13th ICCRTS: C2 for Complex Endeavors

"SOVEREIGN: An Autonomous Neural System for Goal-oriented Mobile Robot Navigation"

Topics: (3) Modeling and Simulation, (8) C2 Architectures,
and (9) Collaborative Technologies for Network-Centric Operations

William P. Gnadt
Lockheed Martin Corporation
MS2- Syracuse
497 Electronics Parkway
EP5-MD58
Syracuse, NY 13221
Telephone: +1 315 456 4460
william.gnadt@lmco.com

SOVEREIGN: An Autonomous Neural System for Goal-oriented Mobile Robot Navigation

Abstract

Command and control (C2) systems employ autonomous platforms for such tasks as force protection, surveillance, and search and rescue. The result is a growing need for adaptive control systems to focus data collection efforts, and reduce operator workloads, while addressing navigational demands. The SOVEREIGN (Self-Organizing, Vision, Expectation, Recognition, Emotion, Intelligent, Goal oriented Navigation) neural model embodies these capabilities, and is tested in a simplified 3D virtual reality environment. SOVEREIGN includes several interacting subsystems which model complementary properties of cortical *What* and *Where* processing streams. As a mobile robot, or animat, explores an environment, visual inputs are processed by networks that are sensitive to visual form and motion in the *What* and *Where* streams, respectively. Position-invariant and size invariant recognition categories are learned by real-time incremental learning in the *What* stream. Estimates of target position relative to the animat are computed in the *Where* stream, and can activate approach movements toward the target. Motion cues from animat locomotion can elicit head-orienting movements to bring a new target into view. Movement and visual category sequences are stored in working memories which can trigger learning of sensory and motor sequence categories, or plans. Rewarded plans effect a transition to efficient goal-oriented planned control.

Keywords: Adaptive Resonance Theory, Animat, Goal-oriented Behavior, Mobile Robot, Navigation, Neural networks, Self-organization, What and Where processing

1. Introduction

1.1. Three Basic Design Themes.

This article describes the SOVEREIGN (Self-Organizing, Vision, Expectation, Recognition, Emotion, Intelligent, Goal-oriented Navigation) neural model to clarify how an animal, or animat, can learn to reach valued goal objects through planned sequences of navigational movements. The SOVEREIGN model embodies a self-organizing control system that attempts to learn and perform such behaviors autonomously. As the name SOVEREIGN indicates, this control system unifies visual, recognition, cognitive, cognitive-emotional, and motor competences. We believe that this is the first neural model that embodies and coordinates such a wide range of behavioral competences in an autonomous self-organizing control system that can operate in real time. Such capabilities represent key unsolved problems in mobile robotics, including tasking related to Command and Control (C2). These results have been briefly reported in Gnadt and Grossberg (2005a, 2005b, 2006) and reported in detail in Gnadt (2007).

SOVEREIGN contributes to three large themes about how the brain works. The first theme concerns how brains learn to balance between reactive and planned behaviors. During initial exploration of a novel environment, many reactive movements occur in response to unexpected and unfamiliar environmental cues (Leonard and McNaughton, 1990). These movements may initially appear to be locally random, as an animal orients toward and

approaches many local stimuli. As such an animal becomes familiar with its surroundings, it learns to discriminate between objects likely to yield a reward and those that lead to punishment. Such approach-avoidance behavior is often learned via a perception-cognition-emotion-action cycle in which an action and its consequences elicit sensory cues that are associated with them. Rewards and punishments affect the likelihood that the same actions will be repeated in the future. When objects are out of direct sensory range, multiple reactive exploratory movements may be needed to reach them. Eventually, reactive exploratory behaviors are replaced by more efficient planned sequential trajectories within a familiar environment. One of the main goals of SOVEREIGN is to explain how erratic reactive exploratory behaviors can give rise to organized planned behaviors, and how both reactive and planned behaviors may remain balanced so that planned behaviors can be carried out where appropriate, without losing the ability to respond quickly to novel reactive challenges.

