persists within the Where Stream during active scanning of an object. The shroud protects the view-invariant category from getting reset, even while view-specific categories are reset, as the eyes explore an object. The shroud does this by inhibiting the ITa reset mechanism (see inhibition from Spatial Attention to Category Reset in Fig. 4).

How does the shroud persist during active scanning of an object? A *surface-shroud resonance* arises due to positive feedback interactions between a surface representation (e.g., in cortical area V4) and spatial attention (e.g., in posterior parietal cortex, or PPC), and focuses spatial attention upon the object to be learned

(Fig. 4). When the shroud collapses, the Category Reset stage is disinhibited, giving rise to a transient burst of inhibition that resets the active invariant object category. The collapse of the shroud also enables the eyes to move to another surface, whereupon new view-specific and view-invariant object categories can be learned. The cycle can then repeat itself.

17.4. Human and monkey data support shroud reset properties

Chiu and Yantis (2009) used rapid event-related MRI in humans to provide evidence for the ARTSCAN prediction of how a surface-shroud resonance in the Where stream protects an emerging view-invariant category from being prematurely reset in the What stream when each of the view-specific categories that are associated with it is reset. These authors found that a shift of spatial attention evokes a transient signal in the medial superior parietal lobule that corresponding to a shift in categorization rules. In ARTSCAN, collapse of an attentional shroud (spatial attention shift) disinhibits the parietal reset mechanism (transient signal) that leads to collapse of the previous view-invariant object category and instatement of a new one (shift in categorization rules).

Cao et al. (2011) have used the positional ARTSCAN (pARTSCAN) extension of the ARTSCAN model to simulate neurophysiological data of Li and DiCarlo (2008; see also Li and DiCarlo (2010)) showing that views from different objects can be merged within intertemporal categories when monkeys are presented with an object that is swapped with another object during eye movements to foveate the original object. Why does not such a merging of recognition categories occur all the time, thereby leading to catastrophic forgetting of learned recognition categories? The model quantitatively simulates the swapping data by showing how the swapping procedure fools the spatial attention reset mechanism by instating the swap before the animal can shift its spatial attention.

This result, among others, may be used to develop new ways to test how humans and animals learn to pay attention to important targets while they learn invariant object categories with which to recognize and predict them. As noted above, Chiu and Yantis (2009) used rapid event-related MRI in humans to support the model's prediction that reset is mediated by a transient burst of activation in the parietal cortex which is activated by a shift of spatial attention, also in the parietal cortex. One new experiment in humans would be to repeat the swapping experiment in humans and use MRI to test if there is a transient parietal burst during the swap. The prediction is that there would not be a burst when there is rapid merging of the second object into the category. As the delay between the initial target and the swap increases, a reset should occur when merged categories are no longer learned.

Many other paradoxical data may also be explained by these concepts, including how spatial attention can increase the perceived brightness of a surface (Carrasco, Penpeci-Talgar, & Eckstein, 2000; Reynolds & Desimone, 2003), how predictive remapping of eye position occurs (Duhamel, Colby, & Goldberg, 1992; Gottlieb, Kusunoki, & Goldberg, 2005; Melcher, 2007), how the eyes can prefer to move within the same object for awhile (Theeuwes, Mathot, & Kingstone, 2010), and what sort of category invariance can be learned (Grossberg et al., 2011; Zoccolan et al., 2007).

17.5. Conscious perception of surface-shroud resonances: linking perception to recognition

ARTSCAN also provides new insights into basic issues such as: What do we consciously see? How is seeing related to recognition, and how does recognition of individual objects fail during conditions of perceptual crowding (Green & Bavelier, 2007;




Fig. 5. The LAMINART model clarifies how bottom-up, horizontal, and top-down interactions within and across cortical layers in V1 and V2 interblob and pale stripe regions, respectively, carry out bottom-up adaptive filtering, horizontal grouping, and top-down attention. Similar interactions seem to occur in all six-layered cortices. See text for details.

Source: Reprinted with permission from Raizada and Grossberg (2001).

it unified two major streams of research activity. The two streams of research activity are:

(1) ART as a theory of category learning and prediction. This stream emphasized bottom-up and top-down interactions within higher-level cortical circuits, such as cortical areas V4, inferotemporal cortex, and prefrontal cortex, during the learning of visual recognition categories;

(2) FACADE (Form-And-Color-And-DEpth) as a theory of 3D vision and figure-ground perception (Cao & Grossberg, 2005; Fang & Grossberg, 2009; Grossberg, 1994, 1997; Grossberg, Bullock, & Dranias, 2008; Grossberg & McLoughlin, 1997; Grossberg & Swaminathan, 2004; Grossberg & Yazdanbakhsh, 2005; Grossberg, Yazdanbakhsh, Cao, & Swaminathan, 2008). This stream emphasized bottom-up and horizontal interactions for completion of boundaries during perceptual grouping, and for filling-in of surface brightness and color. These interactions were proposed to occur in lower cortical processing areas such as V1, V2, and V4.

19. Laminar vision, speech, and cognition models: LAMINART, cARTWORD, LIST PARSE

The unification of these two research streams in LAMINART proposed how all cortical areas combine bottom-up, horizontal, and top-down interactions, thereby beginning to functionally clarify why all granular neocortex has a characteristic architecture with six main cell layers (Felleman & Van Essen, 1991), and how these laminar circuits may be specialized to carry out different types of biological intelligence. In particular, this unification suggested how variations of a shared laminar cortical design could be used to explain psychological and neurobiological data about vision, speech, and cognition:

Vision. LAMINART integrates bottom-up and horizontal processes of 3D boundary formation and perceptual grouping, surface filling-in, and figure-ground separation with top-down attentional matching in cortical areas such as V1, V2, and V4 (Cao & Grossberg, 2005; Grossberg, 1999; Grossberg & Raizada, 2000; Grossberg & Swaminathan, 2004; Grossberg & Yazdanbakhsh, 2005; Raizada & Grossberg, 2001).

Speech. cARTWORD models how bottom-up, horizontal, and top-down interactions within a hierarchy of laminar cortical processing stages, modulated by the basal ganglia, can generate a conscious speech percept that is embodied by a resonant wave of activation that occurs between acoustic features, acoustic item chunks, and list chunks (Grossberg & Kazerounian, 2011, see Sections 55–56). Chunk-mediated gating allows speech to be heard in the correct temporal order, even when what is consciously heard depends upon using future context to disambiguate noise-occluded sounds, as occurs during phonemic restoration.

Cognition. LIST PARSE models how bottom-up, horizontal, and top-down interactions within the laminar circuits of lateral prefrontal cortex may carry out working memory storage of event sequences within layers 6 and 4, how unitization of these event sequences through learning into list chunks may occur within layer 2/3, and how these stored sequences can be recalled at variable rates that are under volitional control by the basal ganglia (Grossberg & Pearson, 2008, see Sections 50–52). In particular, the model uses variations of the same circuitry to quantitatively simulate human cognitive data about immediate serial recall and free recall, and monkey neurophysiological data from the prefrontal cortex obtained during sequential sensory-motor imitation and planned performance.

This emerging unified theory of how variations of a shared laminar neocortical design can carry out multiple types of biological intelligence is also of interest in technology, where having a unified VLSI chip set for multiple types of biological intelligence would revolutionize computer science in general, and the design of autonomous adaptive mobile robots in particular. The DARPA SyNAPSE program is currently pursuing such a possibility (<http://en.wikipedia.org/wiki/SyNAPSE>).

20. Invisible boundaries, conscious visibility, surface-shroud resonance, and parietal neglect

While on the topic of boundaries and surfaces, it is instructive to mention some basic properties of perceptual boundaries and surfaces that are relevant to the ART prediction that “all conscious states are resonant states” (Sections 1 and 5). A basic discovery and prediction about perceptual boundaries and surfaces is that they obey complementary computational laws (Grossberg, 1994):

Boundaries are completed *inwardly* between pairs of similarly oriented and collinear cell populations (the so-called *bipole grouping* property; see Section 28). This inward and oriented boundary process enables boundaries to complete across partially occluded object features. Boundaries also pool inputs from opposite contrast polarities, so are *insensitive* to contrast polarity. This pooling process enables boundaries to form around objects that are seen in front of backgrounds whose contrast polarities with respect to the object reverse around the object’s perimeter.

In contrast, surface brightness and color fill-in *outwardly* in an *unoriented* manner until they reach object boundaries or dissipate due to their spread across space (Grossberg & Todorovic, 1988). This filling-in process is *sensitive* to individual contrast polarities. These computational properties of boundaries and surfaces are manifestly complementary.

Boundaries form within the cortical stream that goes through interblobs in cortical area V1 to pale stripes in V2 and beyond to

V4. Surfaces form within the cortical stream through blobs in V1 to thin stripes in V2 and beyond to V4.

Because boundaries pool inputs from opposite contrast polarities, they do not carry a visible brightness or color signal. *All boundaries are invisible*, or amodal, within the boundary cortical stream. Visibility is a property of surfaces, which do not pool over contrast polarities. Given the ART prediction that “all conscious states are resonant states”, this fact raises the question: What sort of resonance supports a conscious percept of surface brightness or color? As noted above, I have predicted that such a consciously visible percept is supported by a *surface-shroud resonance* between visual cortical areas and the parietal cortex that is predicted to play a role in learning invariant object categories (Section 17). This prediction illustrates how ART clarifies mechanistic relationships between the CLEARs processes of consciousness, learning, expectation, attention, resonance, and synchrony.

This prediction also clarifies how, for example, parietal neglect may occur (Driver & Mattingley, 1998; Mesulam, 1999); that is, how a parietal lesion leads a patient to consciously ignore the corresponding region of space. ART proposes that this happens because, despite the fact that the lesion may leave the visual cortex intact, the corresponding surface-shroud resonances cannot form to support conscious percepts of object surfaces. The predicted link between surface-shroud resonances and the control of saccadic eye movements for searching and learning about object surfaces is consistent with the expectation that abnormal search with saccadic eye movements may occur during parietal neglect. Husain et al. (2001) reported a failure to remember which locations have already been examined through saccadic movements in a patient with left neglect following a right parietal infarction, leading to repeated re-fixation of locations. Understanding such results in detail would require an analysis of whether the parietal reset mechanism was also damaged, and whether other parietally-linked mechanisms for accumulating evidence over eye movements were also affected by the lesion (Huang & Grossberg, 2010, Section 54).

21. ART Matching Rule in laminar cortex: folded feedback

Fig. 5 summarizes how the ART Matching Rule is proposed to be realized within the laminar circuits of visual cortex, and by extension within all granular neocortical circuits. As noted in Fig. 5(b), corticocortical feedback axons tend preferentially to originate in layer 6 of a higher cortical area (e.g., V2) and to terminate in layer 1 of the lower cortex (e.g., V1) (Pandya & Yeterian, 1985; Salin & Bullier, 1995). This pathway activates layer 1 apical dendrites of layer 5 cells, which relay them to layer 6 cells in V1 (Cauller & Connors, 1994; Rockland, 1994). The triangle in Fig. 5(b) represents such a layer 5 pyramidal cell. Several other routes through which feedback can pass into V1 layer 6 exist (see Raizada and Grossberg (2001) for a review). Having arrived in layer 6, the feedback is then “folded” back up into layer 4 by feedforward connections that form an on-center off-surround pathway from layers 6 to 4 (Bullier et al., 1996). The on-center in this pathway is predicted to be modulatory. In support of this hypothesis, Hupé et al. (1997, p. 1031) note: “feedback connections from area V2 modulate but do not create center-surround interactions in V1 neurons”. Thus top-down feedback from layer 6 of V2 is predicted to be able to supraliminally activate layer 6 of V1 but not layer 4 of V1. This hypothesis is consistent with neurophysiological data from ferret visual cortex showing that the layer 6-to-4 circuit is functionally weak (Wittmer, Dalva, & Katz, 1997).

In summary, top-down attention may be realized by an inter-cortical, top-down, modulatory on-center, off-surround network that projects from layer 6 in a higher cortical area to layer 6 in a lower cortical area before being “folded” back into layer 4 via a modulatory on-center, off-surround network from layers 6 to 4.

22. Attention–Preattention Interface: a decision network in the cortical deeper layers

LAMINART circuits illustrate how bottom-up, horizontal, and top-down interactions may all interact with the modulatory on-center, off-surround networks that exist within and between cortical layers 6 and 4. Fig. 5 schematizes these circuits for the first few stages of visual processing: the lateral geniculate nucleus (LGN) and cortical areas V1 and V2.

22.1. Bottom-up contrast normalization

Fig. 5(a) shows that the LGN provides bottom-up activation to layer 4 via two routes. Any engineer or VLSI chip designer would immediately ask why two routes are needed. Why waste “extra wire”? LAMINART proposes the following answer: One connection from LGN to layer 4 is via layer 6, which activates layer 4 via the modulatory on-center, off-surround network. This route is not sufficient to activate layer 4 from LGN because the on-center is only modulatory. Hence, a driving connection exists directly from LGN to layer 4. Taken together, these connections form a driving on-center off-surround network that contrast-normalizes the inputs that layer 4 receives from LGN (Grossberg, 1973, 1980; Grossberg & Mingolla, 1985; Heeger, 1992).

22.2. Top-down intercortical attention

As noted above and diagrammed in Fig. 5(b), (e), folded feedback can carry attentional signals from layer 6 of higher cortical areas to layer 4 of lower cortical areas via the layer 6-to-4 modulatory on-center, off-surround network, thereby realizing the ART Matching Rule. Fig. 5(d) shows that such a network also exists top-down from layer 6 of V1 to the LGN. Here, the on-center feedback selectively enhances LGN cells that are consistent with the activation that they cause (Sillito et al., 1994), and the off-surround contributes to length-sensitive (endstopped) responses that facilitate grouping perpendicular to line ends.

22.3. Horizontal interlaminar feedback

As shown in Fig. 5(c), layer 2/3 possesses long-range horizontal connections that are used for perceptual grouping of contours, textures, and shading (see Section 28). These layer 2/3 cells are activated by the deeper cortical layers. In particular, like-oriented layer 4 simple cells that are sensitive to opposite contrast polarities compete (not shown) before generating half-wave rectified outputs that converge onto layer 2/3 complex cells in the column above them. Because the complex cells pool inputs across opposite contrast polarities, they are the earliest cortical stage that ensure “boundaries are invisible” (Section 20). Just like attentional signals from higher cortex, groupings that form within layer 2/3 also send activation into the folded feedback path, to enhance their own positions in layer 4 beneath them via the layer 6-to-4 on-center, and to suppress input to other groupings via the layer 6-to-4 off-surround. There exist direct connections from layer 2/3 to 6 in macaque V1, as well as indirect routes via layer 5.

22.4. A hierarchical design: hierarchical propagation of priming

Fig. 5(e) shows that V2 repeats the laminar pattern of V1 circuitry, but at a larger spatial scale. In particular, the horizontal layer 2/3 connections have a longer range in V2, allowing above-threshold perceptual groupings between more widely-spaced inducing stimuli to form (Amir, Harel, & Malach, 1993). V1 layer 2/3 projects up to V2 layers 6 and 4, just as LGN projects to

layers 6 and 4 of V1. Higher cortical areas send feedback into V2 which ultimately reaches layer 6, just as V2 feedback acts on layer 6 of V1 (Sandell & Schiller, 1982). Feedback paths from higher cortical areas straight into V1 (not shown) can complement and enhance feedback from V2 into V1. Top-down attentional signals to layer 1 may also directly modulate groupings via the apical dendrites in layer 1 of excitatory and inhibitory layer 2/3 cells (Lund & Wu, 1997; Rockland & Virga, 1989). By activating both excitatory and inhibitory cells in layer 2/3 (Fig. 5(e)), the inhibitory cells may balance the excitatory cell activation, thereby creating a modulatory attentional response of grouping cells in layer 2/3 via another modulatory on-center off-surround attentional network. Roelfsema et al. (1998) and Wanning, Stanisor, and Roelfsema (2011) have demonstrated this kind of attentional spreading along perceptual groupings (see Section 27).

Because the ART Matching Rule is realized by a circuit with a modulatory on-center, and top-down signals can jump from layer 6 of a higher cortical level to layer 6 of a lower cortical level, a top-down signal from a much higher cortical level (e.g., the prefrontal cortex) can modulate, or prime, all the cortical levels below it (e.g., V4, V2, V1, LGN) with a task-selective processing constraint, albeit possibly to different degrees (Fig. 5(e)). The entire hierarchy is then “ready” to process incoming inputs constrained by the processing bias embodied in the prime.

23. Three kinds of object-based attention: boundary, prototype, and surface attention

The manner in which top-down attention and pre-attentive perceptual grouping are interfaced within the cortical layers (Fig. 5(e)) enables attention to focus on an entire object boundary, thereby not only influencing what objects are selectively attended, but also what groupings may be perceived. This is true because the same layer 6-to-4 competition, or selection, circuit may be activated by pre-attentive grouping cells in layer 2/3 (Fig. 5(c)), as well as by top-down attentional pathways (Fig. 5(b)). Layer 4 cells can then, in turn, activate the layer 2/3 cells where perceptual groupings are initiated (Fig. 5(c)). In all, a top-down attentional prime can enhance the activation of layer 4 cells whose features are attended; the layer 4 cells can enhance the grouping in layer 2/3 that is starting to form above them; and feedback from layer 2/3 to layers 6-then-4-then-2/3 can enable this attentional advantage to propagate along the entire grouping. In summary, when ambiguous and complex scenes are being processed, intracortical but interlaminar folded feedback enables stronger groupings that are starting to form in layer 2/3 to inhibit weaker groupings, whereas intercortical folded feedback from top-down attention enables higher-order processing constraints to bias which groupings will be selected.

The kind of attention that was just summarized may be called *boundary attention*. The kind of attention whereby ART categorization networks prime a learned critical feature pattern may be called *prototype attention*. Both of these kinds of attention tend to operate within the cortical What stream. A third kind of attention also exists, namely *surface attention*, which clarifies how the brain orients and attends to the spatial location of objects as part of surface-shroud resonances.

A general conclusion of this summary is that using the phrase “object attention” is insufficient to distinguish the functional and mechanistic differences between boundary, prototype, and surface attention. ART and LAMINART provide precise functional distinctions between these different types of attention, and shows how they may interact during perception and recognition.

24. The link between attention and learning

Various experiments have provided data supporting the ART prediction that top-down feedback can modulate plasticity. Psychophysically, the role of attention in controlling adult plasticity during perceptual learning was demonstrated by Ahissar and Hochstein (1993). Gao and Suga (1998) reported physiological evidence that acoustic stimuli cause plastic changes in the inferior colliculus (IC) of bats only when the IC received top-down feedback from auditory cortex. These authors also reported that plasticity is enhanced when the auditory stimuli were made behaviorally relevant, consistent with the ART proposal that top-down feedback allows attended, and thus relevant, stimuli to be learned, while suppressing unattended irrelevant ones. Evidence that cortical feedback controls thalamic plasticity in the somatosensory system has been reported by Krupa, Ghazanfar, and Nicolelis (1999) and by Parker and Dostrovsky (1999). These findings are reviewed by Kaas (1999).

Models of intracortical grouping-activated feedback and intercortical attention-activated feedback have shown that either type of feedback can rapidly synchronize the firing patterns of higher and lower cortical areas (Grossberg & Grunewald, 1997; Grossberg & Somers, 1991). ART puts this result into a larger perspective by suggesting how resonance may lead to attentive synchronization, which may, in turn, trigger cortical learning by enhancing the probability that “cells that fire together wire together”. An excellent discussion of top-down cortical feedback, synchrony, and their possible relations to the ART model is given by Engel et al. (2001).

25. How can early development and perceptual learning occur without attention?

LAMINART solves a basic design problem that early non-laminar versions of ART did not handle. The problem is called the Attention–Preattention Interface Problem because it may be solved by the way in which laminar circuits combine pre-attentive and attentive processing constraints, as summarized above (Grossberg, 1999). This problem may be stated in several ways.

One way is to ask how early cortical development can occur in a stable way through time. The hypothesis that attentional feedback exerts a controlling influence over plasticity in sensory cortex does not imply that unattended stimuli can never be learned. Indeed, plasticity can take place during early development, before top-down attentional circuits may have developed between cortical areas. Grossberg (1999) noted that, were this not possible, an infinite regress could be created, since a lower cortical level like V1 might then not then be able to stably develop unless it received attentional feedback from V2, but V2 itself could not develop unless it had received reliable bottom-up signals from V1. How do the laminar circuits of visual cortex avoid this infinite regress, without causing massive instability, in response to stimuli that occur with high statistical regularity in the environment (e.g., Grossberg and Williamson (2001))? How does this process continue to fine-tune sensory representations in adulthood even in cases where focused attention may not be explicitly allocated, and slow perceptual learning may take place without conscious awareness (Pilly, Grossberg, & Seitz, 2010; Seitz & Watanabe, 2003; Watanabe, Nanez, & Sasaki, 2001)?

26. Reconciling visual boundary completion with the ART Matching Rule

26.1. Attention–Preattention Interface Problem

Another way to state the Attention–Preattention Interface Problem is to note that, despite experimental and mathematical

support for the ART Matching Rule, basic data about visual perception, at least at first glance, seem to conflict with its properties. In particular, the ART Matching Rule says that top-down attention typically has a modulatory on-center, so that bottom-up input to cells in the on-center is needed to fire them above threshold. However, if the ART Matching Rule holds, then how can pre-attentive groupings, such as illusory contours, form over positions that receive no bottom-up inputs? Moreover, these illusory contours are formed via horizontal connections in visual cortex that develop through a process that includes visually-driven learning, and thus need to solve the stability–plasticity dilemma (see Grossberg and Williamson (2001) for a review). If the ART Matching Rule is necessary to prevent catastrophic forgetting, then how can we see illusory contours without destabilizing cortical development and learning?

This issue is so important that it is worth stating it in greater detail. The ART Matching Rule has three aspects: first, incoming sensory signals that receive matching top-down excitatory feedback are enhanced; second, non-matching inputs that do not receive excitatory feedback are suppressed; and third, top-down feedback on its own is modulatory; that is, unable to produce above-threshold activity in the lower area in the absence of incoming bottom-up signals. The conceptual challenge is this: If ART matching is needed to stabilize cortical development and learning, and if ART matching requires that suprathreshold activation can occur only where there are bottom-up inputs, then does not the existence of illusory contours contradict the ART Matching Rule, since such groupings form over positions that receive no bottom-up inputs, and yet do not seem to destabilize cortical development or learning? If the brain had not solved this problem, anyone could roam through the streets of a city and destabilize pedestrians' visual systems simply by showing them images of Kanizsa squares! The absurdity of this possibility indicates the fundamental nature of this issue.

26.2. When top-down signals are driving: volitional control by the basal ganglia

Before going on, it should also be noted that top-down signals can drive their target cells to suprathreshold activation during percepts of visual imagery or internal thought and planning. This property is consistent with the ART Matching Rule, because a volitional signal from the basal ganglia can change the balance between excitation and inhibition within the modulatory on-center to favor excitation. How this may happen has been described in Grossberg (2000a); also see Sections 7 and 44. As noted in Section 7, a similar basal ganglia volitional mechanism has been predicted to determine when sequences of events are stored in a prefrontal cognitive working memories (Grossberg & Pearson, 2008) and how the span of spatial attention may be volitionally altered during visual search tasks (Foley et al., 2012).

27. A pre-attentive grouping is its own attentional prime

The LAMINART model proposes how the brain uses its laminar circuits to solve the Attention–Preattention Interface Problem in an ingenious, parsimonious, and simple way. Here is where the laminar cortical circuit that combines pre-attentive and attentive processing constraints plays a key role: Both *intercortical* attentional feedback and *intracortical* grouping feedback share the same competitive selection circuit from layer 6-to-4. In particular, when a horizontal grouping starts to form in layer 2/3, it activates the intracortical feedback pathway from layer 2/3-to-6, which activates the modulatory on-center off-surround network from layer 6-to-4. This feedback pathway helps to select which cells will remain active to participate in a winning grouping. But this is the same network that realizes the ART Matching Rule

when it stabilizes cortical development and learning using top-down attention from a higher cortical area. In other words, the layer 6-to-4 selection circuit, which in the adult helps to choose winning groupings via *intercortical* top-down attentional signals, also helps to assure in the developing brain, using *intracortical* but *interlaminar* grouping feedback, that the ART Matching Rule holds at every position along a grouping. Because the ART Matching Rule holds, only the correct combinations of cells can “fire together and wire together”, and hence stability is achieved. *Intracortical* feedback via layers 2/3-to-6-to-4-to-2/3 can realize this selection process even before *intercortical* attentional feedback can develop. This property is sometimes summarized with the phrase: “The pre-attentive grouping is its own attentional prime” (Grossberg, 1999).

In summary, by joining together bottom-up (*interlaminar*) adaptive filtering, horizontal (*intralaminar*) grouping, top-down *intracortical* (but *interlaminar*) pre-attentive feedback, and top-down *intercortical* (and *interlaminar*) attentive feedback, the LAMINART model shows how some developmental and learning processes can occur without top-down attention, by using *intracortical* feedback processes that computationally realize the same stabilizing effects that top-down *intercortical* attentional processes were earlier predicted to realize. Because of this intimate link between *intracortical* and *intercortical* feedback processes, attention can modulate and thereby enhance pre-attentive groupings as they unfold, as reported in neurophysiological experiments for real contour inputs by Roelfsema et al. (1998) and for illusory contour inputs by Wanning et al. (2011), and simulated using LAMINART by Grossberg and Raizada (2000) for real contours and by Raizada and Grossberg (2001) for illusory contours.

28. Balancing excitation and inhibition: contrast normalization, grouping, and attention

Within the cortical circuits that realize these grouping and attentional processes, there needs to be a *balance* between excitatory and inhibitory interactions (see Eq. (1)). In particular, a balance between excitation and inhibition within layer 2/3 is needed to carry out perceptual grouping. Perceptual groupings can start to form within layer 2/3 due to direct excitatory interactions among cells with long-range horizontal connections, and short-range disynaptic inhibitory interneurons (Fig. 5(c), (e)). The balance between these excitatory and inhibitory inputs helps to ensure that perceptual groupings can form inwardly between pairs or greater numbers of (almost) like-oriented and (almost) co-linear inducers, but not outwardly from a single inducer. This combination of properties is called *bipole grouping* (Grossberg & Mingolla, 1985; Grossberg & Raizada, 2000). In general, a bipole cell in layer 2/3 fires under one of the following circumstances (Fig. 5(c), (e)): it gets direct bottom-up input; or (almost) co-linear *intralaminar* inputs from pairs (or greater numbers) of bipole cell populations with similar orientational tuning on both sides of the cell; or bottom-up input plus input from one or both sides of the cell.

A balance between excitation and inhibition is also required in the on-center of the circuit from layers 6-to-4 so that it can provide excitatory modulation of cell activities in layer 4, but cannot fire them fully. The layers 6-to-4 on-center off-surround circuit helps to do several things: it achieves contrast-normalization of bottom-up inputs (Fig. 5(a), (e)); it helps to choose the strongest grouping in layer 2/3 via the 2/3-to-6-to-4 feedback loop (Fig. 5(c), (e)); and it enables top-down attention, via the folded feedback circuit, from layer 6 in a higher cortical level to layers 6-to-4 in a lower cortical level, to be modulatory at layer 4 (Fig. 5(b), (d), (e)).

29. Balancing excitation and inhibition: the road to synchrony

29.1. Without inputs, balanced circuits spike intermittently

The existence of balanced excitation and inhibition in various kinds of brain circuits makes intuitive sense: If inhibition is too weak, then cells can saturate their activities and excitation can propagate uncontrollably (see Section 6). If inhibition is too strong, then cells cannot get sufficiently activated to process anything. Grossberg and Williamson (2001) modeled how such balanced circuits can develop within the laminar circuits of visual cortex, and showed how the developed circuits could simulate properties of perceptual grouping found in neurobiological data from animals and in psychophysical data from human adult observers. Grossberg and Raizada (2000) and Raizada and Grossberg (2001) went on to simulate more neurobiological data about perceptual grouping, as well as data about attention, using the interaction kernels that developed in the Grossberg and Williamson (2001) model.

The success of these simulations presents us with a paradox, since neural circuits with balanced excitatory and inhibitory connections have also been used to explain the observed sparseness and variability in the number and temporal distribution of spikes emitted by cortical neurons (Shadlen & Newsome, 1998; van Vreeswijk & Sompolinsky, 1998). These spiking patterns are quite inefficient in firing cortical cells. Given the LAMINART model proposal that such variability may reflect mechanisms that are needed to ensure stable development and learning by cortical circuits – that is, “stability implies variability” – the cortex is faced with the difficult problem of how to overcome the inefficiency of variable spiking patterns in driving responses from cortical neurons.

29.2. With inputs, balanced circuits rapidly synchronize: on the cusp of excitability

The LAMINART model shows how these balanced excitatory and inhibitory connections overcome the inefficiency of intermittent spiking when they are driven by inputs. They can then rapidly resynchronize desynchronized signals that belong to the same object, thereby enabling the cortex to process them efficiently. In other words, the process that enables cortical cells to respond selectively to input patterns – namely, balanced excitation and inhibition – allows the cortex to exist at a “cusp of excitability” in the resting state, from which cortical cells can fire vigorously and synchronously in response to input patterns that are selected by cortical bottom-up filtering, horizontal grouping, and top-down attention processes. Rather than think of sparse, intermittent spiking as a problem of inefficiency, it seems more appropriate to view it as one that keeps the cortex ready to fire until driven to do so, while also providing enough activity in the resting state to support activity-dependent homeostatic processes.

30. A new way to compute: digital and binary, feedforward and feedback, analog coherence

As illustrated by the above comments, LAMINART models represent a breakthrough in computing that identifies new principles and processes that embody novel computational properties with revolutionary implications. LAMINART models embody a new type of hybrid between *feedforward* and *feedback* computing, and also between *digital* and *analog* computing for processing distributed data (Grossberg, 2003a). These properties go beyond the types of Bayesian models that are so popular today. They underlie the fast but stable self-organization that is characteristic of cortical development and life-long learning.