The second design theme illustrates the hypothesis that advanced brains are organized into parallel processing streams with complementary properties (Grossberg, 2000a). Each stream's properties are related to those of a complementary stream much as a lock fits its key, or two pieces of a puzzle fit together. The mechanisms that enable each stream to compute one set of properties prevent it from computing a complementary set of properties. As a result, each of these streams exhibits complementary strengths and weaknesses. How, then, do these complementary properties get synthesized into a consistent behavioral experience? It is proposed how interactions between these processing streams overcome their complementary deficiencies and generate behavioral properties that realize the unity of conscious experiences. In this sense, pairs of complementary streams are the functional units because only through their interactions can key behavioral properties be competently computed. SOVEREIGN clarifies how these complementary properties interact together to control goal-orienting sequences of navigational behaviors. For example, it is well-known that there are What and Where (or Where/How) cortical processing streams (Goodale and Milner, 1992; Mishkin, Ungerleider and Macko, 1983; Ungerleider and Mishkin, 1982). In particular, key properties of the What and Where cortical processing streams seem to be complementary.

A third design theme underlying the SOVEREIGN model is that brains use homologous circuits to compute navigational and hand/arm movements. In other words, movements of the body and of the hand/arms are controlled by circuits that share many properties. This proposed homology clarifies how navigational and arm movements can be coordinated when a body moves with respect to a goal object with the intention of grasping or otherwise manipulating it using the hand/arm system.

A considerable body of neural modeling of arm movement trajectory control (e.g., the VITE model: Bullock and Grossberg, 1988; Bullock, Cisek, and Grossberg, 1998) suggests that cortical arm movement control circuits compute a representation of where the arm wants to move (i.e., a target position) and compare this with an outflow representation of where the arm is now (i.e., the present position) by computing a difference vector between target position and present position representations. The difference vector represents the direction and distance that the arm needs to move to realize its goal position. Basal ganglia volitional signals of various kinds, such as a GO signal, translate the difference vector into a motor trajectory of variable speed. Additional cortical, spinal, and cerebellar circuitry is needed to ensure that the brain generates the forces that are needed to actually carry out such a commanded trajectory (e.g., the FLETE model: Bullock and Grossberg, 1991; Contreras-Vidal, Grossberg, and Bullock, 1997).

A key difference between navigation and hand/arm movement control concerns how present position is calculated. Because the arm is attached to the body, present position of the arm

can be directly computed using outflow, or corollary discharge, movement commands that explicitly code the commanded arm position. In contrast, when a body moves with respect to the world, no such immediately available present position command is available. This difference requires more elaborate brain machinery to compute present position of the body in the world during navigational movements. The brain needs to use a variety of sensory cues, both proprioceptive and visual, to create a representation of present position that can be compared with representations of target position, so that a difference vector and volitional commands can move the body towards desired goal objects. In summary, both navigational movement in the world and movement of limbs with respect to the body use a difference vector computational strategy.

1.2. What SOVEREIGN Does.

SOVEREIGN's perceptual competences include on-line, albeit simplified, visual representations of a 3D virtual reality environment in which the model controls navigation. SOVEREIGN computes, in parallel, both visual form and motion information about the world. As in the brain, the visual form of objects is computed within the What cortical processing stream, whereas visual motion is computed within the Where cortical processing stream. In this way, the brain can process both what objects are and where and how to track and act upon them.

SOVEREIGN uses the visual form information to incrementally learn spatially-invariant and size-invariant object recognition categories to recognize visually perceived objects in the world. These recognition categories, in turn, learn to read out top-down attentive expectations of the visual objects that they code. Object categories in the What stream are spatially-invariant and size-invariant to prevent a combinatorial explosion from occurring in which each position and size of an object would need its own representation. The Where stream represents the spatial locations of these objects. In particular, visual motion information is used to guide reactive orienting movements and attention shifts to locations at which changes occur in SOVEREIGN's visual world. What-Where inter-stream interactions are needed to enable both recognition and acquisition of desired goal objects. These parallel streams help SOVEREIGN to balance between reactive and planned behaviors, in a manner that is further discussed below.