30.1. Fast feedforward vs. slower feedback processing of unambiguous vs. ambiguous data

The LAMINART synthesis of feedback and feedback processing can be understood from the following example: When an unambiguous scene is processed, the LAMINART model can quickly group the scene in a fast feedforward sweep of activation that passes directly through layers 4 to 2/3 and then on to layers 4 to 2/3 in subsequent cortical areas (Fig. 5(c), (e)). This property clarifies how recognition can be so fast in response to unambiguous scenes; e.g., Thorpe, Fize, and Marlot (1996). On the other hand, if there are multiple possible groupings in a scene, say in response to a complex textured scene, then competition among these possibilities due to inhibitory interactions in layers 4 and 2/3 (black cells and synapses in Fig. 5) can cause all cell activities to become smaller. This happens because the competitive circuits in the model are *self-normalizing*; that is, they tend to conserve the total activity of the circuit. As noted in Section 6, this self-normalizing property is related to the ability of the shunting on-center off-surround networks that realize the competitive circuits to process input contrasts over a large dynamic range without saturation (Douglas, Koch, Mahowald, Martin, & Suarez, 1995; Grossberg, 1973, 1980; Heeger, 1992).

30.2. Real-time probabilities that run as fast as they can

In other words, these self-normalizing circuits carry out a type of real-time probability theory in which the amplitude of cell activity covaries with the certainty of the network's selection, or decision, about a grouping. Amplitude, in turn, is translated into processing speed and coherence of cell activations. Low activation slows down feedforward processing in the circuit because it takes longer for cell activities to exceed output threshold and to activate subsequent cells above threshold. In the model, network uncertainty is resolved through feedback: Active layer 2/3 grouping cells feed back signals to layers 6-then-4-then-2/3 to close a cortical feedback loop that contrast-enhances and amplifies the winning grouping to a degree and at a rate that reflect the amount of statistical evidence for that grouping. As the winner is selected, and weaker groupings are suppressed, its cells become more active and synchronous, hence can again rapidly send the cortical decision to subsequent processing stages.

In summary, the LAMINART circuit “runs as fast as it can”: it behaves like a real-time probabilistic decision circuit that operates as quickly as possible, given the evidence. It operates in a fast feedforward mode when there is little uncertainty, and automatically switches to a slower feedback mode when there is uncertainty. Feedback selects a winning decision that enables the circuit to speed up again, since activation amplitude, synchronization, and processing speed both increase with certainty.

30.3. Combining the stability of digital with the sensitivity of analog

The LAMINART model also embodies a novel kind of hybrid computing that simultaneously realizes the stability of digital computing and the sensitivity of analog computing. This is true because the feedback loop between layers 2/3-6-4-2/3 that selects or confirms a winning grouping (Fig. 5(c), (e)) has the property of *analog coherence* (Grossberg, 1999; Grossberg, Mingolla, & Ross, 1997; Grossberg & Raizada, 2000); namely, this feedback loop can synchronously choose and store a winning grouping without losing analog sensitivity to amplitude differences in the input pattern. The coherence that is derived from synchronous storage in the feedback loop provides the stability of digital computing – the feedback loop exhibits hysteresis that can preserve the stored pattern against external perturbations – while preserving the sensitivity of analog computation.

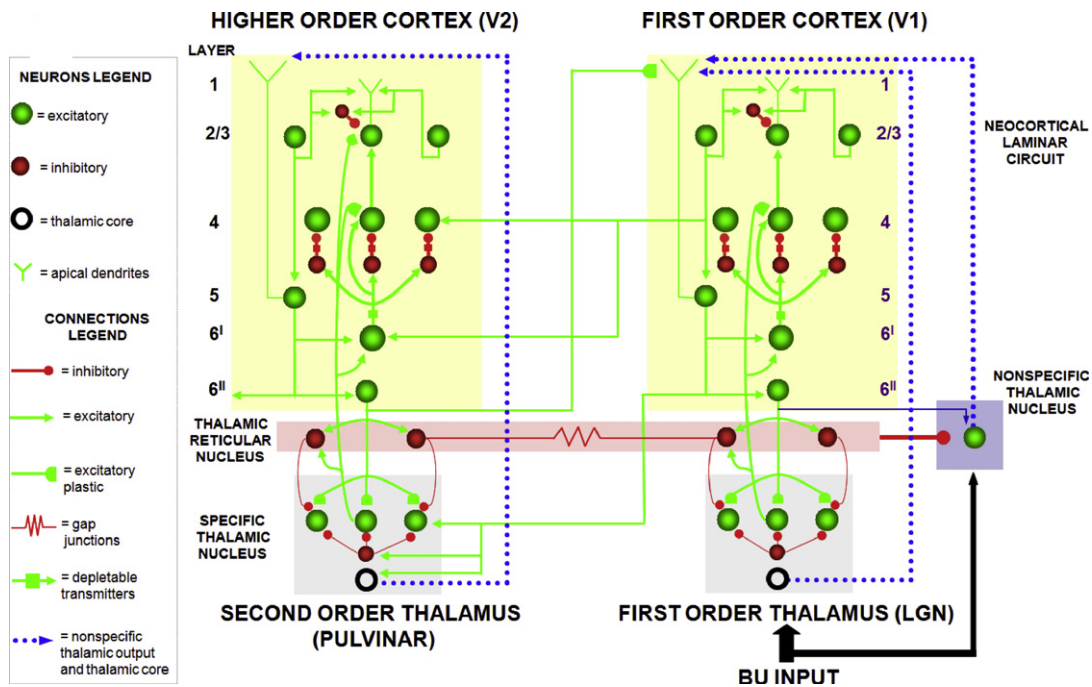


Fig. 6. The SMART model clarifies how laminar neocortical circuits in multiple cortical areas interact with specific and nonspecific thalamic nuclei to regulate learning on multiple organizational levels, ranging from spikes to cognitive dynamics. The thalamus is subdivided into specific first-order and second-order nuclei, nonspecific nucleus, and thalamic reticular nucleus (TRN). The first-order thalamic matrix cells (shown as an open ring) provide nonspecific excitatory priming to layer 1 in response to bottom-up input, priming layer 5 cells and allowing them to respond to layer 2/3 input. This allows layer 5 to close the intracortical loop and activate the pulvinar (PULV). V1 layer 4 receives inputs from two parallel bottom-up thalamocortical pathways: a direct LGN \rightarrow 4 excitatory input, and a $6^I \rightarrow 4$ modulatory on-center, off-surround network that contrast-normalizes the pattern of layer 4 activation via the recurrent $4 \rightarrow 2/3 \rightarrow 5 \rightarrow 6^I \rightarrow 4$ loop. V1 activates the bottom-up V1 \rightarrow V2 corticocortical pathways from V1 layer 2/3 to V2 layers 6^I and 4, as well as the bottom-up corticothalamocortical pathway from V1 layer 5 to the PULV, which projects to V2 layers 6^I and 4. In V2, as in V1, the layer $6^I \rightarrow 4$ pathway provides divisive contrast normalization to V2 layer 4 cells. Corticocortical feedback from V2 layer 6^{II} reaches V1 layer 1, where it activates apical dendrites of layer 5 cells. Layer 5 cells, in turn, activate the modulatory $6^I \rightarrow 4$ pathway in V1, which projects a V1 top-down expectation to the LGN. TRN cells of the two thalamic sectors are linked via gap junctions, which synchronize activation across the two thalamocortical sectors when processing bottom-up stimuli. The nonspecific thalamic nucleus receives convergent bottom-up excitatory input from specific thalamic nuclei and inhibition from the TRN, and projects to layer 1 of the laminar cortical circuit, where it regulates mismatch-activated reset and hypothesis testing in the cortical circuit. Corticocortical feedback connections from layer 6^{II} of the higher cortical area terminate in layer 1 of the lower cortical area, whereas corticothalamic feedback from layer 6^{II} terminates in its specific thalamus and on the TRN. This corticothalamic feedback is matched against bottom-up input in the specific thalamus.

Source: Reprinted with permission from Grossberg and Versace (2008).

31. SMART: spikes, synchrony, and attentive learning in laminar thalamocortical circuits

The Synchronous Matching ART (SMART) model (Fig. 6) predicts how finer details about CLEARs processes may be realized by multiple levels of brain organization. In particular, SMART incorporates spiking dynamics and hierarchical thalamocortical and corticocortical interactions into the LAMINART model. SMART hereby provides a unified functional explanation of single cell properties, such as spiking dynamics, spike-timing-dependent plasticity (STDP), and acetylcholine modulation; hierarchical laminar thalamic and cortical circuit designs and their interactions; aggregate cell recordings, such as current-source densities and local field potentials; and single cell and large-scale inter-areal oscillations in the gamma and beta frequency domains.

SMART embodies attentive category learning requirements into a hierarchy of laminar neocortical circuits interacting with first-order (e.g., LGN) and higher-order (e.g., the pulvinar nucleus; Sherman and Guillery (2001) and Shipp (2003)) specific thalamic nuclei, and nonspecific thalamic nuclei (van Der Werf, Witter, & Groenewegen, 2002). As illustrated in the SMART macrocircuit, corticothalamocortical pathways work in parallel with corticocortical routes (Maunsell & Van Essen, 1983; Salin & Bullier, 1995; Sherman & Guillery, 2002). Specific first-order thalamic nuclei (e.g., LGN) relay sensory information to the cerebral cortex (e.g., V1), whereas specific second-order thalamic nuclei (e.g., pulvinar) receive their main input from lower-order cortical areas (e.g., V1), notably from layer 5, and relay this

information to higher-order cortical areas (Rockland, Andresen, Cowie, & Robinson, 1999; Sherman & Guillery, 2002). Terminations arising from layer 5 are similar to retinogeniculate driving connections, and are often found in more proximal segments of the dendrites. This pattern of connectivity seems to be constant across species (Rouiller & Welker, 2000).

32. Top-down match, gamma oscillations, and STDP in laminar thalamocortical circuits

The SMART model clarifies how a match between cortical and thalamic inputs at the level of specific first-order and higher-order thalamic nuclei might trigger fast stable learning of neural representations in the thalamocortical system. In particular, suppose that, at a specific thalamic nucleus, a sufficiently good match occurs between a bottom-up input pattern and a top-down expectation from layer 6 of its corresponding cortical area. In SMART, such a match can trigger fast synchronized gamma oscillations (γ , 20–70 Hz) in the shunting activation equations (see Eq. (1)), whose short period enables synchronized spikes to drive learning via a spike-timing-dependent plasticity (STDP; Levy and Steward (1983), Markram, Lubke, Frotscher, and Sakmann (1997) and Bi and Poo (2001)) learning rule. STDP is maximal when pre- and post-synaptic cells fire within 10–20 ms of each other, and thus favors learning in match states whose synchronous fast oscillations fall within the temporal constraints of STDP (Traub et al., 1998; Wespatat, Tennykeit, & Singer, 2004). In contrast, mismatched

cells undergo slower beta oscillations (β , 4–20 Hz), whose spikes do not fall within the STDP learning window.

SMART hereby brings the new features of synchronized oscillation frequency and STDP into the discussion of how learning may be selectively regulated by ART dynamics. Aggregate and single-cell recordings from multiple thalamic and cortical levels of mammals have shown high- and low-frequency rhythmic synchronous activity correlated with cognitive, perceptual and behavioral tasks, and large-scale neuronal population models have been proposed to model oscillatory dynamics (Bazhenov, Timofeev, Steriade, & Sejnowski, 1998; Destexhe, Contreras, & Steriade, 1999; Lumer, Edelman, & Tononi, 1997; Siegel, Kording, & König, 2000). However, these models have not linked brain spikes, oscillations, and STDP with the brain states that subserve cognitive information processing.

33. Corticocortical and thalamocortical laminar circuits realize the ART Matching Rule

Fig. 6 summarizes how laminar circuits in the SMART model embody the ART Matching Rule in multiple brain areas, and thereby allow laminar circuits of multiple cortical and thalamic areas to carry out attentive visual learning and information processing. In particular, ART top-down modulatory on-center, off-surround networks occur in both corticocortical and corticothalamic circuits.

Corticocortical top-down attentive matching in SMART refines the explanation by LAMINART of the cortical layers that participate in this process by subdividing layer 6 into sublamina: layer 6^{II} of cortical area V2 sends top-down outputs to cortical area V1 via layer 1, where they activate apical dendrites of layer 5 cells. Layer 5, in turn, activates layer 6^I of V1, which sends modulatory on-center, off-surround signals to layer 4, thereby realizing the ART Matching Rule in V1.

As an example of corticothalamic top-down matching, first consider V1-to-LGN matching: layer 5 cells in V1 activate layer 6^{II}, which sends top-down modulatory on-center, off-surround signals to LGN, thereby realizing the ART Matching Rule in LGN. These pathways help to regulate stable learning between LGN and V1.

As an example of corticothalamic top-down matching in higher-order specific thalamic nuclei, consider how V2 interacts with the pulvinar: A top-down expectation from layer 6^{II} of V2 is matched in the pulvinar against the output pattern from layer 5 of V1, similar to the way in which retinal inputs to the LGN are matched by top-down signals from layer 6^{II} of V1. If a sufficiently good match occurs, then synchronized gamma oscillations can be triggered in the pulvinar and V2, leading to learning of the critical features that are part of the matched pattern.

34. Thalamocortical mismatch, nonspecific thalamic nucleus, and layer 1 mediated reset

If the match is not good enough in visual cortical areas such as V1 and V2, then the nonspecific thalamic nucleus, which is part of the orienting system, gets activated. Mismatch-activated outputs from the nonspecific thalamus carry out arousal and reset functions similar to those summarized in the ART search cycle of Fig. 2. How this works is anatomically more precisely characterized in the SMART circuit diagram of Fig. 6. Nonspecific thalamic activation is broadcast as an arousal signal to many cortical areas via diffuse inputs across layer 1. Apical dendrites in layer 1 of layer 5 cells receive this arousal input. If some of these layer 5 cells are active when the arousal burst occurs, their firing rate is enhanced in response to the arousal input. This enhancement of layer 5 cell firing triggers a selective reset of cortical and thalamic cells in the following way:

Layer 5 cells project to layer 4 via layer 6 (Fig. 6). The signals from layer 6 to 4 are gated by habituating transmitters (see Section 13). Habituation occurs in the pathways associated with layer 4 cells that are active just prior to the arousal burst. When the arousal burst occurs, these previously active cells are disadvantaged relative to cells that were not active. A reset event that is caused by the arousal burst inhibits the previously active cells as it selects new cells with which to better code the novel input, as in the ART memory search cycle of Fig. 2.

35. Slower beta oscillations initiated in lower cortical layers during mismatch reset

As noted above, SMART predicts that thalamocortical mismatches may cause cortical reset via the deeper cortical layers 6 and 4. Model simulations show that such mismatches lead to slower beta oscillations. Putting these two properties together leads to the prediction that the deeper layers of neocortex may express beta oscillations more frequently than the superficial layers. Such a property has been experimentally reported (Buffalo, Fries, Landman, Buschman, & Desimone, 2011); also see Section 38. It remains to test whether the observed experimental property is related to the SMART reset prediction.

35.1. Does varying novelty change the amount of beta?

Two issues may be noted in this regard. One concerns how the prediction may be tested: One possible test would be to carry out a series of experiments on the same animal in which the animal is exposed to environments with progressively more novel events. More novel events should cause more cortical resets. Do more cortical resets per unit time cause more beta oscillations in the lower cortical layers and thereby decrease the ratio of gamma to beta power per unit time?

35.2. Do superficial and deeper layers synchronize and desynchronize during resonance and reset?

The second issue notes that the differences between the oscillation frequencies in the deeper and more superficial cortical layers are averages over time. It is also predicted that interlaminar intracortical feedback loops synchronize all the cortical layers during a match event (Yazdanbakhsh & Grossberg, 2004). Indeed, these are the intracortical feedback loops whereby “a pre-attentive grouping is its own attentional prime”, and thus enable neocortex to develop its circuits, without a loss of stability, even before intercortical attentional circuits can develop. Do the cortical layers desynchronize into superficial-layer gamma and deeper-layer beta during a reset, and then resynchronize into gamma during an attentive resonance?

36. Vigilance control by acetylcholine release from the nucleus basalis of Meynert

As in all ART models, the generality of learned recognition codes in SMART is proposed to be controlled by a vigilance process. SMART predicts how vigilance may be altered by acetylcholine when the nucleus basalis of Meynert is activated via the nonspecific thalamus (Kraus, McGee, Littman, Nicol, & King, 1994; van Der Werf et al., 2002) which, in turn, is activated by corticothalamic mismatches with one or more specific thalamic nuclei (Fig. 6). In general, it is known that cholinergic modulation is an essential ingredient in cortical plasticity (e.g., Kilgard and Merzenich (1998)). Saar, Grossman, and Barkai (2001) have shown, in addition, that ACh release reduces the after-hyperpolarization (AHP) current and increases cell excitability in layer 5 cortical

cells. In SMART, this increased layer 5 excitability due to predictive mismatch may cause reset via the layer 5-to-6¹-to-4 circuit, even in cases where top-down feedback may earlier have sufficiently matched bottom-up input, which is a key property of vigilance control. The increase of ACh might hereby promote search for finer recognition categories in response to environmental feedback, even when bottom-up and top-down signals have a pretty good match in the nonspecific thalamus based on similarity alone.

While ACh is often considered to boost plasticity simply via excitability, recent evidence suggests that fluctuations in cortical ACh are not necessary for simple associative learning, but are necessary for learning mainly when there is featural overlap during perceptual categorization (Atri et al., 2004; Botly & De Rosa, 2007, 2009; Chiba, Bucci, Holland, & Gallagher, 1995; Hata, Kumai, & Okaichi, 2007; Winters, Bartko, Saksida, & Bussey, 2007), consistent with the idea that it can increase vigilance to achieve better categorical separation. Further support for this idea comes from lesions in rats of the nucleus basalis of Meynert which have little impact on learning rate, except when there is a high degree of interference between the categories to be learned; that is, when categories share the same features in a certain dimension (Botly & De Rosa, 2007, 2009). Similarly, studies in humans show that scopolamine, by competitively binding muscarinic receptors, diminishes learning of overlapping word pairs more than non-overlapping pairs (Atri et al., 2004). Associative learning studies in rats with combinations of light and tone has shown that the concentration of released ACh increases more during discrimination learning experiments in which an individual stimulus (A; e.g., light) signals reward and a compound stimulus (AB; e.g., light + tone) signals no reward, than during elemental discrimination, in which one stimulus (A; e.g., light) signals reward and another stimulus (B; e.g., tone) signals no reward (Hata et al., 2007). Finally, donepezil, which increases cortical ACh by inhibiting its degradation by Acetylcholinesterase (AChE), has been shown by fMRI to reduce the expanse of response in V1 from a pulsating visual stimulus (Silver, Shenav, & D’Esposito, 2008).

Taken together, these data suggest that increased ACh (and attention) refines perceptual representations by adding specificity. Palma, Grossberg, and Versace (submitted for publication) have carried out simulations of recurrent on-center off-surround networks composed of spiking shunting neurons to illustrate how ACh may modulate the transformation and STM storage of input patterns in a manner compatible with vigilance control.

37. Vigilance diseases may include autism and medial temporal amnesia

ART has been used to explain how symptoms of some mental disorders may be due to vigilance being stuck at either too high or too low values.

37.1. High vigilance and hyperspecific category learning in autism

High vigilance has been predicted to cause symptoms of hyperspecific category learning and attentional deficits in some autistic individuals (Grossberg & Seidman, 2006). Psychophysical experiments have been done to test this prediction in high-functioning autistic individuals (Church et al., 2010; Vladusich, Lafe, Kim, Tager-Flusberg, & Grossberg, 2010). It is also known that there is abnormal cholinergic activity in the parietal and frontal cortices of autistic individuals that is correlated with abnormalities in the nucleus basalis (Perry et al., 2001), consistent with the predicted role of the nucleus basalis and ACh in regulating vigilance.

37.2. Low vigilance without memory search during medial temporal amnesia

Low vigilance has been predicted in individuals with medial temporal amnesia. A hippocampal lesion removes the orienting system during cortico-hippocampal interactions (Fig. 2). By preventing memory search, the ability to learn new categories is degraded. Such a lesion, in effect, keeps vigilance equal to zero, and any learning that can occur without mismatch-mediated reset and memory search can only form very general categories (Carpenter & Grossberg, 1993). Relevant data from amnesic individuals have been reported by Knowlton and Squire (1993), who showed that amnesic subjects and normal subjects perform equally well on easy categorization tasks, but the performance of amnesic subjects drops significantly for more demanding tasks.

Knowlton and Squire (1993) posited that two separate memory systems are needed to explain these data. In contrast, Zaki, Nosofsky, Jessup, and Unversagt (2003) quantitatively fit these data with a single exemplar-based model whose sensitivity parameter was chosen lower for amnesic than for normal subjects. This exemplar model, which is usually expressed in terms of formal algebraic equations, may be interpreted as a real-time dynamic process undergoing only locally defined interactions. Such an interpretation shows that the model implicitly posits prototypes and top-down processes akin to ART. A low sensitivity parameter c in this exemplar model (see their Eq. (4)) plays a role similar to that played by a low vigilance parameter ρ in an ART model (Amis, Carpenter, Ersoy, & Grossberg, submitted for publication).

38. Gamma and beta oscillations during attentive resonance and mismatch reset

38.1. Shared gamma/beta oscillatory dynamics in visual cortex, frontal eye fields, and hippocampus

As noted in Section 32, one of the SMART predictions concerns how brain oscillations may contribute to learning. SMART demonstrates through computer simulations that sufficiently good top-down *matches* may cause fast gamma oscillations that support attention, resonance, learning, and consciousness, whereas sufficiently bad *mismatches* inhibit learning by causing slower beta oscillations while triggering attentional reset and hypothesis testing that are predicted to be initiated in the deeper cortical layers. The predicted use of ART Matching Rule top-down circuits across all brain systems that solve the stability–plasticity dilemma suggests that this gamma/beta dichotomy may occur in multiple brain systems. At least three kinds of recent data currently support this prediction:

(1) Buffalo et al. (2011) have reported more gamma oscillations in the superficial layers of visual cortex and more beta oscillations in deeper layers of the cortex.

(2) Buschman and Miller (2009) have reported beta oscillations during spatial attention shifts in the frontal eye fields of monkeys. In fact, this discovery was made after Earl Miller was told the SMART prediction. These investigators then reanalyzed their spatial attention data, aligning them in time with respect to the attention shifts, and thereby found the underlying beta oscillations.

(3) Berke, Hetrick, Breck, and Green (2008) have reported beta oscillations during hippocampal place cell learning in novel environments.

The Berke et al. (2008) results nicely illustrate how beta oscillations may be related to ART category learning. Place cells in the hippocampus fire when an animal or human is in a particular location, or “place”, in a spatial environment. Place cells hereby play a critical role in spatial navigation. Place cell receptive field selectivity can develop as an animal navigates within seconds

to minutes, and can remain stable for months (Frank, Stanley, & Brown, 2004; Muller, 1996; Thompson & Best, 1990; Wilson & McNaughton, 1993). Place cell learning thus seems to solve the *stability–plasticity dilemma*.

39. Inverted-*U* in beta power through time during hippocampal place cell learning

Are place cells learned using ART dynamics? The Berke et al. (2008) data are consistent with this hypothesis. They showed that, paradoxically, beta power was very low as a mouse traversed a lap for the first time in a novel environment, grew to full strength on the second and third laps, became low again after two minutes of exploration, and remained low on subsequent days. Beta oscillation power also correlated with the rate at which place cells became spatially selective, and did not correlate with theta oscillations. Given the rapidity with which place cell learning occurred, and the sharp increase in beta activity during the second exposure to the environment, it would seem that a highly selective learning mechanism is at work.

These data can be explained as follows (Grossberg, 2009a): In any ART system, the top-down adaptive weights that represent learned expectations need to be broadly distributed before learning occurs, so that they can match whatever input pattern first initiates learning of a new category (Carpenter & Grossberg, 1987). Indeed, when a new category is first activated, it is not known at the category level what pattern of features caused the category to be activated. *Whatever* feature pattern was active needs to be matched by the top-down expectation on the first learning trial, so that resonance and weight learning can begin. Hence the need for the initial values of top-down weights to be sufficiently large and broadly distributed to match any feature pattern. The low beta power on the first lap of exploration can be explained by the initial top-down match.

Given that top-down weights are initially broadly distributed, the learning of top-down expectations is a process of *pruning* weights on subsequent learning trials, and uses mismatch-based reset events to discover categories capable of representing the environment. Beta power on subsequent laps can be explained by mismatch reset events that correlate with the rate at which place cells become spatially selective. After learning stabilizes, there are no more mismatches, so beta power subsides.

Such an inverted-*U* in beta power through time is thus a signature of ART category learning in any environment.

40. Entorhinal grid cell and hippocampal place cell learning as an ART system

Is there additional evidence that hippocampal place cells are learned as part of an ART system? In order to understand how such learning happens, we need to know what sorts of inputs activate place cells. Hafting, Fyhn, Molden, Moser, and Moser (2005) made a remarkable discovery that has greatly enhanced the understanding of how place cells may form. Grid cells in the superficial layers of medial entorhinal cortex (MEC) provide inputs to hippocampal place cells. Grid cells are so called because each of them, unlike a place cell, fires at multiple spatial positions that form a regular hexagonal grid during navigation in an open field. Grid cells also exhibit a gradient of spatial scales along the dorsoventral axis of the MEC, with anatomically neighboring cells sharing similar grid spacings and orientations but having different spatial phases that are not topographically organized.

40.1. Grid and place cells are learned from most frequent co-activations in self-organizing maps

The GridPlaceMap model (Pilly & Grossberg, *in press*) has shown how a hierarchy of self-organizing maps, each obeying the same laws, can respond to realistic rat trajectories by learning grid cells with hexagonal grid firing fields of multiple spatial scales, and place cells with unimodal firing fields (see also Gorchetchnikov and Grossberg (2007) and Mhatre, Gorchetchnikov, and Grossberg (2012)). These learned receptive fields fit neurophysiological data about grid and place cells and their development in juvenile rats (Langston et al., 2010; Wills, Cacucci, Burgess, & O'Keefe, 2010). Due to the inputs from multiple scales of developing grid cells, the learned hippocampal place fields can represent much larger spaces than the grid cells, indeed spaces whose spatial scale may be the least common multiple of grid cell scales. These least common multiple scales are large enough to let place cells support navigational behaviors. Despite their difference appearances, both grid cell and place cell receptive fields are learned by the same self-organizing map equations in the GridPlaceMap model. At both the grid and place cell levels, the self-organizing maps amplify and learn to categorize the most energetic and frequent co-activations of their inputs.

In this conception, place cells are spatial category cells that are activated by multiple scales of entorhinal grid cells in a self-organizing map. However, it is known that self-organizing maps cannot solve the stability–plasticity dilemma (Section 1). Indeed, when Grossberg (1976a, 1978a) introduced the modern laws for competitive learning and self-organizing maps, he proved that they can learn well in response to sparse input environments, indeed with Bayesian properties, but that they exhibit catastrophic forgetting in response to dense non-stationary environments. Grossberg (1976b) introduced ART as an enhanced model capable of dynamically stabilizing category learning in a self-organizing map that can be learned from arbitrary environments.

40.2. Stabilizing grid and place cell learning using top-down attentive ART feedback

If, in fact, grid and place cell learning occur in self-organizing maps, then, as in all self-organizing map models, one expects that grid cell and place cell learning are dynamically stabilized by ART top-down attentive matching mechanisms. Such matching is already indirectly supported by the Berke et al. (2008) data on how beta power changes through time when place cells are learned in a novel environment. Are there anatomical and neurophysiological data that more directly support this hypothesis?

The anatomy of the hippocampal system supports the possibility that such attentive feedback exists, since feedback pathways exist from the hippocampal CA1 region to the entorhinal cortex. Neurophysiological data also support the predicted role of attention in hippocampal learning. For example, Kentros, Agnietri, Streater, Hawkins, and Kandel (2004) showed that “conditions that maximize place field stability greatly increase orientation to novel cues. This suggests that storage and retrieval of place cells is modulated by a top-down cognitive process resembling attention and that place cells are neural correlates of spatial memory” (p. 283). In like manner, it has been proposed that learning of place cell receptive fields reflects an “automatic recording of attended experience” (Morris & Frey, 1997, p. 1489). These experiments clarify that cognitive processes like attention play a role in hippocampal learning and memory stability, and interact with NMDA receptors to mediate long-lasting hippocampal place field memory in novel environments (Kentros et al., 1998). Thus, the learning of entorhinal grid cells and hippocampal place cells may be viewed as part of a specialized ART system for learning spatial categories as an animal or human navigates in its environment.

41. The ART of cognitive–emotional interactions: the feeling of what happens

41.1. Cognitive–emotional resonances link knowing to feeling and motivation

ART recognition categories can be activated when objects are experienced but, without further processing, the amount of attention that they attract does not reflect the emotional value of these objects through time. The first adaptive resonance to be discovered was, in fact, a cognitive–emotional resonance that links cognitive representations to the emotional representations that select and maintain motivated attention upon them as valued actions are carried out (Grossberg, 1975). This resonance was discovered during an analysis of how attentional blocking and unblocking may occur during learning (Grossberg, 1975, 1982, 1984b; Pavlov, 1927). How such a cognitive–emotional resonance may arise is outlined below as part of the CogEM, or Cognitive–Emotional–Motor model (Dranijs et al., 2008; Grossberg, 1971, 1972a, 1972b, 1982, 1984b; Grossberg, Bullock et al., 2008), that combines cognitive and emotional mechanisms within ART. Such a cognitive–emotional resonance clarifies how cognitive and emotional constraints may be harmonized during decision-making. When this resonance was first described, cognitive science and Artificial Intelligence focused almost entirely on propositional calculations and affective neuroscience did not exist. Today, such a cognitive–emotional resonance may be used to clarify how cognitive science and affective neuroscience may be unified to form a more comprehensive theory.

41.2. The feeling of what happens arises from a CogEM resonance: core consciousness

In *The Feeling of What Happens*, Damasio (1999, p. 178, Fig. 6.1) derived from his clinical data a heuristic circuit that is very similar to rigorously defined CogEM circuits that are described in Section 42. Damasio used his circuit to explain how “core consciousness” arises. In his schematic model, the first sensory stage is called the “map of object x ” and the second sensory stage is called the “second-order map”. The CogEM drive representation is called the “proto-self”. As in CogEM, conjoint inputs from the “map of object” and “proto-self” activate the “second-order map” which, in turn, attentionally enhances the “map of object” via top-down feedback. Damasio also noted that these structures combine processes of homeostasis, emotion, attention, and learning (see pp. 272–273) that the CogEM model had proposed twenty years before. The Damasio (1999) discussions of how core consciousness may be altered in clinical patients may be translated into how the cognitive–emotional resonances that support normal behaviors in the CogEM model are altered in various mental disorders. CogEM may also be used to explain symptoms of mental disorders such as schizophrenia (Grossberg, 1984a, 2000c) and autism (Grossberg & Seidman, 2006). Indeed, autistic individuals may exhibit emotional symptoms in addition to cognitive symptoms such as hyperspecific categorization (Section 37).

42. CogEM model: unifying recognition, emotion, motivated attention, and action

42.1. The amygdala is a drive representation

Because of its spatially compact representation, a view- and spatially-invariant object category can be associated through reinforcement learning with one or more drive representations, which are brain regions, such as the amygdala and hypothalamus




Fig. 7. The START model combines the reinforcement learning and motivated attention and action processes of the Cognitive–Emotional–Motor (CogEM) model with the adaptively timed modulation by Spectral Timing of ART recognition learning and search. Adaptively timed learning maintains motivated attention (pathway $D \rightarrow S_i^{(2)} \rightarrow S_i^{(1)} \rightarrow D$) while it inhibits activation of the orienting system (pathway $D \rightarrow A$). See text for details.
Source: Reprinted with permission from Grossberg and Merrill (1992).

(Aggleton, 1993; LeDoux, 1993), that represent internal drive states and emotions. Activation of a drive representation by an invariant object category can, in turn, trigger emotional reactions and motivational feedback that can amplify the activity of the valued category, thereby drawing motivated attention to it, facilitating its competition with other possible categories for selection, and influencing its ability to control behavioral outputs. Because a single invariant object category can be activated by a wide range of input exemplars, all these input exemplars can, via the invariant category, trigger the choice and release of actions that realize valued goals in a context-sensitive way.

42.2. Conditioned reinforcer and incentive motivational learning

In Fig. 7, visually perceived objects are called conditioned stimuli (CS_i ; see I_i in Fig. 7). The invariant object categories that they activate are called sensory representations (S_{CS_i} ; denoted $S_i^{(1)}$ in Fig. 7) which, in turn, activate drive representations (D). Reinforcement learning can convert an event or object (say CS_1) that activates an invariant object category ($S_{CS_1}^{(1)}$) into a *conditioned reinforcer* by strengthening associative links in the inferotemporal-to-amygdala pathways from the category to the drive representation (D). In addition to its amygdala projections, the invariant object category can also send excitatory projections to regions of prefrontal cortex ($S_{CS_2}^{(1)}$), such as orbitofrontal cortex. The amygdala (D) also sends projections to orbitofrontal cortex (Barbas, 1995; Grossberg, 1975, 1982). When these orbitofrontal cells receive converging inferotemporal and

amygdala inputs during reinforcement learning, the corresponding amygdala-to-orbitofrontal pathways can be selectively strengthened. This kind of conditioned input from the amygdala is said to provide *incentive motivation* to the orbitofrontal representations. Both conditioned reinforcer learning and incentive motivational learning obey gated steepest descent learning laws (Eqs. (3) and (4)).

42.3. Orbitofrontal object–value cells

Orbitofrontal representations fire most vigorously when they receive convergent inputs from inferotemporal categories and amygdala incentive motivation (Baxter, Parker, Lindner, Izquierdo, & Murray, 2000; Schoenbaum, Setlow, Saddoris, & Gallagher, 2003), that is, they are object–value cells. The incentive outputs from the amygdala are sensitive to the drive state of the individual, and become desensitized when the corresponding drive is satisfied. In this way, an object that represents a valued goal object can vigorously activate its orbitofrontal representation through the combination of direct inferotemporal-to-orbitofrontal connections and indirect, value-sensitive inferotemporal-to-amygdala-to-orbitofrontal connections. The latter connections withdraw their support of orbitofrontal firing when the corresponding amygdala drive state is satisfied.

42.4. Cognitive–emotional inferotemporal–amygdala–orbitofrontal resonances

Orbitofrontal cells ($S_{CS_2}^{(1)}$), in turn, send top-down feedback to sensory cortex ($S_{CS_1}^{(1)}$) to enhance sensory representations that are motivationally salient (Fig. 7). Competition among inferotemporal categories chooses those with the best combination of sensory and motivational support. An inferotemporal–amygdala–orbitofrontal feedback loop is then activated which supports a cognitive–emotional resonance that supports learning of conditioned reinforcer and incentive motivational adaptive weights, core consciousness of goals and feelings (Damasio, 1999; Grossberg, 1975, 2000a), and releases learned action commands from prefrontal cortex via downstream circuitry ($S_{CS_2}^{(1)} \rightarrow M$) to achieve valued goals. The interactions that are schematically summarized above constitute the main processes of the CogEM, or Cognitive–Emotional–Motor, model. CogEM was progressively developed to functionally explain and predict behavioral and neurobiological data about these processes with increasing precision since its introduction in Grossberg (1972a, 1972b, 1975, 1982). In particular, CogEM top-down prefrontal-to-sensory cortex feedback was the first example to be discovered of ART top-down attentive matching, one that clarifies data about attentional blocking and unblocking (Grossberg, 1975; Grossberg & Levine, 1987; Kamin, 1969; Pavlov, 1927). When this CogEM circuit functions improperly, symptoms of various mental disorders result. For example, hypoactivity of the amygdala or orbitofrontal cortex can cause failures in Theory of Mind processes that may occur in both autism and schizophrenia (Grossberg, 2000c; Grossberg & Seidman, 2006). In addition, when top-down predictive matching processes become driving rather than modulatory, say due to abnormal tonic basal ganglia activity, then properties similar to schizophrenic hallucinations may emerge (Grossberg, 2000a).

43. MOTIVATOR: complementary roles of amygdala and basal ganglia regulate resonance

43.1. Complementary roles of amygdala and basal ganglia in reinforcement learning and action

The MOTIVATOR model (Dranijs et al., 2008; Grossberg, Bullock et al., 2008) further develops the CogEM model, just as SMART further develops ART; see Fig. 8. The MOTIVATOR model unifies

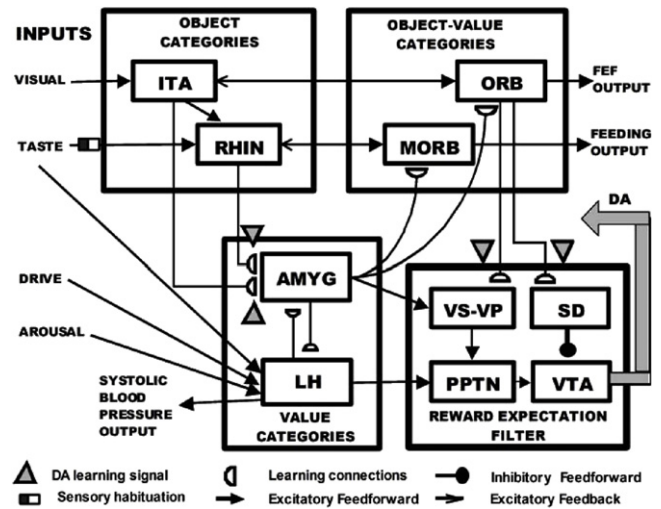


Fig. 8. The MOTIVATOR model: brain areas in the MOTIVATOR circuit can be divided into four regions that process information about conditioned stimuli (CSs) and unconditioned stimuli (USs): Object Categories represent visual or gustatory inputs, in anterior inferotemporal (ITA) and rhinal (RHIN) cortices. Value Categories represent the value of anticipated outcomes on the basis of hunger and satiety inputs, in amygdala (AMYG) and lateral hypothalamus (LH). Object-Value Categories resolve the value of competing perceptual stimuli in medial (MORB) and lateral (ORB) orbitofrontal cortex. The Reward Expectation Filter involves basal ganglia circuitry that responds to unexpected rewards. Source: Reprinted with permission from Dranijs et al. (2008).

the CogEM and TELOS models (Section 44; Brown, Bullock, & Grossberg, 1999; Brown et al., 2004). MOTIVATOR includes the cognitive–emotional inferotemporal–amygdala–orbitofrontal resonances of the CogEM model, but also incorporates mechanisms from TELOS that model how the basal ganglia interact with these brain regions. The roles of the amygdala and basal ganglia in processing valued expected and unexpected events are often complementary.

TELOS models how the dopaminergic cells of the substantia nigra pars compacta (SNc) of the basal ganglia respond to unexpected rewards (Brown et al., 1999; Schultz, 1998), thereby regulating associative learning in neuronal connections that support perceptual, cognitive, and cognitive–emotional resonances and the actions that they control. TELOS also proposes how a different part of the basal ganglia (e.g., the substantia nigra pars reticulata, or SNr), gates selection and release of these learned plans and actions. In this way, the basal ganglia work together with the amygdala to provide motivational support, focus attention, and release contextually appropriate actions to achieve valued goals.

43.2. Amygdala and basal ganglia responses to expected vs. unexpected events

In particular, the amygdala contains *value categories* that are learned in response to autonomic hypothalamic input patterns (Fig. 8). These amygdala categories interact with sensory and prefrontal, notably orbitofrontal, cortex during cognitive–emotional resonances to maintain motivated attention on the prefrontal categories that control actions aimed at acquiring valued goals that are *expected* in a given environment. In contrast, cells in the substantia nigra pars compacta (SNc) of the basal ganglia generate dopamine bursts or dips in response to *unexpected* rewarding signals (Brown et al., 1999; Hollerman & Schultz, 1998; Ljungberg, Apicella, & Schultz, 1992; Mirenovic & Schultz, 1994; Schultz, 1998; Schultz, Apicella, & Ljungberg, 1993; Schultz et al., 1995).

These bursts and dips generate widespread, or nonspecific, activations of dopaminergic inputs to many brain regions, where they act as Now Print signals to facilitate learning or unlearning, respectively, of currently active associative linkages. As this learning leads to expected consequences, the Now Print signals ebb away, leaving strong learned connections, such as the conditioned reinforcer and incentive motivation pathways among sensory cortices, amygdala, and orbitofrontal cortex (Figs. 7 and 8), to direct sustained motivated attention and action to achieve expected goals.

44. TELOS balances reactive and planned behaviors: frontal-parietal resonance opens gates

The basal ganglia also play a role during the processing of expected events. This function helps the brain to balance between reactive and planned behaviors. The complementary orienting vs. attentional systems in ART interact to discover and learn new recognition codes. They are subsumed by a larger complementary brain system that balances reactive vs. planned behaviors.

44.1. How does the brain know before it knows: gating reactive behaviors during plan selection

Rapid reactive movements are needed to ensure survival in response to unexpected dangers. Planned movements, that involve focused attention, often take longer to select and release. How does the brain prevent reactive movements from being triggered prematurely in situations where a more slowly occurring planned movement would be more adaptive? Movement *gates* can prevent the reactive movement from being launched until the planned movement can effectively compete with it. All movement gates tonically inhibit movement commands. When a specific gate is inhibited, the cells that control the corresponding movement command can be activated. Thus, the brain needs to keep the movement gate active that would otherwise be inhibited by a reactive cue until an appropriate plan can be chosen. Then the winning planned movement command can open its gate and launch its movement. The substantia nigra pars reticulata (SNr) regulates this sort of gating process.

These gates need to cope with the following challenging problem: When a cue occurs, the fastest response would be an orienting response to look at it. For this to happen, the cue needs to open the appropriate basal ganglia gate to enable the reactive movement to occur. However, if the cue is a discriminative cue to do a different action, especially an action that requires rapid execution, then the reactive response is not adaptive. However, as noted above, it may take longer to fully process the cue to determine the adaptive conditional response that it would activate the reactive response. How does the brain know that a plan is being elaborated, even before it is chosen, so that the reactive gate can be kept shut? How does the brain know before it knows? The brain must prevent a reactive movement command from opening its gate before a planned movement command is ready to open a different gate, yet also allow a reactive movement command to open its gate as rapidly as possible when no planned movement command is being formed.

Brown et al. (2004) developed the TELOS model to explain and simulate how the brain may achieve this sort of balance between reactive and planned movements as it controls the learning and performance of saccadic eye movements (Fig. 9). The name TELOS is from the ancient Greek *telos* for goal, end, or completion of a plan, but is also an acronym for the model's full name of TElencephalic Laminar Objective Selector. Similar circuits would be expected in the control of other movements as well.

44.2. Frontal-parietal resonance marks plan choice and leads to planned gate opening

According to TELOS, the brain “knows before it knows” in the following way: The model predicts how the distribution of excitation and inhibition that converges on the basal ganglia when a plan is being elaborated keeps the reactive gate closed. When a movement plan is finally chosen, there is agreement between cells in the frontal eye fields (FEF) and the parietal cortex representation of target position. This agreement changes the excitatory-inhibitory balance and is expressed by a synchronous FEF-parietal resonance. This resonance is predicted to signal attentive consistency between a finally selected movement plan and the location of the corresponding target location. When this happens, the balance of excitation and inhibition enables the appropriate basal ganglia movement gate to open and release the context-appropriate action. Buschman and Miller (2007) have reported such prefrontal-parietal resonances during movement control, and Pasupathy and Miller (2004) have reported that the different time courses of activity in the prefrontal cortex and basal ganglia are consistent with how basal ganglia-mediated gating of prefrontal cortex may be learned.

In further support of this proposal, TELOS model simulations emulate how SNc dopaminergic reward and non-reward signals guide monkeys to learn and perform saccadic eye movements in the fixation, single saccade, overlap, gap, and delay (memory-guided) saccade tasks. After learning occurs, model cell activation dynamics quantitatively simulate, and predict functional roles for, the dynamics of seventeen types of identified neurons during performance of these tasks.

45. How value categories interact with the subjective value of objects

The MOTIVATOR model explains and simulates how cognitive-emotional resonances may occur between higher-order sensory cortices, such as inferotemporal and rhinal cortices, and an evaluative neuraxis composed of the hypothalamus, amygdala, basal ganglia, and orbitofrontal cortex. Given a conditioned stimulus (CS), the model amygdala and lateral hypothalamus interact to calculate the expected current value of the subjective outcome that the CS predicts, constrained by the current state of deprivation or satiation. The amygdala relays the expected value information to orbitofrontal cells that receive inputs from anterior inferotemporal cells, and to medial orbitofrontal cells that receive inputs from rhinal cortex. The activations of these orbitofrontal cells during a cognitive-emotional resonance code the subjective values of objects. These values guide behavioral choices.

The model basal ganglia detect errors in CS-specific predictions of the value and timing of rewards. Excitatory inputs from the pedunclopontine nucleus interact with timed inhibitory inputs from model striosomes in the ventral striatum to regulate dopamine burst and dip responses from cells in the substantia nigra pars compacta and ventral tegmental area. As noted in Section 43, learning and unlearning in cortical and striatal regions is modulated by the widespread broadcast of these dopaminergic bursts and dips.

The MOTIVATOR model has, to the present, been used to simulate tasks that examine food-specific satiety, Pavlovian conditioning, reinforcer devaluation, and simultaneous visual discrimination. Model simulations successfully reproduce neurophysiological properties of known cell types, including signals that predict saccadic reaction times and CS-dependent changes in systolic blood pressure. The model hereby illustrates how cognitive processes in prefrontal cortex can influence both bodily processes such as blood pressure, and actions such as eye movements towards a valued goal object.

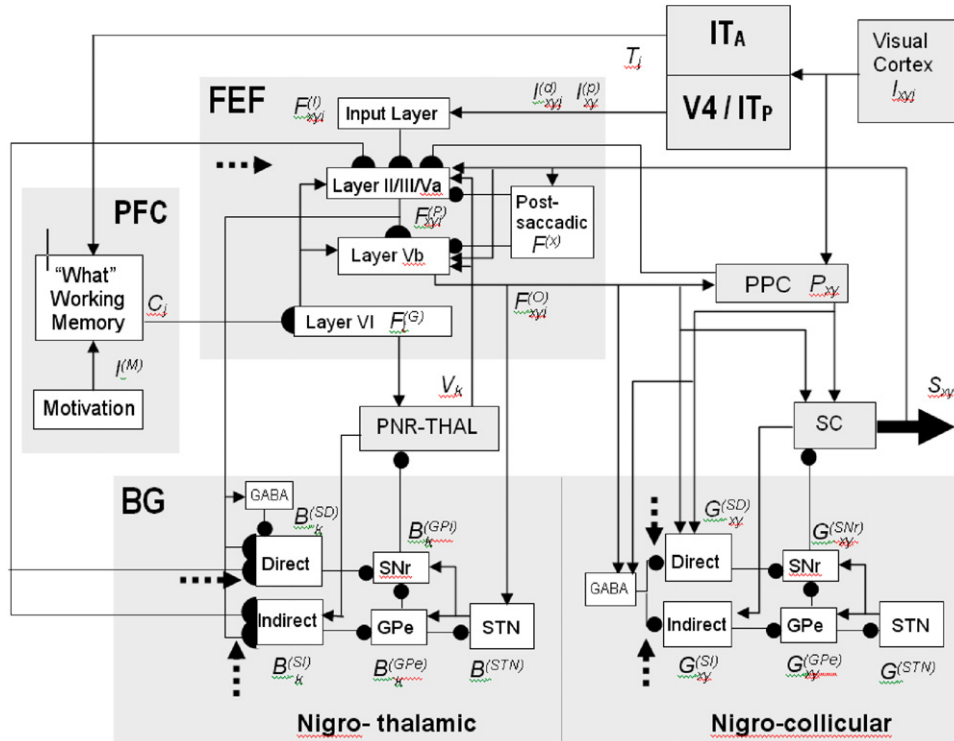


Fig. 9. The TELOS model proposes how the basal ganglia interacts with the laminar circuits of the frontal eye fields (FEF) and the superior colliculus (SC) to learn, plan, and control saccadic eye movements: separate gray-shaded blocks highlight the major anatomical regions whose roles in planned and reactive saccade generation are treated in the model. Excitatory links are shown as arrowheads and inhibitory links as ballheads. Filled semi-circles terminate cortico-striatal and cortico-cortical pathways modeled as subject to learning, which is modulated by reinforcement-related dopaminergic signals (dashed arrows). In the FEF block, Roman numerals I–VI label cortical layers; V_a and V_b , respectively, are superficial and deep layers V . Subscripts xy index retinotopic coordinates, whereas subscript i denotes an FEF zone gated by an associated BG channel. All variables for FEF activities use the symbol F . Processed visual inputs $I_{xyi}^{(c)}$ and $I_{xyi}^{(d)}$ emerging from visual areas including V4 and posterior IT feed into the model FEF input cells and affect activations $F_{xyi}^{(l)}$. Such inputs are predicted to synapse on cells in layer III (and possibly layers II and IV). Visual input also excites the PPC, P_{xy} , and anterior IT, T_k . A PFC motivational signal $I^{(M)}$ arouses PFC working memory activity C_i , which in turn provides a top-down arousal signal to model FEF layer VI cells, with activities $F_i^{(G)}$. The FEF input cell activities $F_{xyi}^{(l)}$ excite FEF planning cells $F_{xyi}^{(p)}$, which are predicted to reside in layers III/Va (and possibly layer II). Distinct plan layer activities represent alternative potential motor responses to input signals, e.g. a saccade to an eccentric target or to a central fixation point. FEF layer VI activities $F_i^{(G)}$ excite the groups/categories of plans associated with gatable cortical zones i and associated thalamic zones k . The BG decide which plan to execute and send a disinhibitory gating signal that allows thalamic activation V_k , which excites FEF layer Vb output cell activities $F_{xyi}^{(o)}$ to execute the plan. The model distinguishes (Kemel et al., 1988) a thalamus-controlling BG pathway, whose variables are symbolized by B , and a colliculus-controlling pathway, whose variables are symbolized by G . Thus, the striatal direct (SD) pathway activities $B_k^{(SD)}$ and $G_{xy}^{(SD)}$ respectively inhibit GPI activities $B_k^{(GPI)}$ and SNr activities $G_{xy}^{(SNr)}$, which respectively inhibit thalamic activities V_k and collicular activities S_{xy} . If the FEF saccade plan matches the most salient sensory input to the PPC, then the basal ganglia disinhibit the SC to open the gate and generate the saccade. However, if there is conflict between the bottom-up input to PPC and the top-down planned saccade from FEF, then the BG-SC gate is held shut by feedforward striatal inhibition (note BG blocks labeled GABA) until the cortical competition resolves. When a plan is chosen, the resulting saccade-related FEF output signal $F_{xyi}^{(o)}$ activates PPC, the STN and the SC (S_{xy}). The SC excites FEF postsaccadic cell activities $F_{xy}^{(x)}$, which delete the executed FEF plan activity. The STN activation helps prevent premature interruption of plan execution by a subsequent plan or by stimuli engendered by the early part of movement. Source: Reprinted with permission from Brown et al. (2004).

46. Adaptively timed resonances: distinguishing expected vs. unexpected disconfirmations

46.1. Balancing consummatory vs. exploratory behaviors

The cognitive-emotional resonances that support reinforcement learning must be adaptively timed, since rewards are often delayed in time relative to actions aimed at acquiring them. On the one hand, if an animal or human could not inhibit its exploratory behavior, then it could starve to death by restlessly moving from place to place, unable to remain in one place long enough to obtain delayed rewards there, such as food. On the other hand, if an animal inhibited its exploratory behavior for too long while waiting for an expected reward, such as food, then it could starve to death if food was not forthcoming. Thus, the survival of a human or animal may depend on its ability to accurately time the delay of a goal object based upon its previous experiences in a given situation. Being able to predict *when* desired consequences occur is often as important as predicting *that* they will occur. Indeed, to control pre-

dictive action, the brain needs to coordinate the What, Why, When, Where, and How of desired consequences by combining recognition learning, reinforcement learning, adaptively timed learning, spatial learning, and sensory-motor learning, respectively.

Adaptive timing, just like amygdala and basal ganglia interactions, requires balancing between reactive and planned behaviors; notably, between *exploratory* behaviors, which may discover novel sources of reward, and *consummatory* behaviors, which may acquire expected sources of reward. A human or animal needs to suppress exploratory behavior and focus attention upon an expected source of reward when the reward is expected. The Spectral Timing model (Brown et al., 1999; Fiala, Grossberg, & Bullock, 1996; Grossberg & Merrill, 1992, 1996; Grossberg & Schmajuk, 1989) accomplishes this by predicting how the brain distinguishes *expected non-occurrences*, or *disconfirmations*, of rewards, which should not interfere with acquiring the delayed goal, from *unexpected non-occurrences*, or *disconfirmations*, of rewards, which can trigger consequences of predictive failure, including reset of working

memory, attention shifts, emotional rebounds, and exploratory behaviors.

46.2. Distinguishing expected vs. unexpected disconfirmations

The following example illustrates this distinction: Suppose that a pigeon in a Skinner box is learning to press a lever for food on a variable-interval reinforcement schedule. Suppose that the pigeon is rewarded, on average, ten seconds after it presses the lever. If the pigeon presses the lever and looks immediately at the place where the food reward will be delivered, it will see no food. Why does the pigeon not treat the non-occurrence of reward as a predictive failure, and thereby become frustrated, shift attention to other activities, unlearn the cognitive–emotional connections that triggered the action, and explore other parts of the Skinner box, rather than remain ten seconds near the food source until the food appears? If the pigeon always treated delayed rewards as predictive failures, then it could only realize immediate gratifications, and could not learn to adaptively time its behaviors to optimize its ability to earn rewards. Somehow the pigeon learns to discount the fact that reward does not occur before ten seconds elapse. This is an example of an *expected* non-occurrence, or disconfirmation, of reward. On the other hand, suppose that the reward is expected after ten seconds and never occurs. This is an *unexpected* non-occurrence, or disconfirmation, of reward, and should be treated as a predictive failure that can free the pigeon from perseverating on a behavior that no longer leads to reward.

If in this situation, reward did come before ten seconds elapsed, then the pigeon could see it and eat it. Thus, adaptive timing does not interfere with the ability to perceptually and cognitively process sensory inputs. Rather, it modulates how mismatched expectations reset ongoing processing, as in ART. In particular, a big enough mismatch in ART can activate the orienting system, and thereby cause category reset and attention shifts, as well as emotional rebounds from drive states that support positive emotions to drive states that support negative drive states, such as frustration. These rebounds can force unlearning, or extinction, of cognitive–emotional connections that support the behavior, and also disinhibit exploratory behaviors (Grossberg & Merrill, 1992).

47. START: spectral timing in the cerebellum, hippocampus, and basal ganglia

47.1. How do neurons code large temporal durations?

How do little nerve cells learn to bridge big time intervals? How do they manage to bridge a temporal interval of hundreds of milliseconds or even several seconds? This problem is similar to the one asked about how little nerve cells, such as hippocampal place cells, manage to represent big spatial regions (Section 40). As will be seen in Section 49, the apparent similarity of these two problems is reflected by a deeper similarity in their temporal and spatial mechanisms.

It is proposed that large temporal intervals can be bridged using a mechanism of *spectral timing* whereby a population “spectrum” of cells, or cell sites, with different reaction rates can learn to match the statistical distribution of expected delays in reinforcement over hundreds of milliseconds or even seconds. Although each of these cells, or cell sites, reacts briefly at different times, their population response as a whole can bridge a much longer time interval, ranging from hundreds of milliseconds to seconds, that can be tuned by learning to match external experimental contingencies. Fiala et al. (1996) have proposed that the brain uses the metabotropic glutamate receptor (mGluR) system to bridge over such long time intervals.




Fig. 10. Adaptively timed cerebellar learning circuit: adaptively timed Long Term Depression at Purkinje cells depresses the level of tonic inhibitory firing of these cells to cerebellar nuclei, thereby disinhibiting nuclear cells and allowing them to express their learned gains in an adaptively timed way. See text for details. Source: Reprinted with permission from Grossberg and Merrill (1996).

The Spectral Timing model clarifies many different aspects of adaptively timed learning, ranging from normal adaptively timed reinforcement learning, motivated attention, and action, via circuits involving basal ganglia, hippocampus, and cerebellum, to abnormal adaptive timing due to cerebellar lesions or autism, during which behaviors may be prematurely released in a context-inappropriate manner that can prevent these behaviors from receiving normal social rewards (Grossberg & Seidman, 2006; Grossberg & Vladusich, 2011; Sears, Finn, & Steinmetz, 1994). The ability to focus motivated attention for an adaptively timed interval is controlled by an adaptively timed cognitive–emotional resonance.

47.2. Adaptively timed learning in the hippocampus and cerebellum

Evidence for adaptive timing occurs during several types of reinforcement learning. For example, classical conditioning is optimal at a range of positive interstimulus intervals (ISI) between the conditioned stimulus (CS) and unconditioned stimulus (US) that are characteristic of the animal and the task. Such conditioning is greatly attenuated at zero and long ISIs. Within this range, learned responses are timed to match the statistics of the learning environment (Smith, 1968). Although the amygdala is a primary site for emotion and stimulus–reward association, the amygdala does not carry out adaptively timed learning. Instead, brain regions such as the hippocampus and cerebellum have been implicated in adaptively timed processing of cognitive–emotional interactions. For example, Thompson et al. (1987) distinguished two types of learning that go on during conditioning of the rabbit nictitating membrane response: Adaptively timed “conditioned fear” learning that is linked to the hippocampus, and adaptively timed “learning of the discrete adaptive response” that is linked to the cerebellum.

A unified explanation of why both hippocampus and cerebellum use adaptively timed learning is given by the START (Spectrally Timed ART) model (Fig. 7), which unifies the ART and CogEM models (Fiala et al., 1996; Grossberg & Merrill, 1992, 1996; Grossberg & Schmajuk, 1989). As noted above, CogEM predicts how salient conditioned cues can rapidly focus attention upon their sensory categories (S) via a cognitive–emotional resonance with their associated drive (D) representations. However, what then prevents the actions (M) that they control from being prematurely released?

For example, suppose that a conditioned stimulus (CS), say via the motor output pathway M in Fig. 7, activates pathways both to a subcortical cerebellar nucleus and to cerebellar cortex parallel fibers that synapse on Purkinje cells (Fig. 10). Suppose that the parallel fibers let the CS activate pathways that respond with a spectrum of differently timed intracellular processes. The unconditioned stimulus (US) activates climbing fibers to provide a teaching signal that converges upon these parallel fiber/Purkinje cell synapses. This teaching signal causes the active synapses within the parallel fiber spectrum to become weaker (Long Term Depression) if they are activated by the CS at times when the US teaching signal is active. Synapses whose CS-activated spectral activity does not overlap the climbing fiber signals become stronger (Long Term Potentiation, or LTP). Because the Purkinje cells tonically inhibit their subcortical target cells, their adaptively timed LTD disinhibits the tonic inhibition by Purkinje cell of cerebellar nuclear cells. In other words, an adaptively timed gate opens and allows the subcortical cells to fire at appropriate times. The model proposes that climbing fibers also control learning of adaptive gains along subcortical pathways through the nuclear cells. Thus, when the adaptively timed Purkinje cell gate opens, the learned gains can be expressed at the correct times and with the correct amplitude to cause a correctly calibrated motor response.