SOVEREIGN also includes cognitive processes, notably mechanisms to temporarily store sequences of events in working memory, and to learn sequential plans, or chunks, of these sequences with which to predict and control future planned behaviors. Parallel object and spatial working memories and sequential chunking networks are modeled. The object working memory and chunking network are in the model's What stream, and the spatial working memory and chunking network are in its Where stream. SOVEREIGN clarifies how these parallel cognitive processes cooperate to acquire desired goal objects that can only be reached through a sequence of actions, and to disambiguate sequential navigational decisions in contexts where only one of them would be insufficient.

Cognitive-emotional mechanisms include the role of rewards and punishments in shaping goal-oriented behaviors. In particular, reinforcement learning can influence which learned cognitive chunks will be attended and selected to elicit behaviors that acquire desired goals within a familiar environment. Learned interactions between cognitive and emotional representations, notably motivationally-mediated attention, play an important role in this context-sensitive selection process.

The SOVEREIGN model thus contributes solutions to three key problems: (1) How an animal, or animat that embodies biologically-inspired designs, learns to balance between reactive and planned behaviors in a task-appropriate way. (2) How plans are learned during erratic

reactive behaviors in such a way that, after learning, they can be read out fluently at the correct times and in the correct spatial contexts. (3) How, in particular, an animat coordinates its reactive and planned behaviors so that its perception-cognition-emotion-action cycles of exploration, real-time vision, learned recognition, sequential working memory storage, learning of sequential plans, reinforcement learning, and planned action sequences are activated when needed as the animat navigates novel and familiar environments.

2. SOVEREIGN Model

2.1. Approach-Orienting Navigation and Learning in a 3D Virtual Reality Environment.

The SOVEREIGN model simulates these processes for an animat that experiences a 3D visual world in a virtual reality environment. This world is the relatively simple spatial format of a plus-maze (Munn, 1950). Although simple, this environment tests in a clear way many of the critical competences that the animat needs to achieve. Much of the problem's difficulty arises because an animat may navigate the maze in different ways, including different speeds and directions of movement, on successive learning trials. Despite this variability of each experience of the maze, the animal can learn to navigate the maze to achieve valued goals in an efficient way. For our purposes, it is sufficient to assume that a learning trial starts after placing the animat in the maze, at the end of one arm. The goal location, in one of the other three arms, is baited with a cue that the animal finds rewarding. By shrouding the top of the maze, only route-based visual and motor cues can be used for navigation (O'Keefe and Nadel, 1978). Thus the model does not attempt to explain how spatial navigation, as supported by hippocampal place cells, is achieved. Only one visual cue is assumed to be visible at a time, at the end of each maze arm, from any location within the maze. A schematic diagram of the experimental setup appears in Figure 1a.

A sequence of images from the 3D virtual reality simulation during reactive approach toward a visual cue appears in Figure 1b. As the animat approaches the choice point, a competitive struggle occurs between the salience of form and motion signals. Suppose that the form signals have led to previous object category learning and have led to positive reinforcement that increases their motivational salience. Such motivational salience enhances the strength of the form representation through attentional feedback. Then the form signals may more effectively compete with the motion signals to determine the animat's momentary behavior. If the form cues win the competition, then the animat can continue to carry out an approach movement that is consistent with its recognition. If the motion signals win the competition, then they may trigger a reactive head-orienting movement to the right or left at a choice point, revealing another source of form signals at the end of an adjacent corridor. The outcome of this form-motion competition is sensitive to navigational variations that change from trial to trial. The sequence of visual scenes that are processed during a typical head-orienting behavior is illustrated in Figure 1c. An alternation of approach and orienting movements is characteristic of the animat's exploration of a novel environment.

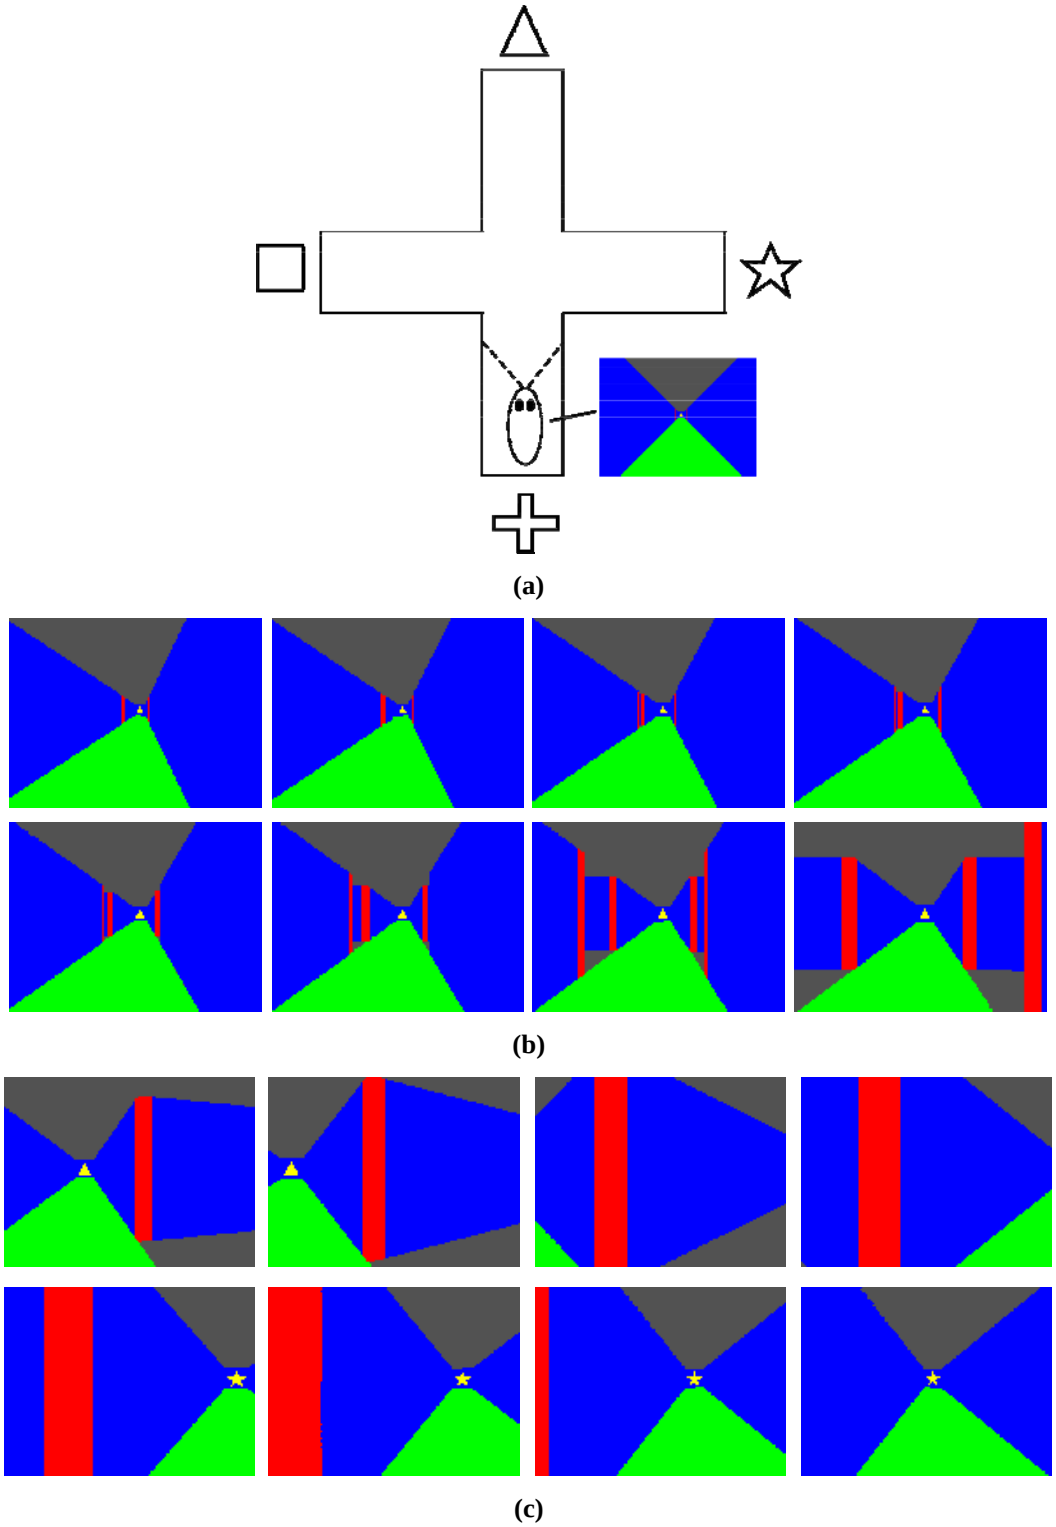

Figure 1. (a) The 3D graphical simulation generates perspective-views from any location within the plus-maze. **(b)** Snapshots from the 3D virtual reality simulation depict changes in the scene during reactive homing toward the triangle cue. **(c)** During reactive approach to the triangle cue, visual motion signals trigger a reactive head orienting movement to bring the star cue into view.

2.2. Parallel Visual and Motor Working Memory and Chunking Networks.

The animat's control system is split into a number of subsystems shown in the macrocircuit of Figure 2. The primary input is via the visual system. The Visual Form and Motion System processes visual cues within the What and Where cortical streams, respectively. The What cortical stream learns position-invariant and size-invariant object categories via on-line incremental learning. The Where stream computes measures of the relative location of visual cues from the animat. In particular, the distance and direction of the animat from a prescribed visual cue are used to cause approach movements towards that visual cue, or from memory. Motion cues result from the animat's self-motion, and determine whether the animat will make a left or right turn, and how big a turn, instead of continuing to approach a target cue.

These visual form and motion signals compete for control of the animat's approach-orienting behaviors within the Visual Form and Motion System. Learned visual categories can be amplified in strength, and thereby more probably attended, by feedback from motivational centers, called Drive Representations, through learned reinforcement-motivational feedback loops that embody their value as events that predict desired rewards. For this to happen, the invariant object categories are amplified by motivational signals that draw attention to them, and amplify, in turn, the approach commands corresponding to that object's position relative to the animat. Such a motivational amplification requires What-Where inter-stream interactions between position-invariant and position-variant information.

When a motivationally-modulated form cue wins, approach persists. When a motion cue wins, an orienting movement often begins. When motion cues are balanced in strength relative to the present gaze direction, the net left vs. right orienting signal is zero, after opponent competition between the opposite directions takes place. A form cue can then win with high probability. However, a suitably strong left/right motion cue difference can win the form-motion competition and direct the Motor Approach and Orienting System to initiate a head-orienting movement in the favored direction. Target position information for approach behaviors, and motion information for head-orienting behaviors, is relayed from the Visual Form and Motion System to the Motor Approach and Orienting System (Figure 2), where they direct body-approach movements or head-and-body orienting movements.