47.3. Adaptively timed learning by the metabotropic glutamate receptor system

Fiala et al. (1996) have developed and simulated a detailed Spectral Timing model of cerebellar adaptive timing. This model builds on Eqs. (1)–(4), linking biochemistry, neurophysiology, neuroanatomy, and behavior to predict how the metabotropic glutamate (mGluR) receptor system may create a spectrum of delays during cerebellar adaptively timed learning. Spectral timing in the hippocampus also presumably has an mGluR realization (Brown et al., 1999; Grossberg & Merrill, 1992, 1996). The Fiala et al. (1996) prediction of a role for mGluR in adaptively timed learning is consistent with data about calcium signaling and mGluR in cerebellar adaptive timing (Finch & Augustine, 1998; Ichise et al., 2000; Miyata et al., 2000; Takechi, Eilers, & Konnerth, 1998). The model simulates both normal adaptively timed conditioning data and premature responding when cerebellar cortex is lesioned (Perret, Ruiz, & Mauk, 1993), thereby eliminating the adaptively timed gating process. Indeed, various individuals with autism, who are known to have cerebellar deficiencies, also perform short-latency responses in the eye-blink paradigm (Grossberg & Seidman, 2006; Sears et al., 1994). The prediction of a role for mGluR in adaptively timed learning raises the question of whether the mGluR system is not functioning normally in some autistic individuals.

48. Adaptively timed attentive resonance and action

48.1. Reconciling fast motivated attention with adaptively timed motivated attention and action

The combination of cerebellar and hippocampal adaptive timing reconciles three important behavioral properties: Fast allocation of attention to motivationally salient events via a cortico-amygdala resonance, maintenance of motivated attention for an adaptively timed duration via a cortico-amygdala-hippocampal resonance, and adaptively timed responding via a cortico-cerebellar circuit during this adaptively timed duration. Indeed, for motivated attention to be effective, it needs to be activated as soon as important motivationally salient cues are detected. However,

for adaptively timed responding to be effective, motivated attention needs to persist long enough to support the read-out of adaptively timed motor commands, and to prevent irrelevant events from prematurely resetting these commands. The START model (Grossberg & Merrill, 1992, 1996) proposes how dentate-CA3 circuits in the hippocampus can regulate the duration of a CogEM cognitive-emotional resonance that includes cortex (S_i^j in Fig. 7), amygdala (D in Fig. 7), and hippocampus (T in Fig. 7).

Without such adaptively timed maintenance of cortical activity, an unexpected event could prematurely reset motor commands that are being read out by the cortex. As noted in Section 46, if a top-down prototype and bottom-up sensory input mismatch too much for resonance to occur, then the orienting system A can reset the active categories that are controlling read-out of the motor command (Fig. 2). The hippocampal system and nonspecific thalamus are proposed to be part of this mismatch-activated orienting system. These mismatches are not, however, sensitive to whether the novel event that caused the mismatch is task-relevant, and whether it is important enough interrupt the ongoing task. The START model clarifies how such mismatches may be modulated by task-relevance in an adaptively timed way. In particular, inhibition from the adaptive timing circuit can prevent a mismatch from causing premature category reset, as explained below.

Figs. 7 and 11 suggest how adaptively timed learning within the dentate-CA3 circuits (T in Fig. 7) of the hippocampus is proposed to inhibit the activation of the orienting system A during an interval wherein a valued and predictable goal is being acted upon. Indeed, dentate-CA3 cell firing reflects the learned delays observed during the rabbit nictitating membrane response (Berger, Berry, & Thompson, 1986). The START model proposes how adaptively timed inhibition of the hippocampal orienting system (Figs. 7 and 11) and adaptively timed disinhibition of cerebellar nuclear cells (Fig. 10) may be coordinated to enable motivated attention to be maintained on a goal while adaptively timed responses are released to obtain a valued goal.

48.2. Interacting contributions of amygdala, hippocampus, and cerebellum during conditioning

The START circuitry that is summarized in Fig. 7 combines CogEM, ART, and Spectral Timing properties in the following way: As noted in Section 42, processing stages $S^{(1)}$ and $S^{(2)}$ in Fig. 7 play the role of sensory cortex and orbitofrontal cortex, respectively. Stage D is an emotional center, or drive representation, like the amygdala. Stage M schematizes motor output pathways. The feedback pathways $D \rightarrow S^{(2)} \rightarrow S^{(1)}$ from a particular drive representation to sensory representations are capable of focusing attention on motivationally consistent events in the world. The excitatory pathways from $S^{(1)} \rightarrow D$ learn the conditioned reinforcer properties of a sensory cue, such as a CS, whereas the pathways $D \rightarrow S^{(2)}$ learn the incentive motivational properties of cues. Representations in $S^{(2)}$ can fire vigorously only if they receive convergent signals from $S^{(1)}$ and D , corresponding to the sensitivity of orbitofrontal cortex cells to converging sensory and reinforcing properties of cues. Once activated, these orbitofrontal cells deliver positive feedback to $S^{(1)}$ and bias the competition among sensory representations to focus attention on valued events while attentionally blocking inhibited features.

Prior to conditioning, a CS can be stored at $S^{(1)}$ and can prime D and $S^{(2)}$ without firing these representations. After conditioning, the CS can trigger strong conditioned $S^{(1)} \rightarrow D \rightarrow S^{(2)} \rightarrow S^{(1)}$ feedback and rapidly draw attention to itself as it activates the emotional representations and motivational pathways controlled by D . Representation D can also inhibit the orienting system A as it focuses attention upon motivationally valued sensory events (Figs. 7 and 11). Here is thus one way in which the CogEM and

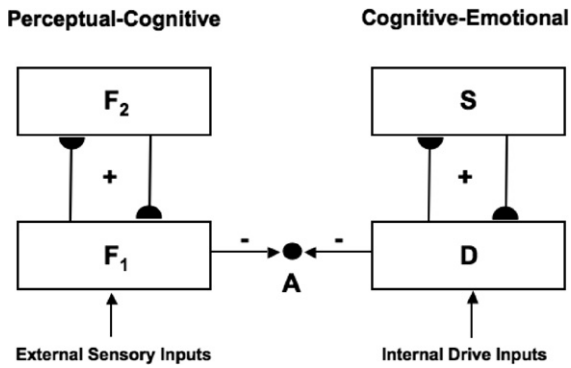


Fig. 11. Mismatch and Spectral Timing circuits both inhibit the orienting system: when a big enough mismatch occurs between distributed features at level F_1 and the top-down expectation from the category level F_2 , then inhibition is reduced at the orienting system (Fig. 2) enough to trigger activation of the orienting system A . However, if adaptively timed inhibition from the cognitive–emotional sensory-drive ($S - D$) resonance also inhibits A , then the orienting system cannot fire until the adaptively timed signal is removed.

ART models interact: Emotionally salient goal objects can inhibit the orienting system and thus prevent irrelevant distractors from attracting attention even when there is an ART mismatch.

This inhibition of the orienting system becomes adaptively timed as follows: The sensory representations $S^{(1)}$ send pathways to a spectral timing circuit T , assumed to be in the dentate–CA3 region of the hippocampus, whose adaptive weights w are trained by a Now Print, or teaching signal, N . The teaching signal N is transiently activated by changes in the activity of the drive representation D that occur when a reinforcing event activates D . Weight learning can also be modulated by dopamine signals from the SNr in the basal ganglia (Brown et al., 2004).

After conditioning of T takes place, adaptively timed read-out from T can enable attention to focus on task-relevant cues by maintaining their cortical representations $S^{(2)}$ in a cortico-amygdala–hippocampal resonance while T inhibits the orienting system A for an adaptively timed duration. In Figs. 7 and 11, the simplest inhibitory path from T -to- D -to- A is depicted. A more complex set of pathways exists *in vivo*.

In summary, the START model enables three behaviorally important properties to coexist:

Fast motivated attention. Rapid focusing of attention on motivationally salient cues occurs from regions like the amygdala to prefrontal cortex (pathway $D \rightarrow S^{(2)}$ in Fig. 7). Without further processing, fast activation of the CS-activated $S^{(2)}$ sensory representations could prematurely release motor behaviors (pathway $S^{(2)} \rightarrow M$ in Fig. 7).

Adaptively timed responding. Adaptively timed read-out of responses via cerebellar circuits (pathway M in Fig. 10) enables learned responses to be released at task-appropriate times, despite the fact that CS representations which activate cerebellum can be quickly activated by fast motivated attention.

Adaptively timed duration of motivated attention and inhibition of orienting responses. Premature reset of active CS representations by irrelevant cues during task-specific delays is prevented by adaptively timed inhibition of mismatch-sensitive cells in the orienting system of the hippocampus (pathway $T \rightarrow D \rightarrow A$ in Fig. 7). This inhibition is part of the competition between consummatory and orienting behaviors (Staddon, 1983). Adaptively timed incentive motivational feedback ($D \rightarrow S^{(2)} \rightarrow S^{(1)}$ in Fig. 7) simultaneously maintains CS activation in short-term memory, so that the CS can continue to read-out adaptively-timed responses until they are complete.

The Contingent Negative Variation, or CNV, event-related potential is predicted to be a neural marker of adaptively timed motivational feedback. Many additional data have been explained

using these circuits, including data from delayed non-match to sample (DNMS) experiments wherein both temporal delays and novelty-sensitive recognition processes are involved (Gaffan, 1974; Mishkin & Delacour, 1975).

49. Neural relativity: spectral timing and spectral spacing in the hippocampus

As noted in the discussions of Sections 39, 40 and 46–48, the hippocampus processes spatial as well as temporal information. This observation raises the question: Why are both space and time both processed in the hippocampus? The fact of this convergence is consistent with data and hypotheses about a possible role of hippocampus in episodic learning and memory, since episodic memories typically combine both spatial and temporal information about particular autobiographical events; e.g., Eichenbaum and Lipton (2008).

As noted in Sections 39–40, multiple scales of entorhinal grid cells can cooperate in a self-organizing map to form place cell receptive fields. These multiple scales form along a dorsoventral spatial gradient in the entorhinal cortex such that grid cells have increasingly large spatial scales (that is, larger spatial intervals between activations in a hexagonal grid) in the ventral direction. Grid cells with several different spatial scales along the dorsoventral gradient can cooperate to form place cells that can represent spaces much larger than those represented by individual grid cells, indeed place cells capable of representing the lowest common multiple of the grid cell scales that activate them (Gorchetnikov & Grossberg, 2007; Pilly & Grossberg, *in press*).

With this background, we can begin to see a similarity in how the entorhinal–hippocampal system deal with both time and space. In the case of temporal representation by Spectral Timing, a spectrum of small time scales can be combined to represent much longer and behaviorally relevant temporal times. In the case of spatial representation by grid cells, a “spectrum” of small grid cell spatial scales can be combined to represent much larger and behaviorally relevant spaces. This homology has led to the name Spectral Spacing for the mechanism whereby grid cells give rise to place cells.

The Spectral Timing model reflects the part of entorhinal–hippocampal dynamics that is devoted to representing objects and events, and includes lateral entorhinal cortex. The Spectral Spacing model reflects a complementary part of entorhinal–hippocampal dynamics that is devoted to representing spatial representations, and includes medial entorhinal cortex. Both of these processing streams are joined in the hippocampus to support spatial navigation as well as episodic learning and memory (Eichenbaum & Lipton, 2008).

This heuristic homology between spatial and temporal representations is supported by rigorous mathematical modeling and data simulations. Grossberg and Pilly (*in press*) have developed the Spectral Spacing model to show that neural mechanisms which enable a dorsoventral gradient of grid cell spatial scales to be learned are formally the same as mechanisms that enable a gradient of temporal scales to form in the Spectral Timing model. Grossberg and Pilly (*in press*) support this claim by quantitatively simulating challenging data about parametric properties of grid cells along the dorsoventral gradient. Thus, it may be that space and time are both in the hippocampus because they both exploit a shared set of computational mechanisms. Briefly said: Space and Time are One. The phrase “neural relativity” tries to celebrate this predicted property of the entorhinal–hippocampal system.

50. ART in time: working memory, list chunks, and sequential performance

The above mechanisms do not explain how the brain responds selectively to *sequences* of events. Predictive behavior depends upon the capacity to think about, plan, execute, and evaluate such event sequences. In particular, multiple events in a specific temporal order can be stored temporarily in *working memory*. As event sequences are temporarily stored, they are grouped, or chunked, through learning into unitized plans, or list chunks, and can later be performed at variable rates under volitional control. Examples of list chunks include unitized representations of familiar words and motor skills. Here, the term working memory is used to describe brain processes that temporarily store the *temporal order of several events*, not merely persistence of individual events. ART categorization mechanisms are used to learn such list chunks, as would be expected from the fact that sequential behaviors must also solve the stability–plasticity dilemma.

50.1. Item and order working memories and learned list chunks

Grossberg (1978a, 1978b) introduced a model of working memory and of list chunking, which proposed how working memories are designed to enable list chunks of variable length to be quickly learned and stably remembered. Inputs to this working memory are unitized *item categories* of individual events or objects (e.g., invariant object categories, acoustically processed letter categories, simple motor synergies) that have enough adaptively timed, incentive motivational support to be persistently stored before being transferred into working memory. Sequences of items are stored in working memory as a temporally evolving spatial pattern of activity across item-selective cells. The *relative activities* of different cell populations code the temporal order in which the items will be rehearsed, with the largest activities rehearsed earliest; hence, the name *Item and Order working memory* for this class of models. A more recent name is *competitive queuing* (Houghton, 1990).

50.2. All working memories obey the LTM Invariance Principle and the Normalization Rule

Grossberg (1978a, 1978b) derived Item and Order working memories from two postulates: the LTM Invariance Principle and the Normalization Rule. The LTM Invariance Principle is the main postulate. It makes precise the idea that there is no point in storing novel sequences of events in working memory if the brain cannot learn to unitize the sequences for future skillful performance. The LTM Invariance Principle claims that working memories are *designed* to enable such stable list chunking to occur. In particular, it demands that all working memories enable a novel superset list chunk (up to a certain maximal length) to be learned without forcing catastrophic forgetting of familiar subset chunks. For example, the LTM Invariance Principle ensures that a novel superset word like MYSELF can be learned without forcing forgetting of the familiar subwords MY, SELF, and ELF. As a result, as new items are stored through time in working memory, subset list chunks can continue to activate their familiar list chunks until they are inhibited by contextually more predictive superset list chunks; e.g., until MY is supplanted by competition from MYSELF through time. The *learning* of chunk MY within its bottom-up filter is not undermined, but the *current activation* of the chunk MY can be inhibited by MYSELF. Mathematically, this boils down to the following property: Activities of items in working memory tend to preserve their *relative* activations, or *ratios*, throughout the time that they are stored in working memory, even if the storage of new

items through time might change the absolute amount of activity with which each item is stored. This property enables the adaptive filter that converts the distributed pattern of stored items into list chunks (see Fig. 2(a)) to activate already learned list chunks in response to the individual sublists in working memory.

The Normalization Rule assumes that the *total activity* of the working memory network has a maximum that is (approximately) independent of the total number of actively stored items. In other words, working memory has a *limited capacity* and activity is redistributed, not just added, when new items are stored.

50.3. Item and order working memories are recurrent shunting on-center off-surround networks

How can brain evolution be smart enough to discover the laws of something so seemingly sophisticated as a working memory? Remarkably, an Item and Order working memory that satisfies the LTM Invariance Principle and Normalization Rule can be realized by an on-center off-surround network whose cells obey the shunting, or membrane, equations of neurophysiology and which interact via *recurrent* on-center off-surround connections (see Eq. (1)). Recurrent shunting on-center off-surround networks are ubiquitous in the brain, as they provide a solution of the noise–saturation dilemma (Section 6). Recurrence is needed because the positive feedback from a cell population to itself in the recurrent on-center stores the evolving input pattern, while the recurrent competition contrast-normalizes the stored activities across the network. The shunting, or multiplicative, properties of the membrane equations, combined with the on-center off-surround interactions, enable the network to compute *ratios* of cell activities across the network, as is required by the LTM Invariance Principle.

What sets an Item and Order working memory apart from just any recurrent shunting on-center off-surround network? Such a recurrent network behaves like an Item and Order working memory model when it is equipped with a volitionally-activated nonspecific *rehearsal wave* to initiate read-out of stored activity patterns, and output-contingent self-inhibitory feedback to prevent perseverative performance of the most active stored item, and thereby achieve inhibition of return (Fig. 12).

The fact that an Item and Order working memory can be realized by a recurrent shunting on-center off-surround network is consistent with the ART Matching Rule. When a list chunk reads out a learned top-down expectation into a working memory, the modulatory on-center of the ART Matching Rule is converted into a driving on-center by volitional control from the basal ganglia, and the off-surround is just what is needed to compute the ART Matching Rule; see Section 7.

51. Why primacy and bowed serial properties in human and animal working memory data?

51.1. Stable chunking implies primacy, recency, or bowed gradients in working memory

Why are there limits to the number of items that can be stored in working memory in the correct temporal order? In an Item and Order working memory, it can be mathematically proved that, under constant attentional conditions, the pattern of activation that evolves in an Item and Order working memory is one of following types, depending on how network parameters are chosen (Bradski, Carpenter, & Grossberg, 1992, 1994; Grossberg, 1978a, 1978b):

Primacy gradient. Here, the first item to be stored has the largest activity and the last item to be stored has the smallest activity. A

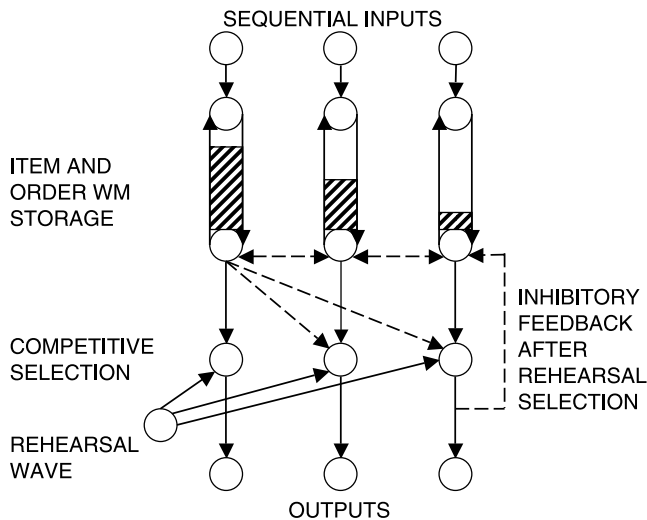


Fig. 12. Schematic of an item and order working memory: a temporal sequence of inputs creates a spatial activation pattern among STM activations, often a primacy gradient (height of hatched rectangles is proportional to cell activity). Relative activation levels among stored items codes both which items are stored and the temporal order in which they are stored. A nonspecific rehearsal wave allows item activations to be rehearsed, with the largest activity being read out first. The output signal from this item also activates a self-inhibitory interneuron that inhibits the item, and thereby enables the next most active item to be performed. The process then repeats itself.

Source: Reprinted with permission from Grossberg and Pearson (2008).

primacy gradient allows the stored items to be rehearsed in their presented order.

Recency gradient. Here the first item is stored with the smallest activity and the last item with the largest activity. Rehearsal of a recency gradient recalls the most recent item first and the first item last.

Bowed gradient. Here, the first and last items to be stored have larger activities, and thus are earlier rehearsed, than items in the middle of the list.

From primacy to bowed gradient. It was also proved that, as more and more items are stored, a primacy gradient becomes a bowed pattern whose recency part becomes increasingly dominant.

This last result predicts why bowed gradients are found in many types of serially ordered behavior: The property of stable learning and memory of list chunks imposes a severe limitation on the number of items that can be recalled in the correct temporal order from working memory, because a bow necessarily occurs in the stored gradient of sufficiently long lists.

51.2. Psychological and neurophysiological data support predicted item and order properties

The properties of Item and Order working memories have been supported by subsequent psychological and neurobiological experiments. For example, Farrell and Lewandowsky (2002) concluded from their human psychophysical data: “Several competing theories of short-term memory can explain serial recall performance at a quantitative level. However, most theories to date have not been applied to the accompanying pattern of response latencies: . . . these data rule out three of the four representational mechanisms. The data support the notion that serial order is represented by a primacy gradient that is accompanied by suppression of recalled items”. Page and Norris (1998) adapted the Item and Order working memory to describe a “primacy model” which they used to fit data about immediate serial recall, notably the effects of word length, list length, and phonological similarity.

The LTM Invariance Principle should be obeyed in all working memories, whether they store verbal, spatial, or motor items, so

that all working memories should exhibit similar data properties, such as error distributions. Jones, Farrand, Stuart, and Morris (1995) reported similar performance characteristics to those of verbal working memory for a spatial serial recall task, in which visual locations were remembered. Agam, Bullock, and Sekuler (2005) reported psychophysical evidence of Item and Order working memory properties in humans as they performed sequential copying movements. Averbek, Chafee, Crowe, and Georgopoulos (2002); Averbek, Crowe, Chafee, and Georgopoulos (2003a, 2003b) reported the first neurophysiological evidence in monkeys that a primacy gradient, together with inhibition of the most active cell after its command is read out, governs the sequential performance of sequential copying movements.

Agam, Galperin, Gold, and Sekuler (2007) reported data consistent with the formation of list chunks as movement sequences are practiced, thereby supporting the prediction that working memory networks are designed to interact closely with list chunking networks. A number of psychophysics experiments have been carried out to successfully test this hypothesis during speech perception (e.g., Auer and Luce (2008), Goldinger and Azuma (2003), McLennan, Conon, and Luce (2005), McLennan, Luce, and Charles-Luce (2003), Luce and McLennan (2008), and Vitevitch and Luce (1999)).

The fact that verbal, spatial, and motor sequences, in both humans and monkeys, seem to obey the same working memory laws provides accumulating evidence for the Grossberg (1978a, 1978b) prediction that all working memories have a similar design to enable stable list chunks to form by ART learning mechanisms.

52. Free recall and the magical numbers four and seven in item-order-rank working memory

The prediction that primacy gradients become bows for longer lists provides a conceptually satisfying explanation of the well-known *immediate memory span* of $7 + / - 2$ items (Miller, 1956). This property of Item and Order working memories was originally used to explain data about free recall (the recall of a once-heard list in any order) and related paradigms in which bowing effects are observed (Grossberg, 1978b). Indeed, because relative activity translates into both *relative order* and *probability of recall* (bigger activities can provide earlier and more reliable recall in a noisy brain), such a model helps to explain why items from the beginning and end of a list in free recall may be recalled earlier and with larger probability (Murdock, 1962). *Transposition errors* also have a natural explanation in such a working memory, since stored items with similar activity levels will transpose their relative activities, and thus their rehearsal order, more easily than items with very different activity levels if noise perturbs these levels through time. Grossberg (1978a, 1978b) also proved that, if attention varies across items, then multi-modal bows, or Von Restorff (1933) effects, also called isolation effects (Hunt & Lamb, 2001), occur when an item in a list “stands out like a sore thumb” and is thus more likely to be remembered than other list items. Von Restorff effects can also be caused by rate and feature-similarity differences across items, factors which also influence bowing in the present modeling framework. Associative and competitive mechanisms that are consistent with the Item and Order working memory model have also been used to explain Von Restorff effects during serial verbal learning (Grossberg, 1969c, 1974).

52.1. Transient and immediate memory spans

The Grossberg (1978a) analysis distinguished between the classical *immediate memory span* (IMS) of Miller (1956) and the model-derived new concept of *transient memory span* (TMS). The TMS was predicted to be the result of short-term working memory

storage recall, without a significant top-down long-term memory (LTM) component from ART mechanisms. The TMS is, accordingly, the longest list length for which a working memory can store a primacy gradient. The IMS was predicted to be the result of combining bottom-up inputs and top-down read-out of list chunk learned expectations at the working memory. How do top-down learned signals alter the relative activities stored in working memory, and thus the temporal order that is recalled? The read-out of list chunk top-down expectations into working memory was mathematically proved to generate an extended primacy gradient in working memory, and thus a longer list that could be recalled in the correct order.

Thus, ART top-down feedback to an Item and Order working memory leads to an IMS that is longer than the TMS. Estimating the IMS at seven, Grossberg (1978a) predicted that the TMS would be around four. Subsequent data, reviewed by Cowan (2001), support the existence of a four plus-or-minus one working memory capacity limit when learning and grouping influences are minimized. Indeed, long-term memory (LTM) does bias working memory towards more primacy dominance (e.g. Knoedler, Hellwig, and Neath (1999)), and its influence can be difficult to limit. Cowan (2001) reviewed proposals for limiting LTM influence, such as using novel sequence orderings of well-learned items that are difficult to group.

52.2. LIST PARSE: laminar cortical circuits for working memory and list chunk learning

The LIST PARSE (Laminar Integrated Storage of Temporal Patterns for Associative Retrieval, Sequencing and Execution) model (Grossberg & Pearson, 2008) built on this foundation to predict how laminar circuits in ventrolateral prefrontal cortex may embody a cognitive Item and Order working memory and list chunk learning network that interacts with a motor working memory in dorsolateral prefrontal cortex, and a basal ganglia adaptively-timed volitional control system (Fig. 13). LIST PARSE is able to store and learn event sequences of variable length and to perform them at variable speeds under volitional control. The LIST PARSE cognitive model uses a variation of the laminar circuits that are used in the LAMINART vision model (Section 19).

Accumulating neurobiological data support the view that visual and verbal object categories may be learned in temporal and ventromedial prefrontal (e.g., orbitofrontal cortex) cortex. As predicted by the CogEM model (Section 42), orbitofrontal cortex responds best to the combination of item and motivational signals. These attended item representations are then loaded into a sequential working memory that codes temporal order information in ventrolateral and/or dorsolateral prefrontal cortex (e.g., Barbas (2000), Goldman-Rakic (1987), Petrides (2005) and Ranganath and D'Esposito (2005)). The temporally evolving working memory patterns are, in turn, categorized by list chunks, or sequential plans, which can be used to predict subsequent expected events.

Grossberg and Pearson (2008) provide a full review of the LIST PARSE model, relevant data, and other models of working memory. Here it suffices to note that LIST PARSE unifies the explanation of cognitive, neurophysiological, and anatomical data from humans and monkeys concerning how predictive, sequentially organized behaviors are controlled. Because its laminar cognitive prefrontal circuits are variations of laminar circuits in visual cortex, both cognition and vision seem to use variations of a shared laminar cortical design to rapidly learn and stably remember, using ART top-down attentive matching mechanisms, recognition categories with which to predict a changing world. In particular, LIST PARSE quantitatively simulates human cognitive data about immediate serial recall and immediate, delayed, and continuous-distractor

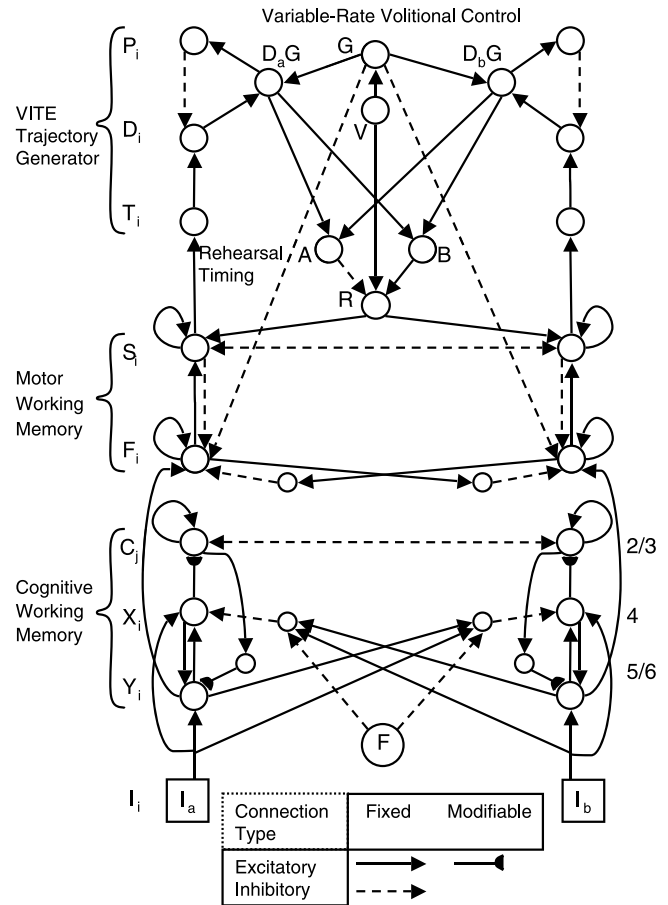


Fig. 13. Circuit diagram of the LIST PARSE model. The item and order working memory is realized by a recurrent shunting on-center off-surround network in layers 4 and 6 of the Cognitive Working Memory, which is assumed to occur in ventrolateral prefrontal cortex. The list chunks are learned in layer 2/3. Outputs from the Cognitive Working Memory to the Motor Working Memory interact with a Vector Integration to Endpoint (VITE) trajectory generator (Bullock & Grossberg, 1988), modulated by the basal ganglia, to perform sequences of variable length at variable speeds. Solid arrows indicate fixed excitatory connections. Solid lines with hemi-disks indicate modifiable (i.e., learned) connections. Dashed arrows indicate fixed inhibitory connections. Only 1-item chunks (C) and their feedback connections (M) within a single Cognitive Working Memory channel are shown, whereas the model uses chunks of various sizes in layer 2/3 and feedback from layers 2/3 to 5/6 of the Cognitive Working Memory is broadly distributed. Also, only the excitatory projections from Cognitive Working Memory to the Motor Plan Field ($Y \rightarrow F$) are shown.

Source: Reprinted with permission from Grossberg and Pearson (2008).

free recall, as well as monkey neurophysiological data from the prefrontal cortex obtained during sequential sensory-motor imitation and planned performance. It hereby clarifies how both spatial and non-spatial working memories may share the same laminar cortical circuit design.