How does the animat know where a target cue is with respect to its current position? As noted in Figure 2, proprioceptive and vestibular signals provide the ground truth upon which the animat's location is estimated relative to its starting point, and with respect to targets in its environment. Proprioceptive and vestibular cues are capable of guiding animat navigation in a familiar environment even in the dark, and can modify movements quickly to cope with uneven or slippery terrains. Visual cues are also used during navigation to estimate body and head position and displacement relative to the animat. These visual signals are associated with, and adaptively calibrated with respect to, the representations that are activated by proprioceptive and vestibular motor signals. These multiple sources of information work together to more accurately guide movements under varying conditions than any one source of positional signals could.

Estimates of spatial displacement compute the NET body displacement and head rotation from a starting point. Sequences of such approach-orienting displacements represent a path that can command an animat to move from a starting location to a goal location in a maze.

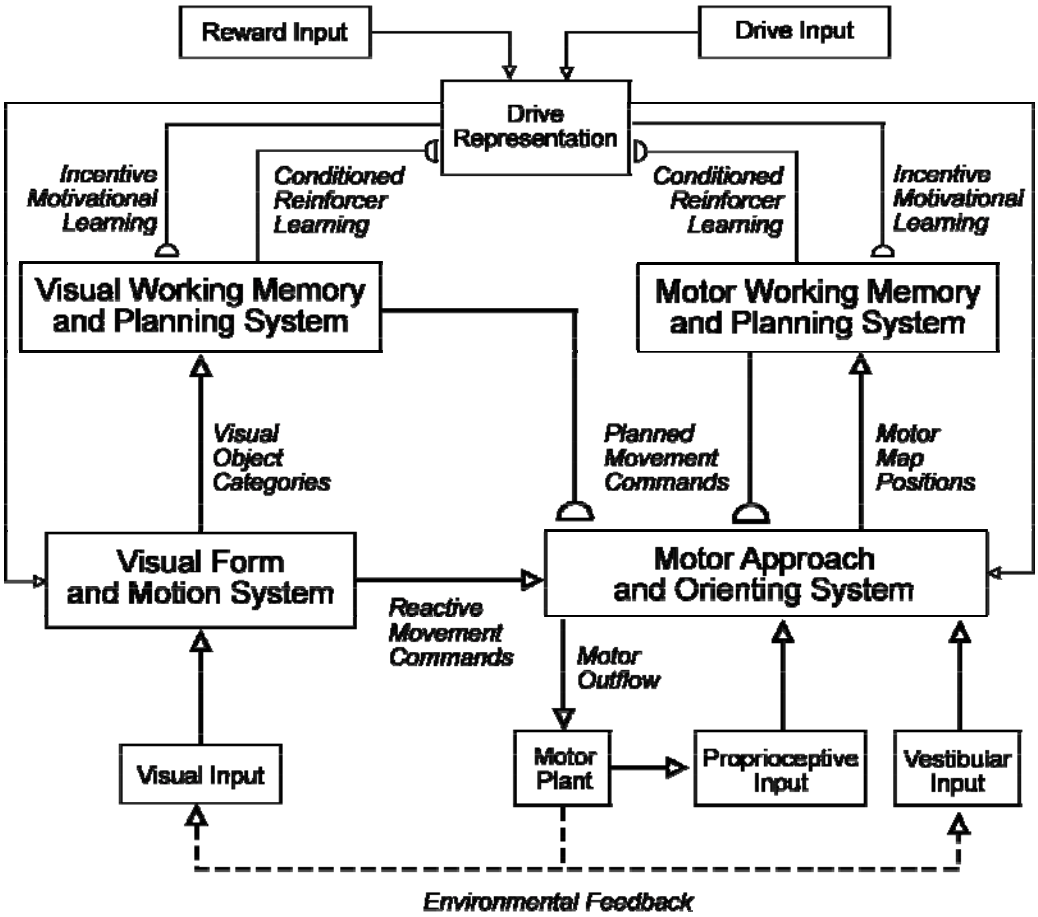

Figure 2. The key interactions between components of the SOVEREIGN model are shown in this flow diagram. See text for details.

The Visual Working Memory and Planning System temporarily stores sequences of visual object categories in short-term working memory. It also categorizes, or chunks, sequences of stored object categories. Chunks are learned that are sensitive to object sequences of variable lengths. These visual list chunks can learn to activate motor commands in the Motor Approach and Orienting System via top-down learning. The motor commands encode approach-orient movements. The Visual Working Memory and Planning System operates in parallel with a Motor Working Memory and Planning System that temporarily stores sequences of motor commands in working memory. It also categorizes, or chunks, sequences of stored motor commands. These motor list chunks can also learn to activate approach-orient commands within the Motor Approach and Orienting System.

Why are visual and motor list chunks both needed? Together these parallel visual and motor working memories can disambiguate decisions that only one of them, acting alone, may find ambiguous. For example, the sequences of approach distances and head turns in two different environments may be the same, but their sequence of visual cues may be different. In a different environment, the sequence of visual cues may be the same, but their sequences of motor actions may differ. The visual and motor working memories induce the learning of list chunks, or sequential planning cells, that are sensitive to their respective object and action sequences, and

can read out a prediction of the next motor command. The sequence that can disambiguate two different environments will typically win over one that cannot.

Rewards and punishments can modulate animat learning and determine what visual representations are attended and what motor plans are executed. Upon receiving reward, the active chunks are associated with active drives and actions. Drive inputs represent the animat's internal motivational state, and reward inputs represent valued inputs from the external environment. Both types of inputs combine in Drive Representations, which are most highly activated when a drive input representing a strong internal need combines with either a primary reward or a conditioned reinforcer input from the Visual Form and Motion System (Figure 2). After such a combination of cognitive and emotional learning occurs, when the animat sees a familiar sensory cue under a prescribed motivational state, it can recall a motivationally-compatible plan to reach the site of previous reward. Repeated, random exploration of the environment hereby effects a gradual transition from reactive to more efficient, planned control that leads the animat to its various motivated goals. Due to the selective role of motivational feedback, the animat is capable of learning to carry out different plans to satisfy different motivational goals even in response to the same sensory cues.

Visual and motor list chunks may learn to activate different approach-orient commands under different motivational states. How can a single chunk give rise to multiple responses? How this occurs can be seen by noting that emotional centers are often organized into *opponent* affective processes, such as fear and relief, and that oppositely valenced rewards can be conditioned to these opponent channels (Grossberg, 1984b, 2000b). These opponent-processing emotional circuits are called *gated dipoles*. In such a circuit, habituative transmitters "gate", or multiply, signal processing in each of the channels of the opponent "dipole." The response amplitude and sensitivity to external reinforcing inputs and internal drive inputs of these opponent-processing emotional circuits are calibrated by an arousal level and chemical transmitters that slowly inactivate, or habituate, in an activity-dependent way.

Sensory and cognitive representations, no less than emotional representations, can be organized into opponent channels with habituative ON and OFF cells. Unexpected events can trigger a burst, or sudden increment, of nonspecific arousal. When such an arousal burst is received on top of the baseline tonic arousal input of a normal dipole, it can cause an antagonistic rebound of activity in the OFF channel. In other words, the sensory, cognitive, or emotional hypothesis that is represented in a dipole's activity can be disconfirmed by an unexpected event. An unexpected event can hereby reset ongoing processing and lead to a shift of attention. SOVEREIGN expands the gated dipole mechanism into a gated multipole, which can select between multiple opponent drive channels. Each channel, whether representing an exploratory or consummatory drive state, can be associated with a particular learned response.