53. Item-order-rank working memory: from numerical to positional coding

53.1. Storing sequences of repeated items

Working memory can store sequences of items that may be repeated several times in the sequence; e.g., *ABACBD*. How does the brain correctly store different ordinal positions for the same item in a sequence? Can an Item and Order working memory do this? For example, how does working memory store multiple target locations to control sequences of eye movements, particularly when the same locations repeat at multiple list positions, or

ranks, in the sequence? Both psychophysical and neural data indicate that rank information may be explicitly coded by the brain. For example, there is an increased likelihood of intrusions from items at similar list positions across trials (reviewed in Henson (1998)). Repetitive presentation of an item in the same absolute serial position in lists across multiple trials with different list orderings leads to a statistically significant increase in the strength of learning of that item (Conway & Christiansen, 2001). There is a statistically small, but significant, class of errors called “protrusions”, in which interchanges between items in similar serial positions of different temporal groupings in a list are more common than interchanges between items in different serial positions (Henson, 1998). That is, when presented with the sequence *ABC* (pause) *DEF*, interchanges between items at *B* and *E* are more common than interchanges between items at *B* and *F*. The well-known “slip of the tongue” error of “spoonerisms” illustrates a similar effect. Here, phonemes or syllables in similar positions in different words are selectively interchanged (e.g., “hissed my mystery lesson”). It is also known that there are cells in the brain which are sensitive to rank. For example, a cell may respond to any target in the initial, middle or final ordinal position in a sequence, or the first, second, third, etc. positions. These cells are said to exhibit *rank order*, also called temporal selectivity or position specificity, while other cells may exhibit selectivity to *conjunctive coding* of item and ordinal position. Such a cell may respond to a specific target presented in a specific ordinal position (Averbeck et al., 2003a; Funahashi, Inoue, & Kubota, 1997; Inoue & Mikami, 2006).

Positional theories (e.g., Burgess (1995), Burgess and Hitch (1992), and Henson (1998)) posit that positional information is explicitly maintained to establish recall order. Such positional information may be absolute, or relative to the beginning and/or end of lists. However, these models all have conceptual weaknesses and do not adequately treat item repetitions; see Grossberg and Pearson (2008, Section 7.2.2) for a review.

53.2. From parietal numerical representations to prefrontal item–order–rank working memory

Grossberg and Pearson (2008) proposed that one source of positional information in the frontal cortex may be derived from the cortical maps of numerical representation that are found in the parietal cortex, with projections to the frontal cortex (Dehaene, 1997; Nieder & Miller, 2003, 2004). LIST PARSE predicts that such numerical maps contribute positional coding signals to the conjunctive coding of list item, order and rank that is found in higher cortical areas.

These numerical maps are, for example, important in foraging behaviors. They enable animals to make decisions about the total number of food items that are available in different locations, thereby guiding navigational and acquisitive behaviors (Buchanan & Bitterman, 1998; Emmerton, Lohmann, & Niemann, 1997; Washburn & Rumbaugh, 1991). They can also support an abstract representation of number, which can represent the total number of events across modalities; e.g., seeing and hearing (Church & Meck, 1984; Starkey, Spelke, & Gelman, 1983). They thus seem to have been involved in planning sequential behaviors for a very long time.

How do such numerical representations arise? Where do numbers come from?! Grossberg and Repin (2003) developed the Spatial Number Network, or SpaN model, to quantitatively simulate psychophysical and neural properties of numerical estimation using variations of more primitive cortical mechanisms for motion perception and spatial localization that are ubiquitous in the Where cortical stream, including the parietal cortex. In the model’s parietal representation, distinct but overlapping cell

populations in an ordered cortical map are activated as an increasing number of events is detected. Thus, as a sequence is stored in working memory, each item in the sequence can activate the numerical map and cause a shift in the location of maximal map activity. These activated positions are projected to prefrontal working memory networks, where they are combined with item and order information to generate cells that are sensitive to item–order–rank information; that is, cells for which activity gradient of order information is spread over cells that are sensitive to a conjunction of item-and-rank information.

53.3. *lisTELOS*: prefrontal item–order–rank working memories explain microstimulation data

Silver, Grossberg, Bullock, Histed, and Miller (2011) have incorporated such an item–order–rank spatial working memory into the *lisTELOS* model, which is a comprehensive model of how sequences of eye movements, which may include repetitions, may be planned and performed (Fig. 14). Similar mechanisms may be expected to control other types of sequences as well, due to shared working memory designs required by the LTM Invariance Principle (Section 51).

The *lisTELOS* model’s name derives from the fact that it unifies and further develops concepts about how item–order–rank working memories store lists of items, and of how *TELOS* properties of the basal ganglia (Section 44) help to balance reactive vs. planned movements by selectively gating sequences of actions through time. Rank-related activity has been observed in many areas including the posterior parietal cortices (PPC), prefrontal cortices (PFC) and supplementary eye fields (SEF). The *lisTELOS* model shows how rank information, originating in PPC, may support item–order–rank PFC working memory representations, and how SEF may select saccades stored in working memory. The model also proposes how SEF may interact with downstream regions such as the frontal eye fields (FEF) during memory-guided sequential saccade tasks, and how the basal ganglia (BG) may control the flow of information. Model simulations reproduce behavioral, anatomical and electrophysiological data under multiple experimental paradigms, including visually- and memory-guided single and sequential saccade tasks. Simulations quantitatively reproduce behavioral data during two SEF microstimulation paradigms (Histed & Miller, 2006; Yang, Heinen, & Missal, 2008). SEF microstimulation altered the order in which eye movements are carried out, but not the target locations to which the eyes move. Quantitative simulations of such data provide strong evidence for item–order–rank coding in spatial working memory; in particular, that item-and-order information are encoded by an activity gradient whose relative activities, and thus the order of recall, are perturbed by microstimulation. The simulations also show how seemingly inconsistent findings about saccade latency of different microstimulation experiments can be given a unified explanation using the model.

54. ARTSCENE search and recognizing scenes: accumulating evidence using working memory

54.1. Searching scenes using contextual cueing

The prediction that all working memories obey the LTM Invariance Principle implies that a similar working memory design should be used in visual, auditory, spatial, and motor working memories; see Section 51. In particular, both visual and spatial working memories, and their learned list chunks, help the brain to learn how to use context to efficiently search a scene for a desired object. For example, humans can learn that a certain combination of objects may define a context for a kitchen and

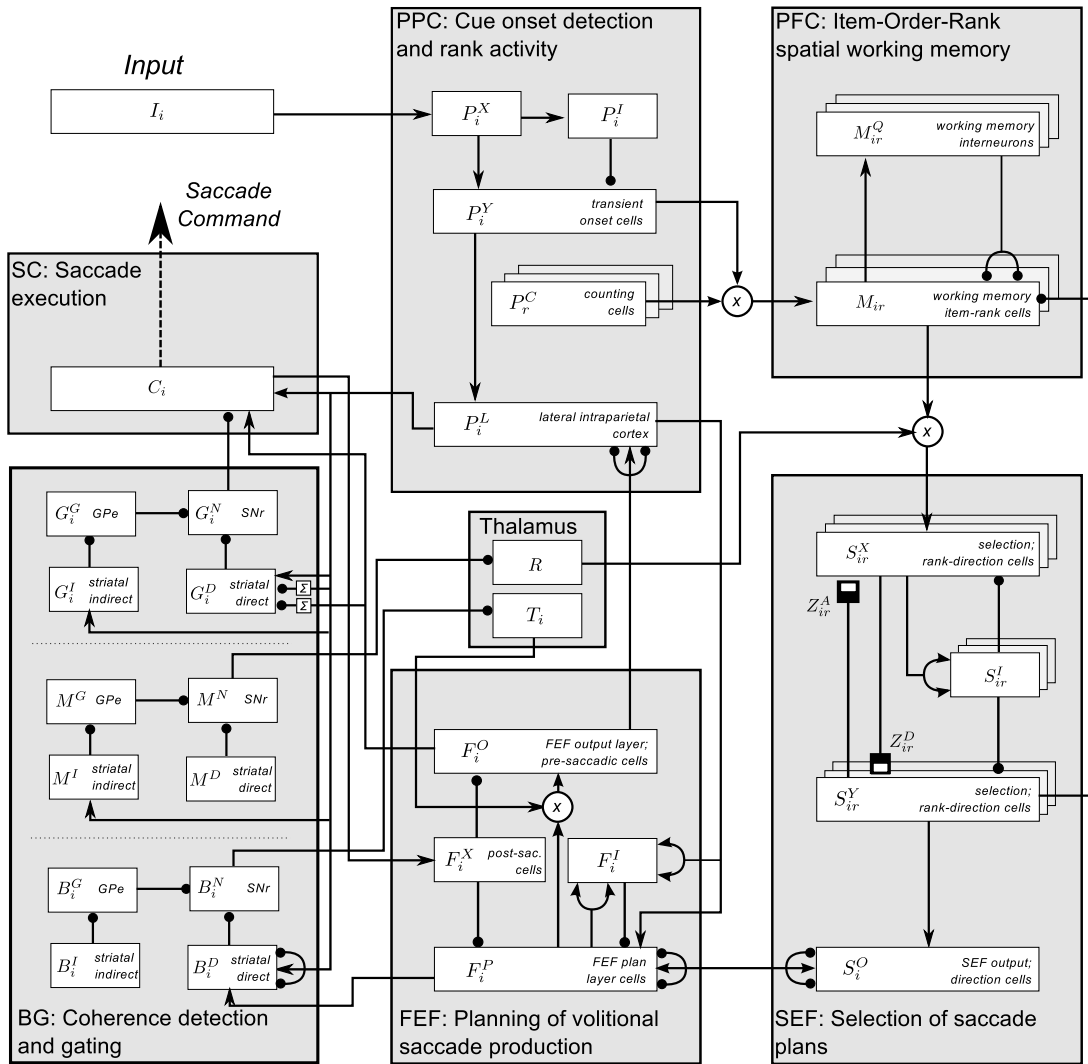


Fig. 14. The listTELOS model macrocircuit: each gray box represents a brain region within which fields of cells, represented by white inset boxes, share similar functional roles, which are summarized in the box. Arrowheads denote excitatory connections between cells, and filled circles represent inhibitory connections. Curved branches at the ends of connections represent one-to-many fan-out connections that impact all other cells in the field. Half-filled boxes at the ends of connections represent habituitive gates which exhibit activity-dependent changes in synaptic efficacy. White circles containing a multiplication sign (\times) represent multiplicative interaction between two signals. Boxes containing a sigma (Σ) represent the sum of outputs from all cells in the field that gave rise to the projection. Stacked field representations denote populations of rank-sensitive cells. SC = superior colliculus; PPC = posterior parietal cortex; PFC = prefrontal cortex; BG = basal ganglia; FEF = frontal eye fields; SEF = supplementary eye fields. Note the three BG loops gating the release of output signals from different brain regions. Source: Reprinted with permission from Silver et al. (2011).

trigger a more efficient search for a typical object, such as a sink, in that context. The ARTSCENE Search model (Figs. 15 and 16; Huang and Grossberg (2010)) was developed to characterize the neural mechanisms of such memory-based context learning and guidance, which is often called *contextual cueing* (Chun, 2000; Chun & Jiang, 1998).

54.2. Scenic gist as a large-scale texture category

ARTSCENE Search simulates challenging behavioral data on positive/negative, spatial/object, and local/distant contextual cueing effects during visual search, as well as related neuroanatomical, neurophysiological, and neuroimaging data. The model proposes how the global scene layout, or *gist* of a scene, can be used at a glance to rapidly form a hypothesis about the environment. Such a gist can be learned as a large-scale ART texture category (Grossberg & Huang, 2009). This hypothesis is then incrementally refined as a scene is sequentially scanned with saccadic eye movements, and the sequences of scanned objects and spatial positions are stored,

and unitized, to learn how to better predict target identity and position as the scene becomes familiar. The ARTSCENE Search model hereby simulates the interactive dynamics of object and spatial contextual cueing and attention in the cortical What and Where streams starting from early visual areas through medial temporal lobe to prefrontal cortex. After learning, model dorsolateral prefrontal cortex (area 46) in the model's Where stream primes possible target locations in posterior parietal cortex based on goal-modulated percepts of spatial scene gist that are represented in parahippocampal cortex. In parallel, model ventral prefrontal cortex (area 47/12) in the model What stream primes possible target identities in inferior temporal cortex based on the history of viewed objects represented in perirhinal cortex.

55. CARTWORD: conscious resonances of speech items, working memory, and list chunks

Item and Order working memories and their list chunks play a critical role in the learning of language. ART has been developed

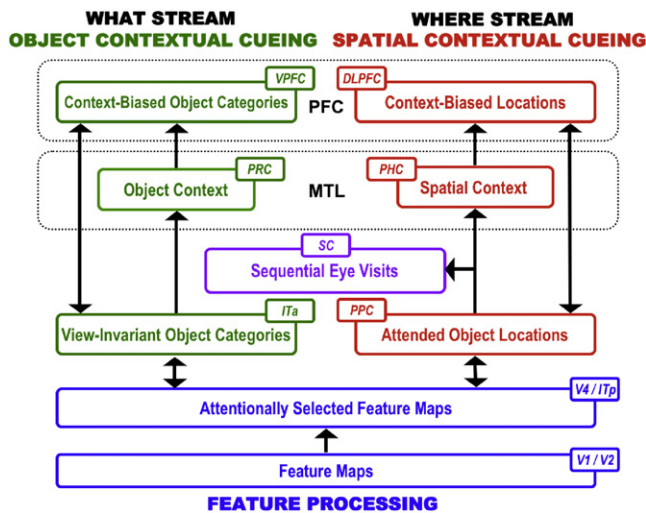


Fig. 15. Macrocircuit of the ARTSCENE search neural model for visual context processing. **V1**= First visual area or primary visual cortex; **V2**= Second visual area; **V4**= Fourth visual area; **PPC**= Posterior parietal cortex; **ITp**= Posterior inferotemporal cortex; **ITa** = Anterior inferotemporal cortex; **MTL**= Medial temporal lobe; **PHC**= Parahippocampal cortex; **PRC**= Perirhinal cortex; **PFC**= Prefrontal cortex; **DLPFC**= Dorsolateral PFC; **VPFC**= Ventral PFC; **SC**= Superior colliculus.

Source: Reprinted with permission from Huang and Grossberg (2010).

in the temporal domain to explain how list chunks resonate with working memory speech items as they are stored through time, thereby driving fast learning of speech and language, dynamically stabilizing these learned memories, and explaining how conscious speech percepts emerge from these resonant states (Grossberg, 1978a, 1986, 2003b; Grossberg & Kazerounian, 2011; Grossberg et al., 1997; Grossberg & Myers, 2000; Kazerounian & Grossberg, submitted for publication).

55.1. Temporal resonances during conscious speech and language perception

Temporal processes such as speech and language provide a particularly clear window into the temporal dynamics with which resonances develop. As in all ART models, a speech or language resonance develops when bottom-up signals that are activated by environmental events interact with top-down expectations, or prototypes, that have been learned from prior experiences. The top-down expectations carry out a matching process that evolves dynamically through time as it selects those combinations of bottom-up features that are consistent with the active learned prototypes while inhibiting those that are not. In this way, an attentional focus concentrates processing on those feature clusters that are deemed important on the basis of past experience. The attended feature clusters, in turn, reactivate the cycle of bottom-up and top-down signal exchange. This reciprocal exchange of signals equilibrates in a resonant state that binds the attended features together. Such resonant states, rather than the activations that are due to bottom-up processing alone, are proposed to be the brain events that regulate fast and stable learning of speech and language, and that give rise to conscious speech and language percepts.

The feedback dynamics of these resonances enable the brain to incorporate both past and future contextual information, often acting over hundreds of milliseconds, into the processing of speech and language, without destroying the perceived temporal order, from past to future, of consciously heard words. Such contextual disambiguation and coherent processing helps humans to understand speech and language during the multi-speaker noisy environments that are characteristic of real-life speech and language experiences.

55.2. Phonemic restoration: how future context influences conscious percepts of past sounds

The conscious ARTWORD (cARTWORD) model of Grossberg and Kazerounian (2011) proposes how the laminar circuits of neocortex are organized to generate conscious speech and language percepts using circuits that are variations of the LAMINART visual circuits and the LIST PARSE cognitive circuits (Fig. 17; see Section 19). cARTWORD simulates how the brain restores linguistic information that is occluded by noise by integrating contextual information over many milliseconds to restore the noise-occluded acoustical signals. It hereby demonstrates how speech and language may be consciously heard in the correct temporal order, despite the influence of contexts that may occur many milliseconds before or after each perceived word. cARTWORD demonstrates such contextual sensitivity by quantitatively simulating a critical example of contextual disambiguation of speech and language; namely, *phonemic restoration* (Miller & Licklider, 1950; Warren, 1970; Warren & Obusek, 1971; Warren & Sherman, 1974; Warren & Warren, 1970).

Consider the following example of phonemic restoration (Warren & Sherman, 1974). Suppose that broadband noise replaces the phonemes /v/ and /b/ in the words “delivery” and “deliberation”, respectively. Despite the initially identical initial portion of these words (“deli-”), if the broadband noise is immediately followed by “ery” or “eration”, listeners hear the /v/ or /b/ as being fully intact and present in the signal. However, if the noise is replaced by silence, then restoration does not occur. Thus, there is a top-down process that can select from the broadband noise those formants which are expected in that context, while suppressing other noise components. But the top-down process, by itself, is merely modulatory, so cannot restore the missing phoneme from silence.

These experimental properties of phonemic restoration match the mathematical properties of top-down matching by the modulatory on-center, off-surround circuits that embody the ART Matching Rule, and which are necessary to solve the stability–plasticity dilemma; that is, to enable fast learning without catastrophic forgetting. Thus, phonemic restoration properties illustrate attentive matching processes that enable speech and language to be learned quickly and stably.

Phonemic restoration also provides an excellent example of a resonance that emerges through time as items are stored in working memory. In particular, why is the noise in “deli-noise-[ery/eration]” not heard before the last portion of the word is even presented? This may be explained by the fact that, if the resonance has not developed fully before the last portion of the word is presented, then this portion can influence the expectations that determine the conscious percept. How then, does the expectation convert the noise in “deli-noise-[ery/eration]” into a percept of [v/-/b]? This occurs due to the top-down matching process that selects expected feature clusters for attentive processing while suppressing unexpected ones. In the “deli-noise-[ery/eration]” example, spectral components of the noise are suppressed that are not part of the expected consonant sound.

This attentive selection process is not merely a process of symbolic inference. It directly influences phonetic percepts. For example, if a reduced set of spectral components is used in the noise, then a correspondingly degraded consonant sound is heard (Samuel, 1981).

A related question concerns how future events can influence past events without smearing over all the events that intervene. This property is related to why, if the /v/ or /b/ in “delivery/deliberation” is replaced by silence, that silence is perceived as silence despite the fact the disambiguating cue would have influenced the percept were these phonemes to be replaced by

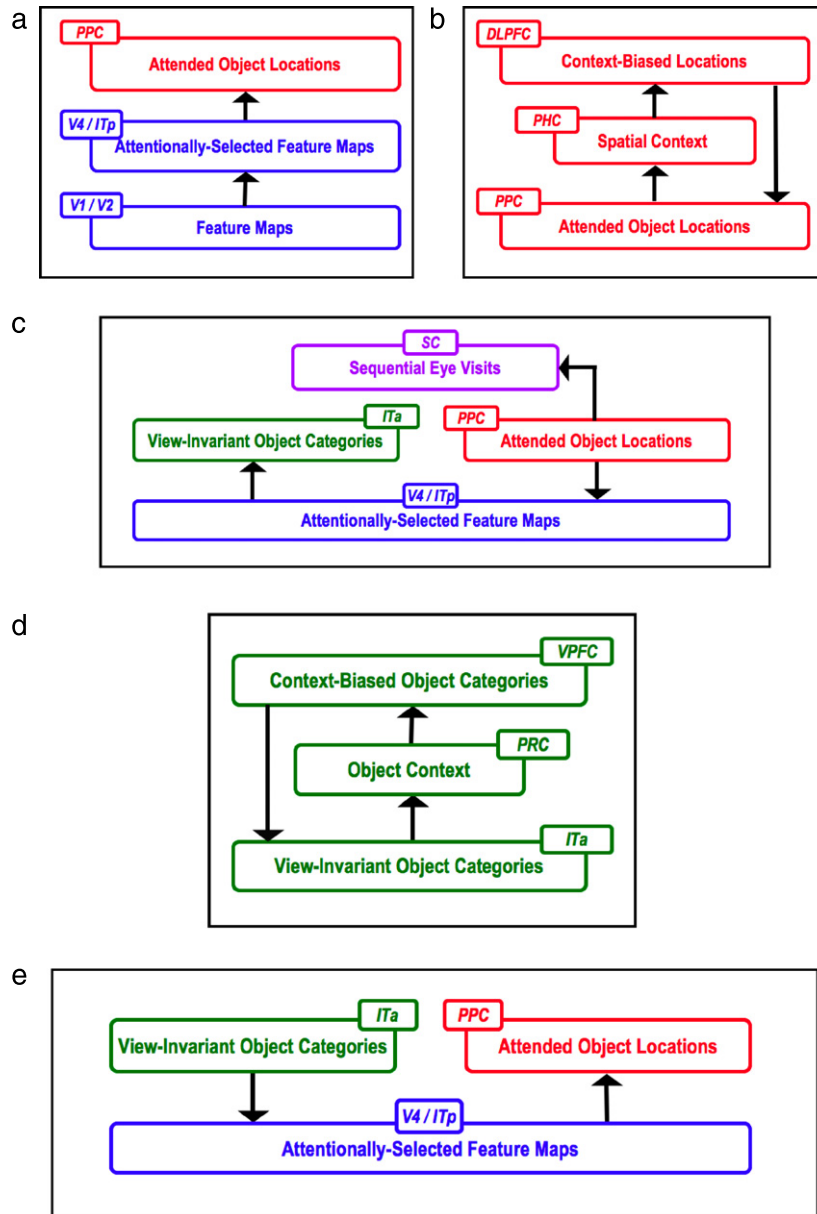


Fig. 16. ARTSCENE search cycle. (a) Step 1: feature maps in cortical areas V1/V2 code different features in the search display. Cortical areas V4/ITp undergo a local competition within each feature map to produce feature-level saliency for each location. Feature-level saliencies are averaged over features into location-level saliencies in PPC. (b) Step 2: the spatial priority map in PPC is cached as a spatial context in PHC, which induces a context-sensitive expectation of target locations in DLPFC. The PPC-PHC-DLPFC-PPC loop then primes possible target locations in PPC via spatial attention. (c) Step 3 (Where to What): PPC drives SC to direct an eye movement, namely overt attention, to the most active location in the spatial priority map. All features at the fixated location are then selected for object recognition in ITa. (d) Step 4: the history of viewed objects is stored in PRC, which builds up a context-induced expectation of target identities in VPFC. The ITa-PRC-VPFC-ITa loop primes possible target categories in ITa via object attention. (e) Step 5 (What to Where): object-attended ITa further primes feature maps in ITp/V4 via feature-based ART attention, and thereby boosts PPC to highlight spatial locations with target-like objects. If the currently fixated and recognized object is not a target, the model goes back to Step 2 for maintaining spatial cueing, and then to Step 3 for inspecting other objects based on the updated priority map in PPC.

Source: Reprinted with permission from Huang and Grossberg (2010).

noise. Here again the nature of the top-down matching process is paramount. This matching process can select feature components that are consistent with its prototype, but it cannot create something out of nothing.

The opposite concern is also of importance. How can sharp word boundaries be perceived even if the sound spectrum that represents the words exhibits no silent intervals between them? The current theory proposes that silence will be heard between words whenever there is a temporal break between the resonances that represent the individual words. In other words, just as conscious speech is a resonant wave, silence is a discontinuity in the rate at which this resonant wave evolves.

There are many examples of the brain's ability to use future events to influence conscious percepts of earlier arriving inputs. An even more striking example of phonemic restoration concerns the fact that "the resolving context may be delayed for two or three, or even more words following the ambiguous word fragment" (Warren & Sherman, 1974, p. 156). In the phrase "[noise]eel is on the —", where the resolving context is given by the last word ("axle", "shoe", "orange" or "table"), listeners "experience the appropriate phonemic restoration ["wheel", "heel", "peel", or "meal"], apparently by storing the incomplete information until the necessary context is supplied so that the required phoneme can be synthesized" (Warren & Warren, 1970, p. 32). In this example,

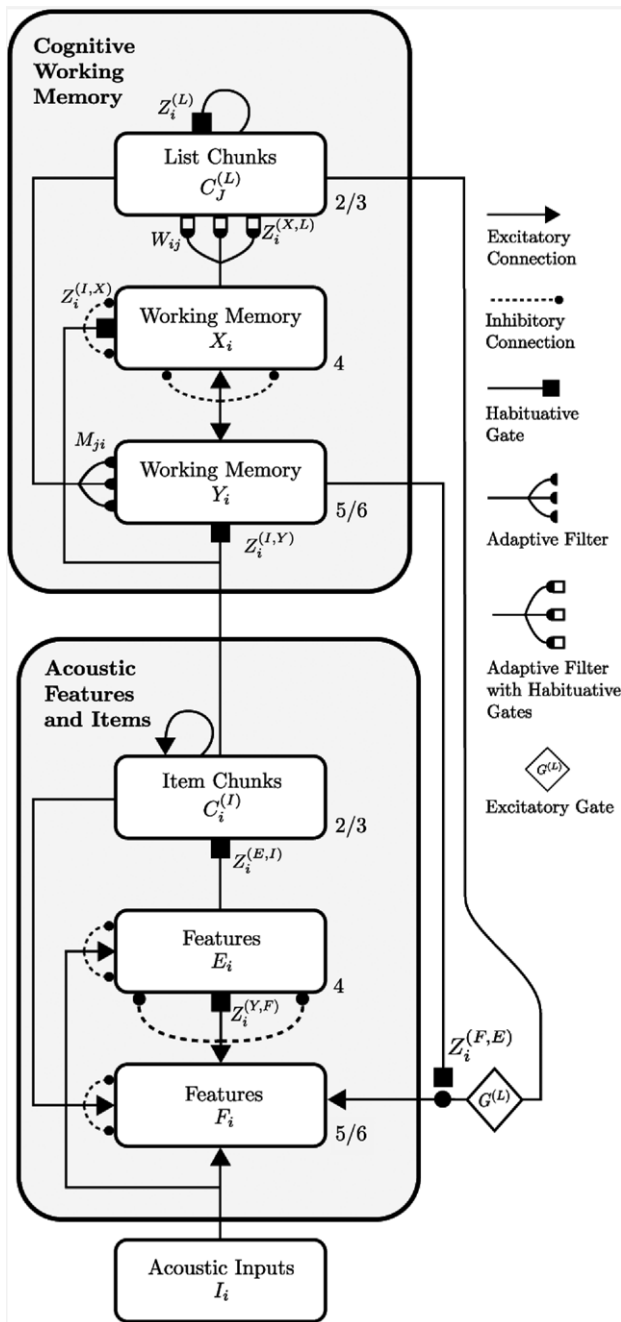


Fig. 17. The cARTWORD model describes a hierarchy of levels responsible for the processes involved in speech and language perception. Each level is organized into laminar cortical circuits, wherein deep layers (6 and 4) are responsible for processing and storing inputs, and superficial layers (2/3) are proposed to group distributed patterns across these deeper layers into unitized representations. The lowest level is responsible for processing acoustic features (cell activities F_i and E_i) and items (cell activities $C_i^{(I)}$), whereas the higher level is responsible for storing of sequences of acoustic items in working memory (activities Y_i and X_i), and representing these stored sequences of these items as unitized, context-sensitive representations by list chunks (activities $C_J^{(L)}$) in a network, called a *masking field*, that is capable of selectively representing lists of variable length. Source: Reprinted with permission from Grossberg and Kazerounian (2011).

the last word contextually disambiguates how the noisy part of the first word is restored. Once again, the later occurring top-down effect of meaning influences the phonemic structure which is consciously perceived as coming earlier in time.

Yet another example of a backwards-in-time effect illustrates how the individual words that we consciously hear may combine

sounds that are separated by 100 ms of silence and that are temporally continuous to more than one word. For example, increasing the silence duration between the words “gray” and “chip” in the phrase “gray chip” may result in the conscious percept “great chip”. Why should separating two words by more silence, which might have been expected to make them more distinctive, increase the probability that the fricative / f / in “chip” will leap backwards-in-time over the intervening silence to join the word “gray” and make it consciously sound like “great”? Moreover, at appropriate durations of the fricative noise / f /, listeners reliably perceive “gray” as “great” and “chip” as “ship”, even at the highest tested silence durations of 100 ms (Repp, Liberman, Eccardt, & Pesetsky, 1978). How does a future noise in “chip” leap backwards-in-time over 100 ms of silence to replace the word “gray” by a conscious percept of “great”? Such examples illustrate that our conscious percepts of words may coherently bind together acoustic sounds that are not necessarily presented contiguously in time. Grossberg and Myers (2000) developed the ARTWORD model, a non-laminar precursor of the cARTWORD model, to quantitatively simulate such percepts using recurrent ART dynamics between an Item and Order working memory and a *masking field* list chunking level that can code words of variable length.

56. Item-list resonance and phonemic restoration: conscious speech is a resonant wave

Due to the conceptual importance of phonemic restoration properties towards understanding ART dynamics during conscious speech perception, they are now summarized all together: During phonemic restoration, a phoneme deleted from a speech stream is perceptually restored when it is replaced by broadband noise, even when the disambiguating context occurs after the phoneme was presented. If the noise is replaced by silence, then restoration does not occur. As noted above, this combination of properties illustrates the ART Matching Rule operating during speech and language perception and understanding. In particular, the top-down expectation has a modulatory on-center. It can select formants that are in its prototype from the broadband noise and incorporate them into a resonance, but cannot by itself fire its target cells, unless the modulatory top-down on-center is converted into a driving on-center by basal ganglia control (Sections 7 and 42), as during internal thought which uses language.

The operation of the ART Matching Rule here reflects the fact that speech and language may be learned quickly, without catastrophic forgetting. The cARTWORD model describes how ART-modulated laminar circuits within a hierarchy of cortical processing stages (Fig. 17) may interact to generate a conscious speech percept that is embodied by a resonant wave of activation that occurs between acoustic features, acoustic item chunks, and list chunks. Such a resonance is often called an *item-list resonance*, for short. In cARTWORD, conscious speech is a resonant wave. Perceived silence is a temporal discontinuity in the rate at which this wave evolves through time. Signals from activated list chunks in the prefrontal cortex (Fig. 16) can open basal ganglia gates (Section 44) – that is, a context-sensitive “decision” is made – thereby activating an item-list resonance that links the entire coding hierarchy, and allowing speech to be consciously heard in the correct temporal order, from past to future, even when what is heard in the present depends upon future context.

57. Resonant auditory streams during source segregation: the auditory continuity illusion

At a lower level of auditory processing than speech and language understanding, ART mechanisms help to select and focus

attention upon and track individual acoustic sources, whether these are human speakers or inanimate sounds in the environment. The existence of resonant dynamics on multiple levels of the auditory system, just as in the visual system (e.g., Section 31) clarifies how conscious auditory percepts may be generated by the synchronous dynamics of multiple brain levels acting together.

Multiple sound sources often contain harmonics that overlap and may be degraded by environmental noise. The auditory system is capable of teasing apart these sources into distinct mental sound sources, or streams, even in situations where multiple sound sources share some of the same frequencies. In accomplishing such an auditory scene analysis (Bregman, 1990), the brain solves the so-called *cocktail party problem*, which arises every time one tries to listen to a friend at a noisy cocktail party or other social occasion. How does the brain separate overlapping frequencies that are due to different acoustic sources into separate streams, so that higher auditory processes, such as those involved in speech and language, can understand the meaning of these streams? How do these streams achieve the coherence that is needed to bridge across spectrally ambiguous durations, including those wherein noise occludes the acoustic signals? ART dynamics have proved useful towards explaining how this may happen.

Are there experimental data about auditory streaming that illustrate ART dynamics as vividly as phonemic restoration does for speech? The *auditory continuity illusion* provides one class of examples (Bregman, 1990). This illusion occurs when a steady tone occurs both before and after a burst of broadband noise. Under appropriate temporal and amplitude conditions, a percept is heard in which the tone appears to continue through the noise, and the residual noise sounds like it occurs in a different stream. The backwards-in-time process that determines this illusion is made clear by the fact that, if a subsequent tone does not follow the noise burst, then the tone does not continue through the noise. Moreover, in the absence of noise, even if a tone is played before and after a silence duration, the silence duration is perceived.

These are properties of the ART Matching Rule acting at the level of auditory scene analysis. The fact that the tone is not completed over silence is easily understood, since top-down expectations are modulatory. But why does the tone continue through the noise when the noise is followed by a subsequent tone, but not otherwise? How does the brain operate backwards-in-time to complete the tone through the noise only if there is a future tone? This property may be explained by how the resonance develops through time: It takes awhile for the resonance to develop in response to the first tone, but much less time for the second tone to boost the resonance and thereby enable it to continue selecting the tone from the noise. The duration of the noise cannot be too long for this future boost to have its influence before the noise interval is perceived.

Given that the ART Matching Rule is at work, what is being learned? As described in Section 58, representations of the pitch of a sound – that is, pitch categories – are learned (among other things) using ART bottom-up filters and top-down learned expectations and attentional matching. Each pitch category can group the sound frequencies corresponding to the pitch of a single acoustic source.

58. Spectral-pitch resonances select auditory streams

The proposal that resonant dynamics helps to select auditory streams raises another basic question as well: What is resonating? In the case of visual perception, surface-shroud resonances support percepts of conscious visual qualia (what we “see”), and object-feature resonances represent the objects to which we attend (what we “know”). In the case of conscious speech perception, the CARTWORD model illustrates how an item-list resonance

between acoustic features, items in working memory, and list chunks can support conscious percepts of speech (Sections 55–56). The acoustic features that resonate in an item-list resonance include frequency-dependent spectral components. How are these acoustic features bound together to represent the heard sounds of acoustic sources?

The ARTSTREAM model Grossberg et al. (2004, Fig. 18)) proposes how a *spectral-pitch resonance* can support both pitch category learning and auditory streaming. ARTSTREAM describes, in particular, how the frequency components that correspond to a given acoustic source may be coherently grouped together into a distinct auditory stream based on pitch and spatial location cues. The model also clarifies how multiple streams may be distinguished and separated by the brain. These streams are formed as spectral-pitch resonances that emerge through feedback interactions between frequency-specific spectral representations of a sound source and categorical representations of pitch. Resonance provides the coherence across frequency, pitch, and time that allows one voice or instrument to be tracked through a noisy multiple source environment. Various popular methods of stream segregation in engineering and technology, such as independent component analysis, do not exhibit this sort of coherence (Comon, 1994; Hyvärinen & Oja, 2000).

In ARTSTREAM, acoustic input signals are first preprocessed by multiple stages that together comprise the Spatial Pitch NETWORK (SPINET) model (Cohen, Grossberg, & Wyse, 1995). SPINET has been used to quantitatively simulate many human psychophysical data about pitch perception. Critically, SPINET provides a *spatial* representation of frequency and pitch, rather than just an autocorrelation measure of these quantities. This spatial representation can be exploited by the brain to break sounds up into multiple streams whose frequencies are grouped by pitch and location cues to separate and resonantly track multiple sound sources.

The final SPINET preprocessing stage is called the Energy Measure layer (Fig. 18). Each of its frequency-specific activities generates output signals that are broadcast in a one-to-many manner to an entire strip of cells in the Spectral Stream layer. This strip is perpendicular to (or at least cuts across) the spatial layout of the frequency spectrum. Since each frequency is represented by an entire strip of cells, the frequency spectrum is redundantly represented across the strips. Multiple streams can be chosen within these strips. The activity within the Spectral Stream layer is contrast-enhanced by a competition that acts within each strip – that is, across streams – at each frequency. Only one stream’s frequency will be able to survive the competition at any time, thereby realizing “exclusive allocation” of that frequency to a single stream (Bregman, 1990). However, which stream can process a given frequency depends on top-down feedback from the Pitch Stream layer.

The Pitch Stream layer gets activated by the Spectral Stream layer via a bottom-up adaptive filter. The cells in this filter are activated by the frequency harmonics of the currently active sounds due to the way in which harmonics are peripherally processed. The adaptive filter thus acts like a “harmonic sieve”. Because harmonics of a single sound highly correlate through time, the adaptive filter can select and learn pitch-sensitive category cells at the Pitch Stream layer. The Pitch Streams also compete within a strip. This competition is asymmetric across streams to break the following kind of symmetry: Other things being equal, if the same frequency spectrum is redundantly activated across multiple streams, its pitch category in the first stream will win the asymmetric competition and inhibit the corresponding pitch categories in the other streams.

The winning pitch category reads out a top-down expectation that obeys the ART Matching Rule. In particular, the top-down

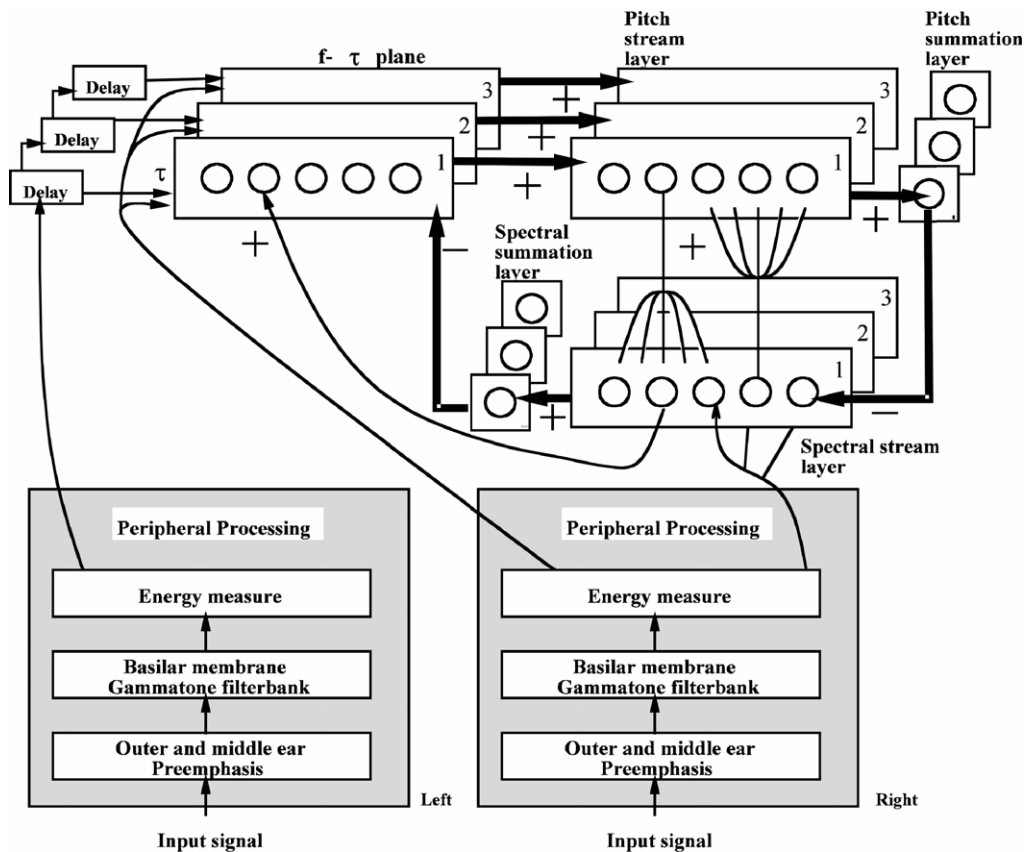


Fig. 18. In the ARTSTREAM model, both pitch and spatial location cues cooperate to select auditory streams. See text for details.
Source: Reprinted with permission from Grossberg et al. (2004).

on-center activates the harmonics of the winning pitch in the corresponding spectral stream. The winning pitch category also activates a nonspecific inhibitory interneuron via the Pitch Summation layer, which inhibits all the frequencies in the spectral stream. Only frequencies that are harmonics of the winning pitch can survive this competition. In this way, the winning pitch in a pitch stream selects and amplifies those harmonics in its stream that are consistent with it. When this happens, the amplified spectral frequencies can inhibit the same frequencies in other streams, thereby achieving “exclusive allocation” of frequencies to one stream at a time (Bregman, 1990).

The first winning pitch also suppresses frequencies at the Spectral Stream layer that are not among its harmonics. These frequencies can then no longer inhibit the same frequencies in other streams. The disinhibited frequencies can then select a pitch in other streams. The asymmetric competition within the Pitch Stream layer then enables a second stream, and only that stream, to represent the pitch of a different sound source. The process continues cyclically in this way. Multiple simultaneously occurring spectral-pitch resonances can hereby emerge.

The stream selection process is also modulated by information about source position using the $f - \tau$ plane (Colburn, 1973, 1977), where individual frequencies f are assigned to a spatial location τ that is defined by a radial direction. See Grossberg et al. (2004) for details.

ARTSTREAM simulates streaming data from psychophysical grouping experiments, such as how a tone sweeping upwards in frequency creates a bounce percept by grouping with a downward sweeping tone due to proximity in frequency, even if noise replaces the tones at their intersection point. Illusory auditory percepts are also simulated, such as the auditory continuity illusion of a tone continuing through a noise burst even if the tone is not

present during the noise, and the scale illusion of Deutsch whereby downward and upward scales presented alternately to the two ears are regrouped based on frequency proximity, leading to a bounce percept. Although these simulations are only a beginning in tackling the rich literature on streaming, they illustrate how pitch categories can help to define multiple coherent streams, and how the streaming, speech, and language levels can all resonate together in generating percepts of conscious speech and language meaning.

59. Strip maps: auditory streaming and speaker-normalized language learning

Auditory signals of speech are speaker-dependent, but representations of language meaning are speaker-independent. The transformation from speaker-dependent to speaker-independent language representations enables speech to be learned and understood from different speakers. The process whereby a speaker-independent representation is formed is called *speaker normalization*. Speaker normalization is a critical transformation between the selection and attentive tracking of an auditory stream that represents a speaker and the ability to understand what the speaker is saying.

In particular, speaker normalization allows language learning to get started. It enables a baby to begin to imitate sounds from adult speakers, notably parents whose spoken frequencies differ significantly from those that the baby can babble. A *circular reaction* (Piaget, 1963) from endogenously babbled to heard sounds enables a baby to learn a map between the auditory representations of its own heard babbled sounds to the motor commands that caused them. Speaker normalization enables sounds from adult caretakers to be filtered by this learned map and to thereby enable the baby

to begin to imitate and refine heard sounds in its own language productions. Learning in such an imitative map needs to remain active for many years in order to enable an individual's changing voice through puberty and adulthood to continue to activate and update this map.

Speaker normalization also enables language meanings that were learned from one teacher's voice to be readily understood when uttered by another speaker. More generally, speaker normalization helps the brain to overcome a combinatorial explosion that would otherwise occur if the brain needed to store every instance of every speaker utterance in order to understand language meaning.

Neural models of how such circular reactions are learned have been used to explain many data about speech production and arm reaching behaviors (Bullock, Grossberg, & Guenther, 1993; Cohen, Grossberg, & Stork, 1988; Grossberg, 1978a; Guenther, 1995; Guenther, Ghosh, & Tourville, 2006). Indeed, the DIRECT arm reaching model of Bullock and Grossberg (1988) and the DIVA speech production model of Guenther (1995) use computationally homologous model circuits.

How does speaker normalization occur in the brain? This question faces us with the same kind of dilemma about how evolution could be smart enough to design a working memory. Section 50 proposed that working memories are just specialized recurrent shunting on-center off-surround networks, which are ubiquitous in the brain because of the need to solve the noise-saturation dilemma. A specialized recurrent shunting network with a rehearsal wave and self-inhibitory interneurons can act like a working memory. Is there a broadly used design that the brain can also exploit to accomplish speaker normalization?

The Neural Normalization Network, or NormNet, model (Ames & Grossberg, 2008) proposes that speaker normalization, just like auditory streaming, depends upon strip maps and asymmetric competition. Thus, the requirements of speaker normalization may have arisen during evolution from an elaboration and specialization of the more primitive requirements of auditory streaming. This would clarify how speaker normalization can transform source data, as soon as it is broken into separate streams, for purposes of speech and language classification and meaning extraction, yet how the frequency content of the streams can be preserved for purposes of speaker identification. Towards this former goal, NormNet learned speaker-normalized vowel categories using an ART classifier. Its simulations used synthesized steady-state vowels from the Peterson and Barney (1952) vowel database, and achieved accuracy rates similar to those achieved by human listeners.

By virtue of speaker normalization, the phonetic processing within item-list resonances would not be expected to be as sensitive to frequency harmonics as are the spectral-pitch resonances which separate and track the speaker sources that phonetic processing attempts to understand. Consistent with this expectation, Remez and his colleagues have shown that harmonics are more important during auditory streaming than during phonetic perception (Remez, 2003; Remez, Pardo, Piorkowski, & Rubin, 2001; Remez, Rubin, Berns, Pardo, & Lang, 1994).

Are there other brain processes that use strip maps? May strip maps be thought of as a design principle that is opportunistically specialized for different purposes? The SpaN model of Grossberg and Repin (2003) shows how strip maps may be used to represent place-value numbers in the brain; e.g., 10, 100, 1000, and so on; see Section 53.

In all these cases, a given input feature can activate a strip, or cluster, of cells that can then be differentiated into subclusters which represent a transformation of the features that are represented at the base of the strip, whether this transformation generates a different auditory stream built out of the same auditory

frequencies, or a normalized frequency representation of a speaker, or a different place-value number with the same base number (e.g., "24" instead of "4"). Indeed, redundant representations of prescribed features are ubiquitous in the brain. They enable high-dimensional information to be multiplexed within the thin, approximately two-dimensional (except for laminar specialization) maps in the neocortex. Perhaps the most familiar example of this redundancy is the ocular dominance columns in visual cortical area V1 that are sensitive to inputs from one eye at each location, but whose constituent cells may be tuned selectively to differently oriented visual features that are seen by that eye at the corresponding location.

60. Conclusion

60.1. ART is not a theory of everything: complementary computing and inhibitory matching

This article provides a heuristic summary of various contributions to ART as a cognitive and neural theory. The summary suggests how ART has enhanced our understanding of many processes of perception, attention, cognition, and cognitive-emotional dynamics. ART has also clarified brain processes that are not summarized here. For example, the 3D FORMOTION model uses ART top-down expectations, hypothesized to occur from cortical area MST to cortical area MT, to select the perceived motion direction of an object, and to thereby solve the aperture problem (e.g., Berzhanskaya, Grossberg, and Mingolla (2007) and Grossberg, Mingolla, and Viswanathan (2001)). The SACCART model clarifies how the deeper layers of the superior colliculus select an attended target position for a saccadic eye movement from auditory, visual, and planned movement signals (Grossberg, Roberts, Aguilar, & Bullock, 1997). Despite this diversity of applications, ART is not a "theory of everything". As noted in Section 3 and Fig. 1, ART-based learning of object representations in the What stream uses excitatory matching and match-based learning to solve the stability-plasticity dilemma. Where stream learning, in contrast, often uses computationally complementary processes of inhibitory matching and mismatch-based learning. Correspondingly, the What stream learns object representations that strive to become positionally-invariant, whereas the Where stream represents the positions and actions that enable recognized objects to be manipulated.

Inhibitory matching often occurs between brain representations of a *target position* and the *present position* of a motor effector. Inhibitory matching computes a *difference vector* that represents the distance and direction of an intended movement. The difference vector is volitionally gated (see Sections 12 and 44) by a basal ganglia GO signal that determines when and how fast the movement will occur (Bullock et al., 1998; Bullock & Grossberg, 1988).

During motor learning, a difference vector can also generate error signals when the same target position and present position are encoded but not properly calibrated. These error signals activate a form of mismatch learning that eliminates the mismatch through time by recalibrating system maps and gains to be consistent. Neural models predict how mismatch learning may tune spatial representations and adaptive motor gains in basal ganglia, cerebellum, motor cortex, parietal cortex, and prefrontal cortex (Brown et al., 1999, 2004; Fiala et al., 1996; Grossberg & Paine, 2000; Guenther, Bullock, Greve, & Grossberg, 1994). Models that carry out computation and learning by difference vectors are often called Adaptive Vector Integration to Endpoint (aVITE) or Vector Associative Map (VAM) models (Gaudio & Grossberg, 1991, 1992).

60.2. Joining ART and VAM: self-stabilizing expertise in evolving bodies

Inhibitory matching and mismatch learning exhibit catastrophic forgetting. However, catastrophic forgetting is a good property for learning sensory–motor maps and gains. In particular, it would be maladaptive to remember for life the maps and gains whereby our brains controlled our infant limbs. Continual recalibration of maps and gains enables us to efficiently control our changing bodies.

The proposal that “all conscious states are resonant states” clarifies why spatial/motor, also called procedural, processes are unconscious: the inhibitory matching process that supports spatial and motor processes cannot lead to resonance. This distinction provides a simple mechanistic explanation of why procedural memories are not conscious, whereas declarative memories can be conscious, as noted in Section 3.

In summary, perceptual/cognitive processes often use ART-like excitatory matching and match-based learning to create self-stabilizing representations of objects and events that enable us to gain increasing expertise as we learn about the world. Complementary spatial/motor processes often use VAM-like inhibitory matching and mismatch-based learning to continually update spatial maps and sensory–motor gains to compensate for bodily changes throughout life. Together these complementary predictive and learning mechanisms create a self-stabilizing perceptual/cognitive front end for intelligently manipulating the more labile spatial/motor processes which enable our changing bodies to act effectively upon a changing world.

Acknowledgment

This research was supported in part by the SyNAPSE program of DARPA (HR0011-09-C-0001).

References

- Abbott, L. F., Varela, K., Sen, K., & Nelson, S. B. (1997). Synaptic depression and cortical gain control. *Science*, 275, 220–223.
- Agam, Y., Bullock, D., & Sekuler, R. (2005). Imitating unfamiliar sequences of connected linear motions. *Journal of Neurophysiology*, 94, 2832–2843.
- Agam, Y., Galperin, H., Gold, B. J., & Sekuler, R. (2007). Learning to imitate novel motion sequences. *Journal of Vision*, 7, <http://dx.doi.org/10.1167/7.5.1>.
- Aggleton, J. P. (1993). The contribution of the amygdala to normal and abnormal emotional states. *Trends in Neurosciences*, 16, 328–333.
- Ahissar, M., & Hochstein, S. (1993). Attentional control of early perceptual learning. *Proceedings of the National Academy of Sciences USA*, 90, 5718–5722.
- Akhbardeh, A., Junnila, S., Koivistoinen, T., & Varri, A. (2007). An intelligent ballistocardiographic chair using a novel SF-ART neural network and biorthogonal wavelets. *Journal of Medical Systems*, 31, 69–77.
- Ames, H., & Grossberg, S. (2008). Speaker normalization using cortical strip maps: a neural model for steady state vowel categorization. *Journal of the Acoustical Society of America*, 124, 3918–3936.
- Amir, Y., Harel, M., & Malach, R. (1993). Cortical hierarchy reflected in the organization of intrinsic connections in macaque monkey visual cortex. *Journal of Comparative Neurology*, 334, 19–46.
- Amis, G., & Carpenter, G. (2007). Default ARTMAP 2. In *Proceedings of the international joint conference on neural networks, IJCNN'07*. (pp. 777–782). Orlando, Florida: IEEE Press.
- Amis, G., & Carpenter, G. (2009). Self-supervised ARTMAP. *Neural Networks*, 23, 265–282.
- Amis, G., Carpenter, G.A., Ersoy, B., & Grossberg, S. (2009). Cortical learning of recognition categories: a resolution of the exemplar vs. prototype debate (submitted for publication).
- Anagnostopoulos, G.C., & Georgopoulos, M. (2000). Hypersphere ART and ARTMAP for unsupervised and supervised incremental learning. In *Neural networks, proceedings of the IEEE-INNS-ENNS international joint conference on neural networks*, vol. 6 (pp. 59–64).
- Anton-Rodriguez, M., Diaz-Pernas, F. J., Diez-Higuera, J. F., Martinez-Zarzuela, M., Gonzalez-Ortega, D., & Boto-Giralda, D. (2009). Recognition of coloured and textured images through a multi-scale neural architecture with orientational filtering and chromatic diffusion. *Neurocomputing*, 72, 3713–3725.
- Atri, A., Sherman, S., Norman, K. A., Kirchhoff, B. A., Nicolas, M. M., Greicius, M. D., Cramer, S. C., Breiter, H. C., Hasselmo, M. E., & Stern, C. E. (2004). Blockade of central cholinergic receptors impairs new learning and increases proactive interference in a word paired-associate memory task. *Behavioral Neuroscience*, 118, 223–236.
- Auer, E. T., Jr., & Luce, P. A. (2008). Probabilistic phonotactics in spoken word recognition. In D. B. Pisoni, & R. E. Remez (Eds.), *The handbook of speech perception*. Wiley Online Library, <http://dx.doi.org/10.1002/9780470757024.ch25>.
- Averbeck, B. B., Chafee, M. V., Crowe, D. A., & Georgopoulos, A. P. (2002). Parallel processing of serial movements in prefrontal cortex. *Proceedings of the National Academy of Sciences, USA*, 99(20), 13172–13177.
- Averbeck, B. B., Crowe, D. A., Chafee, M. V., & Georgopoulos, A. P. (2003a). Neural activity in prefrontal cortex during copying geometrical shapes, I, single cells encode shape, sequence, and metric parameters. *Experimental Brain Research*, 150, 127–141.
- Averbeck, B. B., Crowe, D. A., Chafee, M. V., & Georgopoulos, A. P. (2003b). Neural activity in prefrontal cortex during copying geometrical shapes, II, decoding shape segments from neural ensembles. *Experimental Brain Research*, 150, 142–153.
- Baldi, P., & Hornik, K. (1989). Neural networks and principal component analysis: learning from examples and local minima. *Neural Networks*, 2, 53–58.
- Bar, M., Tootell, R. B. H., Schacter, D. L., Greve, D. N., Fischl, B., Mendola, J. D., Rosen, B. R., & Dale, A. M. (2001). Cortical mechanisms specific to explicit object recognition. *Neuron*, 29, 529–535.
- Barbas, H. (1995). Anatomic basis of cognitive–emotional interactions in the primate prefrontal cortex. *Neuroscience and Biobehavioral Reviews*, 19, 499–510.
- Barbas, H. (2000). Connections underlying the synthesis of cognition, memory and emotion in primate prefrontal cortices. *Brain Research Bulletin*, 52(5), 319–330.
- Bazhenov, M., Timofeev, I., Steriade, M., & Sejnowski, T. J. (1998). Computational models of thalamocortical augmenting responses. *Journal of Neuroscience*, 18, 6444–6465.
- Baxter, M. G., Parker, A., Lindner, C. C., Izquierdo, A. D., & Murray, E. A. (2000). Control of response selection by reinforcer value requires interaction of amygdala and orbital prefrontal cortex. *Journal of Neuroscience*, 20, 4311–4319.
- Ben-Shahar, O., & Zucker, S. (2004). Sensitivity to curvatures in orientation-based texture segmentation. *Vision Research*, 44, 257–277.
- Berger, T. W., Berry, S. D., & Thompson, R. F. (1986). Role of the hippocampus in classical conditioning of aversive and appetitive behaviors. In R. L. Isaacson, & K. H. Pribram (Eds.), *The hippocampus*, vol. 4 (pp. 203–239). New York: Plenum Press.
- Berke, J. D., Hettrick, V., Breck, J., & Green, R. W. (2008). Transient 23–30-Hz oscillations in mouse hippocampus during exploration of novel environments. *Hippocampus*, 18, 519–529.
- Berzhanskaya, J., Grossberg, S., & Mingolla, E. (2007). Laminar cortical dynamics of visual form and motion interactions during coherent object motion perception. *Spatial Vision*, 20, 337–395.
- Bhatt, R., Carpenter, G., & Grossberg, S. (2007). Texture segregation by visual cortex: perceptual grouping, attention, and learning. *Vision Research*, 47, 3173–3211.
- Bi, G. Q., & Poo, M. (2001). Synaptic modification by correlated activity: Hebb's postulate revisited. *Annual Review of Neuroscience*, 24, 139–166.
- Botly, L. C., & De Rosa, E. (2007). Cholinergic influences on feature binding. *Behavioral Neuroscience*, 121, 264–276.
- Botly, L. C., & De Rosa, E. (2009). Cholinergic deafferentation of the neocortex using 192 igg-saporin impairs feature binding in rats. *The Journal of Neuroscience*, 29, 4120–4130.
- Bradski, G., Carpenter, G. A., & Grossberg, S. (1992). Working memory networks for learning temporal order with application to 3-D visual object recognition. *Neural Computation*, 4, 270–286.
- Bradski, G., Carpenter, G. A., & Grossberg, S. (1994). STORE working memory networks for storage and recall of arbitrary temporal sequences. *Biological Cybernetics*, 71, 469–480.
- Brannon, N. G., Seiffert, J. E., Draelos, T. J., & Wunsch, D. C. II (2009). Coordinated machine learning and decision support for situation awareness. *Neural Networks*, 22, 316–325.
- Bregman, A. S. (1990). *Auditory scene analysis: the perceptual organization of sound*. Cambridge, MA: MIT Press.
- Brown, J., Bullock, D., & Grossberg, S. (1999). How the basal ganglia use parallel excitatory and inhibitory learning pathways to selectively respond to unexpected rewarding cues. *Journal of Neuroscience*, 19, 10502–10511.
- Brown, J. W., Bullock, D., & Grossberg, S. (2004). How laminar frontal cortex and basal ganglia circuits interact to control planned and reactive saccades. *Neural Networks*, 17, 471–510.
- Brown, J. M., & Denney, H. I. (2007). Shifting attention into and out of objects: evaluating the processes underlying the object advantage. *Perception & Psychophysics*, 69, 606–618.
- Buchanan, F. M., & Bitterman, M. E. (1998). Learning in honeybees as a function of amount and frequency of reward. *Animal Learning and Behavior*, 16, 247–255.
- Buffalo, E. A., Fries, P., Landman, R., Buschman, T. J., & Desimone, R. (2011). Laminar differences in gamma and alpha coherence in the ventral stream. *Proceedings of the National Academy of Sciences*, 108, 11262–11267.
- Bullier, J., Hupé, J. M., James, A., & Girard, P. (1996). Functional interactions between areas V1 and V2 in the monkey. *Journal of Physiology (Paris)*, 90, 217–220.
- Bullock, D., Cisek, P., & Grossberg, S. (1998). Cortical networks for control of voluntary arm movements under variable force conditions. *Cerebral Cortex*, 8, 48–62.

- Bullock, D., & Grossberg, S. (1988). Neural dynamics of planned arm movements: emergent invariants and speed-accuracy properties during trajectory formation. *Psychological Review*, 95, 49–90.
- Bullock, D., Grossberg, S., & Guenther, F. H. (1993). A self-organizing neural model of motor equivalent reaching and tool use by a multijoint arm. *Journal of Cognitive Neuroscience*, 5, 408–435.
- Burgess, N. (1995). A solvable connectionist model of immediate recall of ordered lists. In G. Tesauró, D. S. Touretzky, & T. K. Leen (Eds.), *Advances in neural information processing systems*, vol. 7. Cambridge, MA: MIT Press.
- Burgess, N., & Hitch, G. J. (1992). Toward a network model of the articulatory loop. *Journal of Memory and Language*, 31, 429–460.
- Buschman, T. J., & Miller, E. K. (2007). Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science*, 315, 1860–1862.
- Buschman, T. J., & Miller, E. K. (2009). Serial, covert shifts of attention during visual search are reflected by the frontal eye fields and correlated with population oscillations. *Neuron*, 63, 386–396.
- Cai, Y., Wang, J.-Z., Tang, Y., & Yang, Y.-C. (2011). An efficient approach for electric load forecasting using distributed ART (adaptive resonance theory) & HS-ARTMAP (Hyper-spherical ARTMAP network) neural network. *Energy*, 36, 1340–1350.
- Cano-Izquierdo, J.-M., Almonacid, M., Pinzolas, M., & Ibarrola, J. (2009). dFasArt: dynamic neural processing in FasArt model. *Neural Networks*, 22, 479–487.
- Cao, Y., & Grossberg, S. (2005). A laminar cortical model of stereopsis and 3D surface perception: closure and da Vinci stereopsis. *Spatial Vision*, 18, 515–578.
- Cao, Y., Grossberg, S., & Markowitz, J. (2011). How does the brain rapidly learn and reorganize view- and positionally-invariant object representations in inferior temporal cortex? *Neural Networks*, 24, 1050–1061.
- Caputo, G., & Guerra, S. (1998). Attentional selection by distractor suppression. *Vision Research*, 38, 669–689.
- Carpenter, G. A. (1997). Distributed learning, recognition, and prediction by ART and ARTMAP neural networks. *Neural Networks*, 10, 1473–1494.
- Carpenter, G. A. (2003). Default ARTMAP. In *Proceedings of the international joint conference on neural networks, IJCNN'03* (pp. 1396–1401).
- Carpenter, G. A., & Gaddam, S. C. (2010). Biased ART: a neural architecture that shifts attention toward previously disregarded features following an incorrect prediction. *Neural Networks*, 23, 435–451.
- Carpenter, G. A., & Grossberg, S. (1987). A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing*, 37, 54–115.
- Carpenter, G. A., & Grossberg, S. (1990). ART 3: hierarchical search using chemical transmitters in self-organizing pattern recognition architectures. *Neural Networks*, 3, 129–152.
- Carpenter, G. A., & Grossberg, S. (1991). *Pattern recognition by self-organizing neural networks*. Cambridge, MA: MIT Press.
- Carpenter, G. A., & Grossberg, S. (1993). Normal and amnesic learning, recognition, and memory by a neural model of cortico-hippocampal interactions. *Trends in Neurosciences*, 16, 131–137.
- Carpenter, G. A., Grossberg, S., Markuzon, N., Reynolds, J. H., & Rosen, D. B. (1992). Fuzzy ARTMAP: a neural network architecture for incremental supervised learning of analog multidimensional maps. *IEEE Transactions on Neural Networks*, 3, 698–713.
- Carpenter, G. A., Grossberg, S., & Reynolds, J. H. (1991). ARTMAP: supervised real-time learning and classification of nonstationary data by a self-organizing neural network. *Neural Networks*, 4, 565–588.
- Carpenter, G. A., Grossberg, S., & Rosen, D. B. (1991). Fuzzy ART: fast stable learning and categorization of analog patterns by an adaptive resonance system. *Neural Networks*, 4, 759–771.
- Carpenter, G. A., Martens, S., & Ogas, O. J. (2005). Self-organizing information fusion and hierarchical knowledge discovery: a new framework using ARTMAP neural networks. *Neural Networks*, 18, 287–295.
- Carpenter, G. A., Milenova, B. L., & Noeske, B. W. (1998). Distributed ARTMAP: a neural network for fast distributed supervised learning. *Neural Networks*, 11, 793–813.
- Carpenter, G. A., & Ravindran, A. (2008). Unifying multiple knowledge domains using the ARTMAP information fusion system. In *Proceedings of the 11th international conference on information fusion*. Cologne, Germany. June 30–July 3.
- Carrasco, M., Penpeci-Talgar, C., & Eckstein, M. (2000). Spatial covert attention increases contrast sensitivity across the CSF: support for signal enhancement. *Vision Research*, 40, 1203–1215.
- Caudell, T. P. (1992). Hybrid optoelectronic adaptive resonance theory neural processor, ART 1. *Applied Optics*, 31, 6220–6229.
- Caudell, T. P., Smith, S. D. G., Johnson, G. C., Wunsch, D. C. II, & Escobedo, R. (1991). An industrial application to neural networks to reusable design. *Neural Networks, International Joint Conference on Neural Networks*, 2, 919.
- Cauler, L. J., & Connors, B. W. (1994). Synaptic physiology of horizontal afferents to layer I in slices of rat SI cortex. *Journal of Neuroscience*, 14, 751–762.
- Chao, H.-C., Hsiao, C.-M., Su, W.-S., Hsu, C.-C., & Wu, C.-Y. (2011). Modified adaptive resonance theory for alarm correlation based on distance hierarchy in mobile networks. *Network operations and management symposium, 2011 13th Asia-Pacific* (pp. 1–4).
- Cherng, S., Fang, C.-Y., Chen, C.-P., & Chen, S.-W. (2009). Critical motion detection of nearby moving vehicles in a vision-based driver-assistance system. *IEEE Transactions on Intelligent Transportation Systems*, 10, 70–82.
- Chiba, A. A., Bucci, D. J., Holland, P. C., & Gallagher, M. (1995). Basal forebrain cholinergic lesions disrupt increments but not decrements in conditioned stimulus processing. *The Journal of Neuroscience*, 15, 7315–7322.
- Chiu, Y. C., & Yantis, S. (2009). A domain-independent source of cognitive control for task sets: shifting spatial attention and switching categorization rules. *Journal of Neuroscience*, 29, 3930–3938.
- Chun, M. M. (2000). Contextual cueing of visual attention. *Trends in Cognitive Sciences*, 4, 170–178.
- Chun, M. M., & Jiang, Y. (1998). Contextual cueing: implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, 36, 28–71.
- Church, B. A., Krauss, M. S., Lopata, C., Toomey, J. A., Thomeer, M. L., Coutinho, M. V., Volker, M. A., & Mercado, E. (2010). Atypical categorization in children with high-functioning autism spectrum disorder. *Psychonomic Bulletin & Review*, 17, 862–868.
- Church, R. M., & Meck, W. H. (1984). The numerical attribute of stimuli. In H. L. Roitblat, T. G. Bever, & H. J. Terrace (Eds.), *Animal cognition*. Hillsdale, NJ: Erlbaum.
- Cohen, M. A., Grossberg, S., & Stork, D. G. (1988). Speech perception and production by a self-organizing neural network. In Y. C. Lee (Ed.), *Evolution, learning, cognition, and advanced architectures* (pp. 217–231). Hong Kong: World Scientific.
- Cohen, M. A., Grossberg, S., & Wyse, L. L. (1995). A spectral network model of pitch perception. *Journal of the Acoustical Society of America*, 98, 862–879.
- Cohen, N. J., & Squire, L. R. (1980). Preserved learning and retention of a pattern-analyzing skill in amnesia: dissociation of knowing how and knowing that. *Science*, 210, 207–210.
- Colburn, H. S. (1973). Theory of binaural interaction based on auditory-nerve data, I, general strategy and preliminary results on interaural discrimination. *Journal of the Acoustical Society of America*, 54, 1458–1470.
- Colburn, H. S. (1977). Theory of binaural interaction based on auditory-nerve data, II, detection of tones in noise. *Journal of the Acoustical Society of America*, 61, 525–533.
- Comon, P. (1994). Independent component analysis: a new concept? *Signal Processing*, 36, 287–314.
- Conway, C. M., & Christiansen, M. H. (2001). Sequential learning in non-human primates. *Trends in Cognitive Science*, 5(12), 539–546.
- Cowan, N. (2001). The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24, 87–185.
- Damasio, A. R. (1999). *The feeling of what happens: body and emotion in the making of consciousness*. Boston, MA: Houghton Mifflin Harcourt.
- Dehaene, S. (1997). *The number sense: how the mind creates mathematics*. New York: Oxford University Press.
- Demetgul, M., Tansel, I. N., & Taskin, S. (2009). Fault diagnosis of pneumatic systems with artificial neural network architectures. *Expert Systems with Applications*, 36, 10512–10519.
- Desimone, R. (1998). Visual attention mediated by biased competition in extrastriate visual cortex. *Philosophical Transactions of the Royal Society of London*, 353, 1245–1255.
- Destexhe, A., Contreras, D., & Steriade, M. (1999). Cortically-induced coherence of a thalamic-generated oscillation. *Neuroscience*, 92, 427–443.
- Douglas, R. J., Koch, C., Mahowald, M., Martin, K. A. C., & Suarez, H. H. (1995). Recurrent excitation in neocortical circuits. *Science*, 269, 981–985.
- Downing, C. J. (1988). Expectancy and visual-spatial attention: effects on perceptual quality. *Journal of Experimental Psychology: Human Perception and Performance*, 14, 188–202.
- Doya, K., Ishii, S., Pouget, A., & Rao, R. P. N. (Eds.) (2007). *Bayesian brain: probabilistic approaches to neural coding*. Cambridge, MA: MIT Press.
- Dranias, M., Grossberg, S., & Bullock, D. (2008). Dopaminergic and non-dopaminergic value systems in conditioning and outcome-specific reevaluation. *Brain Research*, 1238, 239–287.
- Driver, J., & Mattingley, J. B. (1998). Parietal neglect and visual awareness. *Nature Neuroscience*, 1, 17–22.
- Dubin, M. W., & Cleland, B. G. (1977). Organization of visual inputs to interneurons of lateral geniculate nucleus of the cat. *Journal of Neurophysiology*, 40, 410–427.
- Duhamel, J. R., Colby, C. L., & Goldberg, M. E. (1992). The updating of the representation of visual space in parietal cortex by intended eye movements. *Science*, 255, 90–92.
- Dunbar, G. (2012). Adaptive resonance theory as a model of polysemy and vagueness in the cognitive lexicon. *Cognitive Linguistics*, 23, 507–537.
- Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, 113, 501–517.
- Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., & Reitbock, H. J. (1988). Coherent oscillations: a mechanism of feature linking in the visual cortex? *Biological Cybernetics*, 60, 121–130.
- Egley, R., Driver, J., & Rafal, R. D. (1994). Shifting visual attention between objects and locations: evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology General*, 123, 161–177.
- Eichenbaum, H., & Lipton, P. A. (2008). Towards a functional organization of the medial temporal lobe memory system: role of the parahippocampal and medial entorhinal cortical areas. *Hippocampus*, 18, 1314–1324.
- Emmert, J., Lohmann, A., & Niemann, J. (1997). Pigeons' serial ordering of numerosity with visual arrays. *Animal Learning and Behavior*, 25, 234–244.
- Engel, A. K., Fries, P., & Singer, W. (2001). Dynamics predictions: oscillations and synchrony in top-down processing. *Nature Reviews Neuroscience*, 2, 704–716.

- Fang, L., & Grossberg, S. (2009). From stereogram to surface: how the brain sees the world in depth. *Spatial Vision*, 22, 45–82.
- Farrell, S., & Lewandowsky, S. (2002). An endogenous distributed model of ordering in serial recall. *Psychonomic Bulletin & Review*, 9, 59–79.
- Fazl, A., Grossberg, S., & Mingolla, E. (2009). View-invariant object category learning, recognition, and search: how spatial and object attention are coordinated using surface-based attentional shrouds. *Cognitive Psychology*, 58, 1–48.
- Felleman, D. J., & Van Essen, D. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1, 1–47.
- Fiala, J. C., Grossberg, S., & Bullock, D. (1996). Metabotropic glutamate receptor activation in cerebellar Purkinje cells as substrate for adaptive timing of the classically conditioned eye blink response. *Journal of Neuroscience*, 16, 3760–3774.
- Finch, E. A., & Augustine, G. J. (1998). Local calcium signalling by inositol-1, 4, 5-triphosphate in Purkinje cell dendrites. *Nature*, 396, 753–756.
- Foley, N. C., Grossberg, S., & Mingolla, E. (2012). Neural dynamics of object-based multifocal visual spatial attention and priming: object cueing, useful-field-of-view, and crowding. *Cognitive Psychology*, 65, 77–117.
- Francis, G., & Grossberg, S. (1996). Cortical dynamics of boundary segmentation and reset: persistence, afterimages, and residual traces. *Perception*, 35, 543–567.
- Francis, G., Grossberg, S., & Mingolla, E. (1994). Cortical dynamics of feature binding and reset: control of visual persistence. *Vision Research*, 34, 1089–1104.
- Frank, L. M., Stanley, G. B., & Brown, E. N. (2004). Hippocampal plasticity across multiple days of exposure to novel environments. *Journal of Neuroscience*, 24, 7681–7689.
- Funahashi, S., Inoue, M., & Kubota, K. (1997). Delay-period activity in the primate prefrontal cortex encoding multiple spatial positions and their order of presentation. *Behavioral Brain Research*, 84, 203–223.
- Gaffan, D. (1974). Recognition impaired and association intact in the memory of monkeys after transection of the fornix. *Journal of Comparative and Physiological Psychology*, 86, 1100–1109.
- Gao, E., & Suga, N. (1998). Experience-dependent corticofugal adjustment of midbrain frequency map in bat auditory system. *Proceedings of the National Academy of Sciences*, 95, 12663–12670.
- Gaudiano, P., & Grossberg, S. (1991). Vector associative maps: unsupervised real-time error-based learning and control of movement trajectories. *Neural Networks*, 4, 147–183.
- Gaudiano, P., & Grossberg, S. (1992). Adaptive vector integration to endpoint: self-organizing neural circuits for control of planned movement trajectories. *Human Movement Science*, 11, 141–155.
- Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., & Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience*, 2, 1527–1537.
- Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, 233, 1416–1419.
- Goldinger, S. D., & Azuma, T. (2003). Puzzle-solving science: the quixotic quest for units in speech perception. *Journal of Phonetics*, 31, 305–320.
- Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In F. Plum, & V. Mountcastle (Eds.), *Handbook of physiology*, vol. 5 (pp. 373–417). Bethesda: American Physiological Society.
- Gorchetchnikov, A., & Grossberg, S. (2007). Space, time and learning in the hippocampus: how fine spatial and temporal scales are expanded into population codes for behavioral control. *Neural Networks*, 20, 182–193.
- Gori, M., & Tesi, A. (1992). On the problem of local minima in backpropagation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14, 76–86.
- Gottlieb, J., Kusunoki, M., & Goldberg, M. E. (2005). Simultaneous representation of saccade targets and visual onsets in monkey lateral intraparietal area. *Cerebral Cortex*, 15, 1198–1206.
- Gove, A., Grossberg, S., & Mingolla, E. (1995). Brightness perception, illusory contours, and corticogeniculate feedback. *Visual Neuroscience*, 12, 1027–1052.
- Gray, C. M., & Singer, W. (1989). Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proceedings of the National Academy of Sciences USA*, 86, 1698–1702.
- Green, C. S., & Bavelier, D. (2003). Action video game modifies visual selective attention. *Nature*, 423, 534–537.
- Green, C. S., & Bavelier, D. (2007). Action-video-game experience alters the spatial resolution of vision. *Psychological Science*, 18, 88–94.
- Grossberg, S. (1968a). A prediction theory for some nonlinear functional-differential equations, II: learning of patterns. *Journal of Mathematical Analysis and Applications*, 22, 490–522.
- Grossberg, S. (1968b). Some nonlinear networks capable of learning a spatial pattern of arbitrary complexity. *Proceedings of the National Academy of Sciences*, 59, 368–372.
- Grossberg, S. (1968c). Some physiological and biochemical consequences of psychological postulates. *Proceedings of the National Academy of Sciences*, 60, 758–765.
- Grossberg, S. (1969a). On learning and energy-entropy dependence in recurrent and nonrecurrent signed networks. *Journal of Statistical Physics*, 1, 319–350.
- Grossberg, S. (1969b). On the production and release of chemical transmitters and related topics in cellular control. *Journal of Theoretical Biology*, 22, 325–364.
- Grossberg, S. (1969c). On the serial learning of lists. *Mathematical Biosciences*, 4, 201–253.
- Grossberg, S. (1971). On the dynamics of operant conditioning. *Journal of Theoretical Biology*, 33, 225–255.
- Grossberg, S. (1972a). A neural theory of punishment and avoidance, I: qualitative theory. *Mathematical Biosciences*, 15, 39–67.
- Grossberg, S. (1972b). A neural theory of punishment and avoidance, II: quantitative theory. *Mathematical Biosciences*, 15, 253–285.
- Grossberg, S. (1973). Contour enhancement, short-term memory, and constancies in reverberating neural networks. *Studies in Applied Mathematics*, 52, 213–257.
- Grossberg, S. (1974). Classical and instrumental learning by neural networks. In R. Rosen, & F. Snell (Eds.), *Progress in theoretical biology* (pp. 51–141). New York: Academic Press.
- Grossberg, S. (1975). A neural model of attention, reinforcement, and discrimination learning. *International Review of Neurobiology*, 18, 263–327.
- Grossberg, S. (1976a). Adaptive pattern classification and universal recoding, I: parallel development and coding of neural feature detectors. *Biological Cybernetics*, 23, 121–134.
- Grossberg, S. (1976b). Adaptive pattern classification and universal recoding, II: feedback, expectation, olfaction, and illusions. *Biological Cybernetics*, 23, 187–202.
- Grossberg, S. (1978a). A theory of human memory: self-organization and performance of sensory-motor codes, maps, and plans. In R. Rosen, & F. Snell (Eds.), *Progress in theoretical biology*, vol. 5 (pp. 233–374). New York: Academic Press.
- Grossberg, S. (1978b). Behavioral contrast in short-term memory: serial binary memory models or parallel continuous memory models? *Journal of Mathematical Psychology*, 3, 199–219.
- Grossberg, S. (1980). How does a brain build a cognitive code? *Psychological Review*, 87, 1–51.
- Grossberg, S. (1982). Processing of expected and unexpected events during conditioning and attention: a psychophysiological theory. *Psychological Review*, 89, 529–572.
- Grossberg, S. (1984a). Some normal and abnormal behavioral syndromes due to transmitter gating of opponent processes. *Biological Psychiatry*, 19, 1075–1118.
- Grossberg, S. (1984b). Some psychophysiological and pharmacological correlates of a developmental, cognitive, and motivational theory. In R. Karrer, J. Cohen, & P. Tueting (Eds.), *Brain and information: event related potentials* (pp. 58–142). New York: New York Academy of Sciences.
- Grossberg, S. (1986). The adaptive self-organization of serial order in behavior: speech, language, and motor control. In E. C. Schwab, & H. C. Nusbaum (Eds.), *Pattern recognition by humans and machines*, vol. 1: speech perception (pp. 187–294). New York: Academic Press.
- Grossberg, S. (1988). Nonlinear neural networks: principles, mechanisms, and architectures. *Neural Networks*, 1, 17–61.
- Grossberg, S. (1994). 3-D vision and figure-ground separation by visual cortex. *Perception and Psychophysics*, 55, 48–120.
- Grossberg, S. (1997). Cortical dynamics of three-dimensional figure-ground perception of two-dimensional figures. *Psychological Review*, 104, 618–658.
- Grossberg, S. (1999). How does the cerebral cortex work? learning, attention and grouping by the laminar circuits of visual cortex. *Spatial Vision*, 12, 163–186.
- Grossberg, S. (2000a). How hallucinations may arise from brain mechanisms of learning, attention, and volition. *Journal of the International Neuropsychological Society*, 6, 579–588. Invited article for the.
- Grossberg, S. (2000b). The complementary brain: unifying brain dynamics and modularity. *Trends in Cognitive Sciences*, 4, 233–246.
- Grossberg, S. (2000c). The imbalanced brain: from normal behavior to schizophrenia. *Biological Psychiatry*, 48, 81–98.
- Grossberg, S. (2003a). How does the cerebral cortex work? development, learning, attention, and 3D vision by laminar circuits of visual cortex. *Behavioral and Cognitive Neuroscience Reviews*, 2, 47–76.
- Grossberg, S. (2003b). Resonant neural dynamics of speech perception. *Journal of Phonetics*, 31, 423–445.
- Grossberg, S. (2007a). Consciousness CLEARs the mind. *Neural Networks*, 20, 1040–1053.
- Grossberg, S. (2007b). Towards a unified theory of neocortex: laminar cortical circuits for vision and cognition. *Progress in Brain Research*, 165, 79–104.
- Grossberg, S. (2009a). Beta oscillations and hippocampal place cell learning during exploration of novel environments. *Hippocampus*, 19, 881–885.
- Grossberg, S. (2009b). Cortical and subcortical predictive dynamics and learning during perception, cognition, emotion and action. *Philosophical Transactions of the Royal Society of London B Biological Sciences*, 364, 1223–1234.
- Grossberg, S., Bullock, D., & Dranias, M. (2008). Neural dynamics underlying impaired autonomic and conditioned responses following amygdala and orbitofrontal lesions. *Behavioral Neuroscience*, 122, 1100–1125.
- Grossberg, S., Govindarajan, K. K., Wyse, L. L., & Cohen, M. A. (2004). ARTSTREAM: a neural network model of auditory scene analysis and source segregation. *Neural Networks*, 17, 511–536.
- Grossberg, S., & Grunewald, A. (1997). Cortical synchronization and perceptual framing. *Journal of Cognitive Neuroscience*, 9, 117–132.
- Grossberg, S., & Huang, T.-R. (2009). ARTSCENE: a neural system for natural scene classification. *Journal of Vision*, 9(6), 1–19. <http://dx.doi.org/10.1167/9.4.6>. <http://journalofvision.org/9/4/6/>.
- Grossberg, S., & Kazerounian, S. (2011). Laminar cortical dynamics of conscious speech perception: a neural model of phonemic restoration using subsequent context in noise. *Journal of the Acoustical Society of America*, 130, 440–460.

- Grossberg, S., & Levine, D. S. (1987). Neural dynamics of attentionally modulated Pavlovian conditioning: blocking, inter-stimulus interval, and secondary reinforcement. *Applied Optics*, 26, 5015–5030.
- Grossberg, S., Markowitz, J., & Cao, Y. (2011). On the road to invariant recognition: explaining tradeoff and morph properties of cells in inferotemporal cortex using multiple-scale task-sensitive attentive learning. *Neural Networks*, 24, 1036–1049.
- Grossberg, S., & McLoughlin, N. (1997). Cortical dynamics of 3-D surface perception: binocular and half-occluded scenic images. *Neural Networks*, 10, 1583–1605.
- Grossberg, S., & Merrill, J. W. L. (1992). A neural network model of adaptively timed reinforcement learning and hippocampal dynamics. *Cognitive Brain Research*, 1, 3–38.
- Grossberg, S., & Merrill, J. W. L. (1996). The hippocampus and cerebellum in adaptively timed learning, recognition, and movement. *Journal of Cognitive Neuroscience*, 8, 257–277.
- Grossberg, S., & Mingolla, E. (1985). Neural dynamics of form perception: boundary completion, illusory figures, and neon color spreading. *Psychological Review*, 92, 173–211.
- Grossberg, S., Mingolla, E., & Ross, W. D. (1997). Visual brain and visual perception: how does the cortex do perceptual grouping? *Trends in Neurosciences*, 20, 106–111.
- Grossberg, S., Mingolla, E., & Viswanathan, L. (2001). Neural dynamics of motion integration and segmentation within and across apertures. *Vision Research*, 41, 2521–2553.
- Grossberg, S., & Myers, C. W. (2000). The resonant dynamics of speech perception: interword integration and duration-dependent backward effects. *Psychological Review*, 107, 735–767.
- Grossberg, S., & Paine, R. W. (2000). A neural model of corticocerebellar interactions during attentive imitation and predictive learning of sequential handwriting movements. *Neural Networks*, 13, 999–1046.
- Grossberg, S., & Pearson, L. (2008). Laminar cortical dynamics of cognitive and motor working memory, sequence learning and performance: toward a unified theory of how the cerebral cortex works. *Psychological Review*, 115, 677–732.
- Grossberg, S., & Pilly, P. (2012). How entorhinal grid cells may learn multiple spatial scales from a dorsoventral gradient of cell response rates in a self-organizing map. *PLoS Computational Biology* (in press).
- Grossberg, S., & Raizada, R. (2000). Contrast-sensitive perceptual grouping and object-based attention in the laminar circuits of primary visual cortex. *Vision Research*, 40, 1413–1432.
- Grossberg, S., & Repin, D. (2003). A neural model of how the brain represents and compares multi-digit numbers: spatial and categorical processes. *Neural Networks*, 16, 1107–1140.
- Grossberg, S., Roberts, K., Aguilar, M., & Bullock, D. (1997). A neural model of multimodal adaptive saccadic eye movement control by superior colliculus. *Journal of Neuroscience*, 17, 9706–9725.
- Grossberg, S., & Schmajuk, N. A. (1989). Neural dynamics of adaptive timing and temporal discrimination during associative learning. *Neural Networks*, 2, 79–102.
- Grossberg, S., & Seidman, D. (2006). Neural dynamics of autistic behaviors: cognitive, emotional, and timing substrates. *Psychological Review*, 113, 483–525.
- Grossberg, S., & Somers, D. (1991). Synchronized oscillations during cooperative feature linking in a cortical model of visual perception. *Neural Networks*, 4, 453–466.
- Grossberg, S., & Swaminathan, G. (2004). A laminar cortical model for 3D perception of slanted and curved surfaces and of 2D images: development, attention and bistability. *Vision Research*, 44, 1147–1187.
- Grossberg, S., & Todorovic, D. (1988). Neural dynamics of 1-D and 2-D brightness perception: a unified model of classical and recent phenomena. *Perception & Psychophysics*, 43, 241–277.
- Grossberg, S., & Versace, M. (2008). Spikes, synchrony, and attentive learning by laminar thalamocortical circuits. *Brain Research*, 1218, 278–312.
- Grossberg, S., & Vladusich, T. (2011). How do children learn to follow gaze, share joint attention, imitate their teachers, and use tools during social interactions? *Neural Networks*, 23, 940–965.
- Grossberg, S., & Williamson, J. R. (2001). A neural model of how horizontal and interlaminar connections of visual cortex develop into adult circuits that carry out perceptual groupings and learning. *Cerebral Cortex*, 11, 37–58.
- Grossberg, S., & Yazdanbakhsh, A. (2005). Laminar cortical dynamics of 3D surface perception: stratification, transparency, and neon color spreading. *Vision Research*, 45, 1725–1743.
- Grossberg, S., Yazdanbakhsh, A., Cao, Y., & Swaminathan, G. (2008). How does binocular rivalry emerge from cortical mechanisms of 3-D vision? *Vision Research*, 48, 2232–2250.
- Guenther, F. H. (1995). Speech sound acquisition, coarticulation, and rate effects in a neural network model of speech production. *Psychological Review*, 102, 594–621.
- Guenther, F. H., Bullock, D., Greve, D., & Grossberg, S. (1994). Neural representations for sensory-motor control, III: learning a body-centered representation of 3-D target position. *Journal of Cognitive Neuroscience*, 6, 341–358.
- Guenther, F. H., Ghosh, S. S., & Tourville, J. A. (2006). Neural modeling and imaging of the cortical interactions underlying syllable production. *Brain and Language*, 96, 280–301.
- Hafting, T., Fyhn, M., Molden, S., Moser, M. B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436, 801–806.
- Hata, T., Kumai, K., & Okaichi, H. (2007). Hippocampal acetylcholine efflux increases during negative patterning and elemental discrimination in rats. *Neuroscience Letters*, 418, 127–132.
- He, H., Caudell, T. P., Menicucci, D. F., & Mammoli, A. A. (2012). Application of adaptive resonance theory neural networks to monitor solar hot water systems and detect existing or developing faults. *Solar Energy*, 86, 2318–2333.
- He, S., Cavanagh, P., & Intriligator, J. (1996). Attentional resolution and the locus of visual awareness. *Nature*, 383, 334–337.
- He, J., Tan, A.-H., & Tan, C.-L. (2000). A comparative study on Chinese text categorization methods. In *Proceedings of PRICA'2000*.
- Healy, M. J., Caudell, T. P., & Smith, S. D. G. (1993). A neural architecture for pattern sequence verification through inferencing. *IEEE Transactions on Neural Networks*, 4, 9–20.
- Hecht-Nielsen, R. (1987). Counterpropagation networks. *Applied Optics*, 26, 4979–4983.
- Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, 9, 181–197.
- Henson, R. N. A. (1998). Short-term memory for serial order: the start-end model of serial recall. *Cognitive Psychology*, 36, 73–137.
- Histed, M. H., & Miller, E. K. (2006). Microstimulation of frontal cortex can reorder a remembered spatial sequence. *Public Library of Science: Biology*, 4(5), e134.
- Ho, C. S., Liou, J. J., Georgiopoulos, M., Heileman, G. L., & Christodoulou, C. (1994). Analogue circuit design and implementation of an adaptive resonance theory (ART) network architecture. *International Journal of Electronics*, 76, 271–291.
- Hollerman, J., & Schultz, W. (1998). Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience*, 1, 304–309.
- Houghton, G. (1990). The problem of serial order: a neural network model of sequence learning and recall. In R. Dale, C. Mellish, & M. Zock (Eds.), *Current research in natural language generation* (pp. 287–319). London: Academic Press.
- Hsieh, K.-L. (2008). The application of clustering analysis for the critical areas on TFT-LCD panel. *Expert Systems with Applications*, 34, 952–957.
- Hsieh, K.-L., & Yang, I.-Ch. (2008). Incorporating PCA and fuzzy-ART techniques into achieve organism classification based on codon usage consideration. *Computers in Biology and Medicine*, 38, 886–893.
- Hsu, S.-C., & Chien, C.-F. (2007). Hybrid data mining approach for pattern extraction from wafer bin map to improve yield in semiconductor manufacturing. *International Journal of Production Economics*, 107, 88–103.
- Huang, T.-R., & Grossberg, S. (2010). Cortical dynamics of contextually cued attentive visual learning and search: spatial and object evidence accumulation. *Psychological Review*, 117, 1080–1112.
- Hunt, R. R., & Lamb, C. A. (2001). What causes the isolation effect? *Journal of Experimental Psychology: Learning, Memory and Cognition*, 27(6), 1359–1366.
- Hupé, J. M., James, A. C., Girard, D. C., & Bullier, J. (1997). Feedback connections from V2 modulate intrinsic connectivity within V1. *Society for Neuroscience Abstracts*, 406(15), 1031.
- Husain, M., Mannan, S., Hodgson, T., Wojcicuk, E., Driver, J., & Kennard, C. (2001). Impaired spatial working memory across saccades contributes to abnormal search in parietal neglect. *Brain*, 124, 941–952.
- Hyvärinen, A., & Oja, E. (2000). Independent component analysis: algorithms and application. *Neural Networks*, 13, 411–430.
- Ichise, T., Kano, M., Hashimoto, K., Yangihara, D., Nakao, K., Shigemoto, R., Katsuki, M., & Aiba, A. (2000). mGluR1 in cerebellar Purkinje cells essential for long-term depression, synapse elimination, and motor coordination. *Science*, 288, 1832–1835.
- Inoue, M., & Mikami, A. (2006). Prefrontal activity during serial probe reproduction task: encoding, mnemonic and retrieval processes. *Journal of Neurophysiology*, 95, 1008–1041.
- Intriligator, J., & Cavanagh, P. (2001). The spatial resolution of visual attention. *Cognitive Psychology*, 43, 171–216.
- Jones, D., Farrand, P., Stuart, G., & Morris, N. (1995). The functional equivalence of verbal and spatial memory in serial short-term memory. *Journal of Experimental Psychology: Learning Memory and Cognition*, 21, 1008–1018.
- Kaas, J. H. (1999). Is most of neural plasticity in the thalamus cortical? *Proceedings of the National Academy of Sciences USA*, 96, 7622–7623.
- Kamin, L. J. (1969). Predictability, surprise, attention and conditioning. In B. A. Campbell, & R. M. Church (Eds.), *Punishment and aversive behavior*. New York: Appleton-Century-Crofts.
- Kastner, S., & Ungerleider, L. G. (2001). The neural basis of biased competition in human visual cortex. *Neuropsychologia*, 39, 1263–1276.
- Kawamura, T., Takahashi, H., & Honda, H. (2008). Proposal of new gene filtering method, BagPART, for gene expression analysis with small sample. *Journal of Bioscience and Bioengineering*, 105, 81–84.
- Kaylani, A., Georgiopoulos, M., Mollaghasemi, M., & Anagnostopoulos, G. C. (2009). AG-ART: an adaptive approach to evolving ART architectures. *Neurocomputing*, 72, 2079–2092.
- Kazerounian, S., & Grossberg, S. (2012). Real-time learning of predictive recognition categories that chunk sequences of items stored in working memory (submitted for publication).
- Kemel, M. L., Desban, M., Gauchy, C., Glowinski, J., & Besson, M. J. (1988). Topographical organization of efferent projections from the cat substantia nigra pars reticulata. *Brain Research*, 455, 307–323.
- Kentros, C. G., Agnietri, N. T., Streater, S., Hawkins, R. D., & Kandel, E. R. (2004). Increased attention to spatial context increases both place field stability and spatial memory. *Neuron*, 42, 283–295.

- Kentros, C., Hargreaves, E., Hawkins, R. D., Kandel, E. R., Shapiro, M., & Muller, R. V. (1998). Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. *Science*, 280, 2121–2126.
- Keskin, G. A., & Ozkan, C. (2009). An alternative evaluation of FMEA: fuzzy ART algorithm. *Quality and Reliability Engineering International*, 25, 647–661.
- Kilgard, M. P., & Merzenich, M. M. (1998). Cortical map reorganization enabled by nucleus basalis activity. *Science*, 279, 1714–1718.
- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 12, 712–719.
- Knoedler, A. J., Hellwig, K. A., & Neath, I. (1999). The shift from recency to primacy with increasing delay. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 25, 474–487.
- Knowlton, B. J., & Squire, L. R. (1993). The learning of categories: parallel brain systems for item memory and category knowledge. *Science*, 262, 1747–1749.
- Kohonen, T. (1984). *Self-organization and associative memory*. New York: Springer-Verlag.
- Kraus, N., McGee, T., Littman, T., Nicol, T., & King, C. (1994). Nonprimary auditory thalamic representation of acoustic change. *Journal of Neurophysiology*, 72, 1270–1277.
- Krupa, D. J., Ghazanfar, A. A., & Nicolelis, M. A. (1999). Immediate thalamic sensory plasticity depends on corticothalamic feedback. *Proceedings of the National Academy of Sciences USA*, 96, 8200–8205.
- Langston, R. F., Ainge, J. A., Couey, J. J., Canto, C. B., Bjerknes, T. L., Witter, M. P., Moser, E. I., & Moser, M. B. (2010). Development of the spatial representation system in the rat. *Science*, 328, 1576–1580.
- LeDoux, J. E. (1993). Emotional memory systems in the brain. *Behavioral Brain Research*, 58, 69–79.
- Levi, D. M. (2008). Crowding—an essential bottleneck for object recognition: a mini-review. *Vision Research*, 48, 635–654.
- Levy, W. B., & Steward, O. (1983). Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience*, 8, 791–797.
- Li, N., & DiCarlo, J. J. (2008). Unsupervised natural experience rapidly alters invariant object representation in visual cortex. *Science*, 321, 1502–1507.
- Li, N., & DiCarlo, J. J. (2010). Unsupervised natural visual experience rapidly reshapes size invariant object represent in inferior temporal cortex. *Neuron*, 67, 1062–1075.
- Liu, L., Huang, L., Lai, M., & Ma, C. (2009). Projective ART with buffers for the high dimensional space clustering and an application to discover stock associations. *Neurocomputing*, 72, 1283–1295.
- Liu, D., Pang, Z., & Lloyd, S. R. (2008). A neural network method for detection of obstructive sleep apnea and narcolepsy based on pupil size and EEG. *IEEE Transactions on Neural Networks*, 19, 308–318.
- Ljungberg, T., Apicella, P., & Schultz, W. (1992). Responses of monkey dopamine neurons during learning of behavioral reactions. *Journal of Neurophysiology*, 67, 145–163.
- Lopes, M. L. M., Minussi, C. R., & Lotufo, A. D. P. (2005). Electric load forecasting using a fuzzy ART & ARTMAP neural network. *Applied Soft Computing*, 5, 235–244.
- Luce, P. A., & McLennan, C. T. (2008). Spoken word recognition: the challenge of variation. In D. B. Pisoni, & R. E. Remez (Eds.), *The handbook of speech perception*. Wiley Online Library. <http://dx.doi.org/10.1002/9780470757024.ch24>.
- Luck, S. J., Chelazzi, L., Hillyard, S. A., & Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology*, 77, 24–42.
- Lumer, E. D., Edelman, G. M., & Tononi, G. (1997). Neural dynamics in a model of the thalamocortical system I, layers, loops and the emergence of fast synchronous rhythms. *Cerebral Cortex*, 7, 207–227.
- Lund, J. S., & Wu, C. Q. (1997). Local circuit neurons of macaque monkey striate cortex: IV, neurons of laminae 1–3A. *Journal of Comparative Neurology*, 384, 109–126.
- Marchiori, S. C., da Silveira, M. do C., Lotufo, A. D. P., Minussi, C. R., & Lopes, M. L. M. (2011). Neural network based on adaptive resonance theory with continuous training for multi-configuration transient stability analysis of electric power systems. *Applied Soft Computing*, 11, 706–715.
- Markram, H., Lubke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 275, 213–215.
- Martin-Guerrero, J. D., Lisboa, P. J. G., Soria-Olivas, E., Palomares, A., & Balaguer, E. (2007). An approach based on the adaptive resonance theory for analyzing the viability of recommender systems in a citizen web portal. *Expert Systems with Applications*, 33, 743–753.
- Massey, L. (2009). Discovery of hierarchical thematic structure in text collections with adaptive resonance theory. *Neural Computation & Applications*, 18, 261–273.
- Maunsell, J. H. R., & Van Essen, D. C. (1983). Anatomical connections of the middle temporal visual area in the macaque monkey and their relationship to a hierarchy of cortical areas. *Journal of Neuroscience*, 3, 2563–2586.
- McLennan, C. T., Conor, T., & Luce, P. A. (2005). Examining the time course of indexical specificity effects in spoken word recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 306–321.
- McLennan, C. T., Luce, P. A., & Charles-Luce, J. (2003). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 539–553.
- Melcher, D. (2007). Predictive remapping of visual features precedes saccadic eye movements. *Nature Neuroscience*, 10, 903–907.
- Mesulam, M.-M. (1999). Spatial attention and neglect: parietal, frontal and cingulate contributions to the mental representation and attentional targeting of salient extrapersonal events. *Philosophical Transactions of the Royal Society B*, 354, 1325–1346.
- Mhatre, H., Gorchetchnikov, A., & Grossberg, S. (2012). Grid cell hexagonal patterns formed by fast self-organized learning within entorhinal cortex. *Hippocampus*, 22, 320–334.
- Miller, G. A. (1956). The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, 63(2), 81–97.
- Miller, G. A., & Licklider, J. C. R. (1950). The intelligibility of interrupted speech. *Journal of the Acoustical Society of America*, 22, 167–173.
- Mirenovic, J., & Schultz, W. (1994). Importance of unpredictability for reward responses in primate dopamine neurons. *Journal of Neurophysiology*, 72, 1024–1027.
- Mishkin, M., & Delacour, J. (1975). An analysis of short-term visual memory in the monkey. *Journal of Experimental Psychology: Animal Behavior Processes*, 1, 326–334.
- Miyata, M., Finch, E. A., Khiroug, L., Hashimoto, K., Hayasaka, S., Oda, S. I., Inouye, M., Takagishi, Y., Augustine, G. J., & Kano, M. (2000). Local calcium release in dendritic spines required for long-term synaptic depression. *Neuron*, 28, 233–244.
- Morris, R. G. M., & Frey, U. (1997). Hippocampal synaptic plasticity: role in spatial learning or the automatic recording of attended experience? *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 350, 1469–1503.
- Mounts, J. R. W. (2000). Evidence for suppressive mechanisms in attentional selection: feature singletons produce inhibitory surrounds. *Perception and Psychophysics*, 62, 969–983.
- Mulder, S. A., & Wunsch, D. C. (2003). Million city traveling salesman problem solution by divide and conquer clustering with adaptive resonance neural networks. *Neural Networks*, 16, 827–832.
- Muller, R. A. (1996). A quarter of a century of place cells. *Neuron*, 17, 813–822.
- Mumford, D. (1992). On the computational architecture of the neocortex, II, the role of corticocortical loops. *Biological Cybernetics*, 66, 241–251.
- Murdock, B. B. (1962). The serial position effect of free recall. *Journal of Experimental Psychology*, 64, 482–488.
- Nieder, A., & Miller, E. K. (2003). Coding of cognitive magnitude: compressed scaling of numerical information in the primate prefrontal cortex. *Neuron*, 37, 149–157.
- Nieder, A., & Miller, E. K. (2004). A parieto-frontal network for visual numerical information in the monkey. *Proceedings of the National Academy of Sciences*, 101, 7457–7462.
- Owega, S., Khan, B.-U.-Z., Evans, G. J., Jervis, R. E., & Fila, M. (2006). Identification of long-range aerosol transport patterns to Toronto via classification of back trajectories by cluster analysis and neural network techniques. *Chemometrics and Intelligent Laboratory Systems*, 83, 26–33.
- Page, M. P. A., & Norris, D. (1998). The primacy model: a new model of immediate serial recall. *Psychological Review*, 105, 761–781.
- Palma, J., Grossberg, S., & Versace, M. (2012). Persistence and storage of activity patterns in spiking recurrent cortical networks: modulation of sigmoid signals by after-hyperpolarization currents and acetylcholine. *Frontiers in Computational Neuroscience* (in press).
- Pandya, D. N., & Yeterian, E. H. (1985). Architecture and connections of cortical association areas. In A. Peters, & E. G. Jones (Eds.), *Cerebral cortex 10*. New York: Plenum Press.
- Parker, J. L., & Dostrovsky, J. O. (1999). Cortical involvement in the induction, but not expression, of thalamic plasticity. *The Journal of Neuroscience*, 19, 8623–8629.
- Parsons, O., & Carpenter, G. A. (2003). ARTMAP neural networks for information fusion and data mining: map production and target recognition methodologies. *Neural Networks*, 16, 1075–1089.
- Pasupathy, A., & Miller, E. K. (2004). Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature*, 433, 873–876.
- Pavlov, I. P. (1927). *Conditioned reflexes*. Oxford University Press.
- Perret, S. P., Ruiz, B. P., & Mauk, M. D. (1993). Cerebellar cortex lesions disrupt learning-dependent timing of conditioned eyelid responses. *Journal of Neuroscience*, 13, 1708–1718.
- Perry, E. K., Lee, M. L. W., Martin-Ruiz, C. M., Court, J. A., Volsen, S. G., Merrit, J., Folly, E., Iversen, P. E., Bauman, M. L., Perry, R. H., & Wenk, G. L. (2001). Cholinergic activity in autism: abnormalities in the cerebral cortex and basal forebrain. *The American Journal of Psychiatry*, 158, 1058–1066.
- Peterson, G. E., & Barney, H. L. (1952). Control methods used in a study of the vowels. *Journal of the Acoustical Society of America*, 24, 175–184.
- Petrides, M. (2005). Lateral prefrontal cortex: architectonic and functional organization. *Philosophical Transactions of the Society of London B Biological Science*, 360, 781–795.
- Piaget, J. (1963). *The origins of intelligence in children*. New York: Norton.
- Pilly, P. K., & Grossberg, S. (2012). How do spatial learning and memory occur in the brain? Coordinated learning of entorhinal grid cells and hippocampal place cells. *Journal of Cognitive Neuroscience* (in press).
- Pilly, P., Grossberg, S., & Seitz, A. (2010). Low-level sensory plasticity during task-irrelevant perceptual learning: evidence from conventional and double training procedures. *Vision Research*, 50, 424–432. (Special Issue on Perceptual Learning).
- Pollen, D. A. (1999). On the neural correlates of visual perception. *Cerebral Cortex*, 9, 4–19.
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32, 3–25.

- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77, 353–363.
- Prasad, V. S. S., & Gupta, S. D. (2008). Photometric clustering of regenerated plants of gladiolus by neural networks and its biological validation. *Computers and Electronics in Agriculture*, 60, 8–17.
- Raizada, R., & Grossberg, S. (2001). Context-sensitive bindings by the laminar circuits of V1 and V2: a unified model of perceptual grouping, attention, and orientation contrast. *Visual Cognition*, 8, 431–466.
- Raizada, R., & Grossberg, S. (2003). Towards a theory of the laminar architecture of cerebral cortex: computational clues from the visual system. *Cerebral Cortex*, 13, 100–113.
- Ramachandran, V. S. (1990). Interactions between motion, depth, color and form: the utilitarian theory of perception. In C. Blakemore (Ed.), *Vision: coding and efficiency*. Cambridge, England: Cambridge University Press.
- Ranganath, C., & D'Esposito, M. (2005). Directing the mind's eye: prefrontal, inferior and medial temporal mechanisms for visual working memory. *Current Opinion in Neurobiology*, 15, 175–182.
- Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive field effects. *Nature Neuroscience*, 2, 79–87.
- Remez, R. E. (2003). Establishing and maintaining perceptual coherence: unimodal and multimodal evidence. *Journal of Phonetics*, 31, 293–304.
- Remez, R. E., Pardo, J. S., Piorkowski, R. L., & Rubin, P. E. (2001). On the bistability of sine wave analogues of speech. *Psychological Science*, 12, 24–29.
- Remez, R. E., Rubin, P. E., Berns, S. M., Pardo, J. S., & Lang, J. M. (1994). On the perceptual organization of speech. *Psychological Review*, 101, 129–156.
- Repp, B., Liberman, A., Eccardt, T., & Pesetsky, D. (1978). Perceptual integration of acoustic cues for stop, fricative, and affricate manner. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 621–637.
- Reynolds, J., Chelazzi, L., & Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *The Journal of Neuroscience*, 19, 1736–1753.
- Reynolds, J. H., & Desimone, R. (2003). Interacting roles of attention and visual salience in V4. *Neuron*, 37, 853–863.
- Reynolds, J. H., & Heeger, D. J. (2009). The normalization model of attention. *Neuron*, 61, 168–185.
- Rockland, K. S. (1994). The organization of feedback connections from area V1(18) to V1(17). In A. Peters, & K. S. Rockland (Eds.), *Cerebral cortex, vol. 4*. New York: Plenum Press.
- Rockland, K. S., Andresen, J., Cowie, R. J., & Robinson, D. L. (1999). Single axon analysis of pulvinocortical connections to several visual areas in the macaque. *Journal of Comparative Neurology*, 406, 221–250.
- Rockland, K. S., & Virga, A. (1989). Terminal arbors of individual 'feedback' axons projecting from area V2 to V1 in the macaque monkey: a study using immunohistochemistry of anterogradely transported phaseolus vulgaris-leucoagglutinin. *Journal of Comparative Neurology*, 285, 54–72.
- Roelfsema, P. R., Lamme, V. A. F., & Spekreijse, H. (1998). Object-based attention in the primary visual cortex of the macaque monkey. *Nature*, 395, 376–381.
- Roggeveen, A., Pilz, K., Bennett, P., & Sekuler, A. (2009). Individual differences in object based attention. *Journal of Vision*, 9, 143.
- Rouiller, E. M., & Welker, E. (2000). A comparative analysis of the morphology of corticothalamic projections in mammals. *Brain Research Bulletin*, 53, 727–741.
- Saar, D., Grossman, Y., & Barkai, E. (2001). Long-lasting cholinergic modulation underlies rule learning in rats. *Journal of Neuroscience*, 21, 1385–1392.
- Salin, P., & Bullier, J. (1995). Corticocortical connections in the visual system: structure and function. *Physiological Reviews*, 75, 107–154.
- Samuel, A. G. (1981). The role of bottom-up confirmation in the phonemic restoration illusion. *Journal of Experimental Psychology: Human Perception and Performance*, 7, 1124–1131.
- Sandell, J. H., & Schiller, P. H. (1982). Effect of cooling area 18 on striate cortex cells in the squirrel monkey. *Journal of Neurophysiology*, 48, 38–48.
- Schoenbaum, G., Setlow, B., Saddoris, M. P., & Gallagher, M. (2003). Encoding predicted outcome and acquired value in orbitofrontal cortex during cue sampling depends upon input from basolateral amygdala. *Neuron*, 39, 855–867.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80, 1–27.
- Schultz, W., Apicella, P., & Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. *Journal of Neuroscience*, 13, 900–913.
- Schultz, W., Romo, R., Ljungberg, T., Mirenowicz, J., Hollerman, J., & Dickinson, A. (1995). Reward-related signals carried by dopamine neurons. In J. Houk, J. Davis, & D. Beiser (Eds.), *Models of information processing in the Basal Ganglia* (pp. 11–27). Cambridge: MIT Press.
- Sears, L. L., Finn, P. R., & Steinmetz, J. E. (1994). Abnormal classical eye-blink conditioning in autism. *Journal of Autism and Developmental Disorders*, 24, 737–751.
- Seitz, A., & Watanabe, T. (2003). Is subliminal learning really passive? *Nature*, 422, 6927.
- Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *The Journal of Neuroscience*, 18, 3870–3896.
- Sherman, S. M., & Guillery, R. (2001). *Exploring the thalamus*. San Diego: Academic Press.
- Sherman, S. M., & Guillery, R. W. (2002). The role of the thalamus in the flow of information to the cortex. *Philosophical Transactions of the Royal Society London B*, 357, 1695–1708.
- Shieh, M.-D., Yan, W., & Chen, C.-H. (2008). Soliciting customer requirements for product redesign based on picture sorts and ART2 neural network. *Expert Systems with Applications*, 34, 194–204.
- Shipp, S. (2003). The functional logic of cortico-pulvinar connections. *Philosophical Transactions of the Royal Society London B*, 358, 1605–1624.
- Siegel, M., Kording, K. P., & König, P. (2000). Integrating top-down and bottom-up sensory processing by Somato-Dendritic interactions. *Journal of Computational Neuroscience*, 8, 161–173.
- Sigala, N., & Logothetis, N. K. (2002). Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature*, 415, 318–320.
- Sillito, A. M., Jones, H. E., Gerstein, G. L., & West, D. C. (1994). Feature-linked synchronization of thalamic relay cell firing induced by feedback from the visual cortex. *Nature*, 369, 479–482.
- Silver, M. R., Grossberg, S., Bullock, D., Histed, M. H., & Miller, E. K. (2011). A neural model of sequential movement planning and control of eye movements: item-order-rank working memory and saccade selection by the supplementary eye fields. *Neural Networks*, 26, 29–58.
- Silver, M. A., Shenhav, A., & D'Esposito, M. (2008). Cholinergic enhancement reduces spatial spread of visual responses in human early visual cortex. *Neuron*, 60, 904–914.
- Smith, M. C. (1968). CS-US interval and US intensity in classical conditioning of the rabbit's nictitating membrane response. *Journal of Comparative and Physiological Psychology*, 3, 679–687.
- Somers, D. C., Dale, A. M., Seiffert, A. E., & Tootell, R. B. (1999). Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. *Proceedings of the National Academy of Sciences USA*, 96, 1663–1668.
- Staddon, J. E. R. (1983). *Adaptive behavior and learning*. New York: Cambridge University Press.
- Starkey, P., Spelke, E. S., & Gelman, R. (1983). Detection of intermodal numerical correspondences by human infants. *Science*, 222, 179–181.
- Steinman, B. A., Steinman, S. B., & Lehmkuhle, S. (1995). Visual attention mechanisms show a center-surround organization. *Vision Research*, 35, 1859–1869.
- Sudhakar Pandian, R., & Mahapatra, S. S. (2009). Manufacturing cell formation with production data using neural networks. *Computers & Industrial Engineering*, 56, 1340–1347.
- Takahashi, H., Murase, Y., Kobayashi, T., & Honda, H. (2007). New cancer diagnosis modeling using boosting and projective adaptive resonance theory with improved reliable index. *Biochemical Engineering Journal*, 33, 100–109.
- Takechi, H., Eilers, J., & Konnerth, A. (1998). A new class of synaptic response involving calcium release in dendritic spines. *Nature*, 396, 757–760.
- Tan, A.-H. (1997). Cascade ARTMAP: integrating neural computation and symbolic knowledge processing. *IEEE Transactions on Neural Networks*, 8, 237–250.
- Tan, T. Z., Quek, C., Ng, G. S., & Razvi, K. (2008). Ovarian cancer diagnosis with complementary learning fuzzy neural network. *Artificial Intelligence in Medicine*, 43, 207–222.
- Tan, A.-H., & Teo, C. (1998). Learning user profiles for personalized information dissemination. *IEEE World Congress on Computational Intelligence*, 1, 183–188.
- Tanaka, K., Saito, H., Fukada, Y., & Moriya, M. (1991). Coding visual images of objects in the inferotemporal cortex of the macaque monkey. *Journal of Neurophysiology*, 66, 170–189.
- Temereanca, S., & Simons, D. J. (2001). Topographic specificity in the functional effects of corticofugal feedback in the whisker/barrel system. *Society for Neuroscience Abstracts*, 393, 6.
- Theeuwes, J., Mathot, S., & Kingstone, A. (2010). Object-based eye movements: the eyes prefer to stay within the same object. *Attention, Perception & Psychophysics*, 72, 597–601.
- Thompson, L. T., & Best, P. J. (1990). Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. *Brain Research*, 509, 299–308.
- Thompson, R. F., Clark, G. A., Donegan, N. H., Lavond, G. A., Lincoln, D. G., Maddon, J., Mamounas, L. A., Mauk, M. D., & McCormick, D. A. (1987). Neuronal substrates of discrete, defensive conditioned reflexes, conditioned fear states, and their interactions in the rabbit. In I. Gormezano, W. F. Prokasy, & R. F. Thompson (Eds.), *Classical conditioning* (3rd ed.) (pp. 371–399). Hillsdale, NJ: Erlbaum Associates.
- Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, 381, 520–522.
- Traub, R. D., Spruston, N., Soltesz, I., Konnerth, A., Whittington, M. A., & Jefferys, G. R. (1998). Gamma-frequency oscillations: a neuronal population phenomenon, regulated by synaptic and intrinsic cellular processes, and inducing synaptic plasticity. *Progress in Neurobiology*, 55, 563–575.
- Tyler, C. W., & Kontsevich, L. L. (1995). Mechanisms of stereoscopic processing: stereoattention and surface perception in depth reconstruction. *Perception*, 24, 127–153.
- van Der Werf, Y. D., Witter, M. P., & Groenewegen, H. J. (2002). The intralaminar and midline nuclei of the thalamus, anatomical and functional evidence for participation in processes of arousal and awareness. *Brain Research*, 39, 107–140.
- Vanduffel, W., Tootell, R. B., & Orban, G. A. (2000). Attention-dependent suppression of meta-bolic activity in the early stages of the macaque visual system. *Cerebral Cortex*, 10, 109–126.

- van Vreeswijk, C., & Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural Computation*, 10, 1321–1371.
- Vitevitch, M. S., & Luce, P. A. (1999). Probabilistic phonotactics and neighborhood activation in spoken word recognition. *Journal of Memory and Language*, 40, 374–408.
- Vladusich, T., Lafe, F., Kim, D.-S., Tager-Flusberg, H., & Grossberg, S. (2010). Prototypical category learning in high-functioning autism. *Autism Research*, 3, 226–236.
- Von Restorff, H. (1933). Über die Wirkung von Bereichsbildungen im Spurenfeld (the effects of field formation in the trace field). *Psychologie Forschung*, 18, 299–334.
- Wanning, A., Stanisor, L., & Roelfsema, P. R. (2011). Automatic spread of attentional response modulation along Gestalt criteria in primary visual cortex. *Nature Neuroscience*, 14, 1243–1244.
- Warren, R. (1970). Perceptual restoration of missing speech sounds. *Science*, 167, 392–393.
- Warren, R., & Obusek, C. (1971). Speech perception and phonemic restorations. *Perception & Psychophysics*, 9, 358–362.
- Warren, R., & Sherman, A. (1974). Phonemic restorations based on subsequent context. *Perception & Psychophysics*, 16, 150–156.
- Warren, R., & Warren, R. (1970). Auditory illusions and confusions. *Scientific American*, 223, 30–36.
- Washburn, D. A., & Rumbaugh, D. M. (1991). Ordinal judgments of numerical symbols by macaques (*Macaca mulatta*). *Psychological Science*, 2, 190–193.
- Watanabe, T., Nanez, J. E., & Sasaki, Y. (2001). Perceptual learning without perception. *Nature*, 413, 844–848.
- Weber, A. J., Kalil, R. E., & Behan, M. (1989). Synaptic connections between corticogeniculate axons and interneurons in the dorsal lateral geniculate nucleus of the cat. *Journal of Comparative Neurology*, 289, 156–164.
- Wespapat, V., Tennigkeit, F., & Singer, W. (2004). Phase sensitivity of synaptic modifications in oscillating cells of rat visual cortex. *Journal of Neuroscience*, 24, 9067–9075.
- Wienke, D., & Buydens, L. (1995). Adaptive resonance theory based neural networks—the “ART” of real-time pattern recognition in chemical process monitoring. *Trends in Analytical Chemistry*, 14, 398–406.
- Wills, T. J., Cacucci, F., Burgess, N., & O’Keefe, J. (2010). Development of the hippocampal cognitive map in preweanling rats. *Science*, 328, 1573–1576.
- Wilson, M. A., & McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science*, 261, 1055–1058.
- Winters, B. D., Bartko, S. J., Saksida, L. M., & Bussey, T. J. (2007). Scopolamine infused into perirhinal cortex improves object recognition memory by blocking the acquisition of interfering object information. *Learning & Memory*, 14, 590–596.
- Wittmer, L. L., Dalva, M. B., & Katz, L. C. (1997). Reciprocal interactions between layers 4 and 6 cells in ferret visual cortex. *Society for Neuroscience Abstracts*, 651(5), 1668.
- Wunsch, D. C., Caudell, T. P., Capps, C. D., Marks, R. J. II, & Falk, R. A. (1993). An optoelectronic implementation of the adaptive resonance neural network. *IEEE Transactions on Neural Networks*, 4, 673–684.
- Xu, Z., Shi, X., Wang, L., Luo, J., Zhong, C.-J., & Lu, S. (2009). Pattern recognition for sensor array signals using fuzzy ARTMAP. *Sensors and Actuators B: Chemical*, 141, 458–464.
- Yang, S., Heinen, S., & Missal, M. (2008). The effects of microstimulation of the dorsomedial frontal cortex on saccade latency. *Journal of Neurophysiology*, 99(4), 1857–1870.
- Yazdanbakhsh, A., & Grossberg, S. (2004). Fast synchronization of perceptual grouping in laminar visual cortical circuits. *Neural Networks*, 17, 707–718.
- Zaki, S. R., Nosofsky, R. M., Jessup, N. M., & Unversagt, F. W. (2003). Categorization and recognition performance of a memory impaired group: evidence for single-system models. *Journal of the International Neuropsychological Society*, 9, 394–406.
- Zeki, S., & Shipp, S. (1988). The functional logic of cortical connections. *Nature*, 335, 311–317.
- Zhang, N., & Kezunovic, M. (2007). A real time fault analysis tool for monitoring operation of transmission line protective relay. *Electric Power Systems Research*, 77, 361–370.
- Zhang, Y., Suga, N., & Yan, J. (1997). Corticofugal modulation of frequency processing in bat auditory system. *Nature*, 387, 900–903.
- Zoccolan, D., Kouh, M., Poggio, T., & DiCarlo, J. J. (2007). Trade-off between object selectivity and tolerance in monkey inferotemporal cortex. *Journal of Neuroscience*, 27, 12292–12307.