SOVEREIGN embodies a system synthesis and further development of biologically-derived neural networks that have been mathematically and computationally characterized elsewhere. These include LAMINART and FORMOTION models for form and motion processing (Berzhanskaya, Grossberg, and Mingolla, 2007; Cao and Grossberg, 2005; Grossberg, Mingolla, and Viswanathan, 2001; Grossberg and Yazdanbakhsh, 2005; Raizada and Grossberg, 2003), ART fast incremental learning classifiers (Carpenter, et. al., 1992), STORE working memories (Bradski, Carpenter, and Grossberg, 1994), Masking Field sequence chunking networks (Cohen and Grossberg, 1986, 1987; Grossberg and Myers, 2000; Grossberg and Pearson, 2007), Gated Dipole opponent processes (Grossberg, 1980, 1984a; Grossberg and Seidman, 2006), CogEM cognitive-emotional circuits for reinforcement learning (Grossberg and Merrill, 1992, 1996; Grossberg, 2000), Spectral Timing circuits for adaptively timed learning

(Grossberg and Merrill, 1992; Fiala, Bullock, and Grossberg, 1996), and volitional (GO) and endogenous (ERG, Endogenous Random Generator) gates to release consummatory and exploratory behaviors, respectively (Bullock and Grossberg, 1988; Gaudiano and Grossberg, 1991; Pribe, Grossberg, and Cohen, 1997). We are not aware of any other autonomous agent that has yet integrated this range of self-organizing biological competences.

2.3. Reactive Exploration.

The following sequence illustrates the functional flow of the visual input system during reactive exploration in the plus–maze of Figure 1. North designates the vertical direction, with South, East and West following accordingly. For definiteness, assume that the animat is placed into the maze and that all extra–maze cues are suppressed. Furthermore, the animat is motivated under both an exploratory and a hunger drive. The drive and reward inputs to the Drive Representation and then into the Visual and Motor Working Memory and Planning Systems are shown in Figure 2. The exploratory drive is assumed to be excited by an Endogenous Random Generator, or ERG, which is an internal arousal source. Such a source is active when the animat is placed into a new environment. The exploratory drive is inhibited by consummatory drive activity that can support realization of a valued goal. The animat receives a reward (e.g., food) upon reaching the goal location, which is located at the end of the West arm. We show how reactive visual signals during exploration eventually lead the animat toward the goal location, and reinforcement signals strengthen the association between stored plan items and the current motivational state. A step-by-step description of the model under reactive visual guidance follows.

Suppose that, by chance, the animat starts the maze shifted to the left of the corridor, with its head facing slightly to the right of the visual cue (Figure 3a). The left shift reduces the distance to motion cues on the left side of the maze. Because of this positional bias, motion signals within the Visual Form and Motion System (Figure 2) will receive a strong leftward bias. These assumptions are used to demonstrate an exploratory trial which ensures that the animat makes its first head-orienting movement toward the goal location. During the experimental trial, the animat moves forward (Figure 3b), turns left (Figure 3c) and approaches the goal location (Figure 3d) under reactive control.

Movement is organized into orienting and approach movements. In particular, a visually-activated motor command from the Visual Form and Motion System triggers a Motor Outflow command (Figure 2) that specifies a head-orienting angle to align the head with the triangle target. The resulting signals activate the Motor Plant (Figure 2), which converts the movement command into a physical displacement. A head-orienting movement towards the triangle target is thereby initiated. The head turn continues until the NET head-orienting displacement equals the commanded displacement angle.

When the animat faces the triangle cue, a Motor Outflow command from the Visual Form and Motion System activates the Motor Plant (Figure 2) to initiate an approach movement toward the triangle cue. When the Motor Plant converts the commanded approach movement into a physical displacement, the animat's body is passively aligned with the head during an approach movement to maintain a stable posture. Such dynamic stability control is assumed to be present, but is beyond the scope of this work.

During the approach movement, the Motor Approach and Orienting System continues to compute the NET head and body displacement toward the visual target cue. In the absence of competing cues, the body-approach movement could continue until the animat reaches the cue. However, the Visual Form and Motion System processes both form and motion signals while the

animat continues to move. A sufficiently strong motion signal in the model's visual periphery can win a competition between Parvo form target locations and Magno motion cues. If a motion cue wins, then it can terminate the approach movement and trigger a reactive head-orienting movement away from the visual target cue.

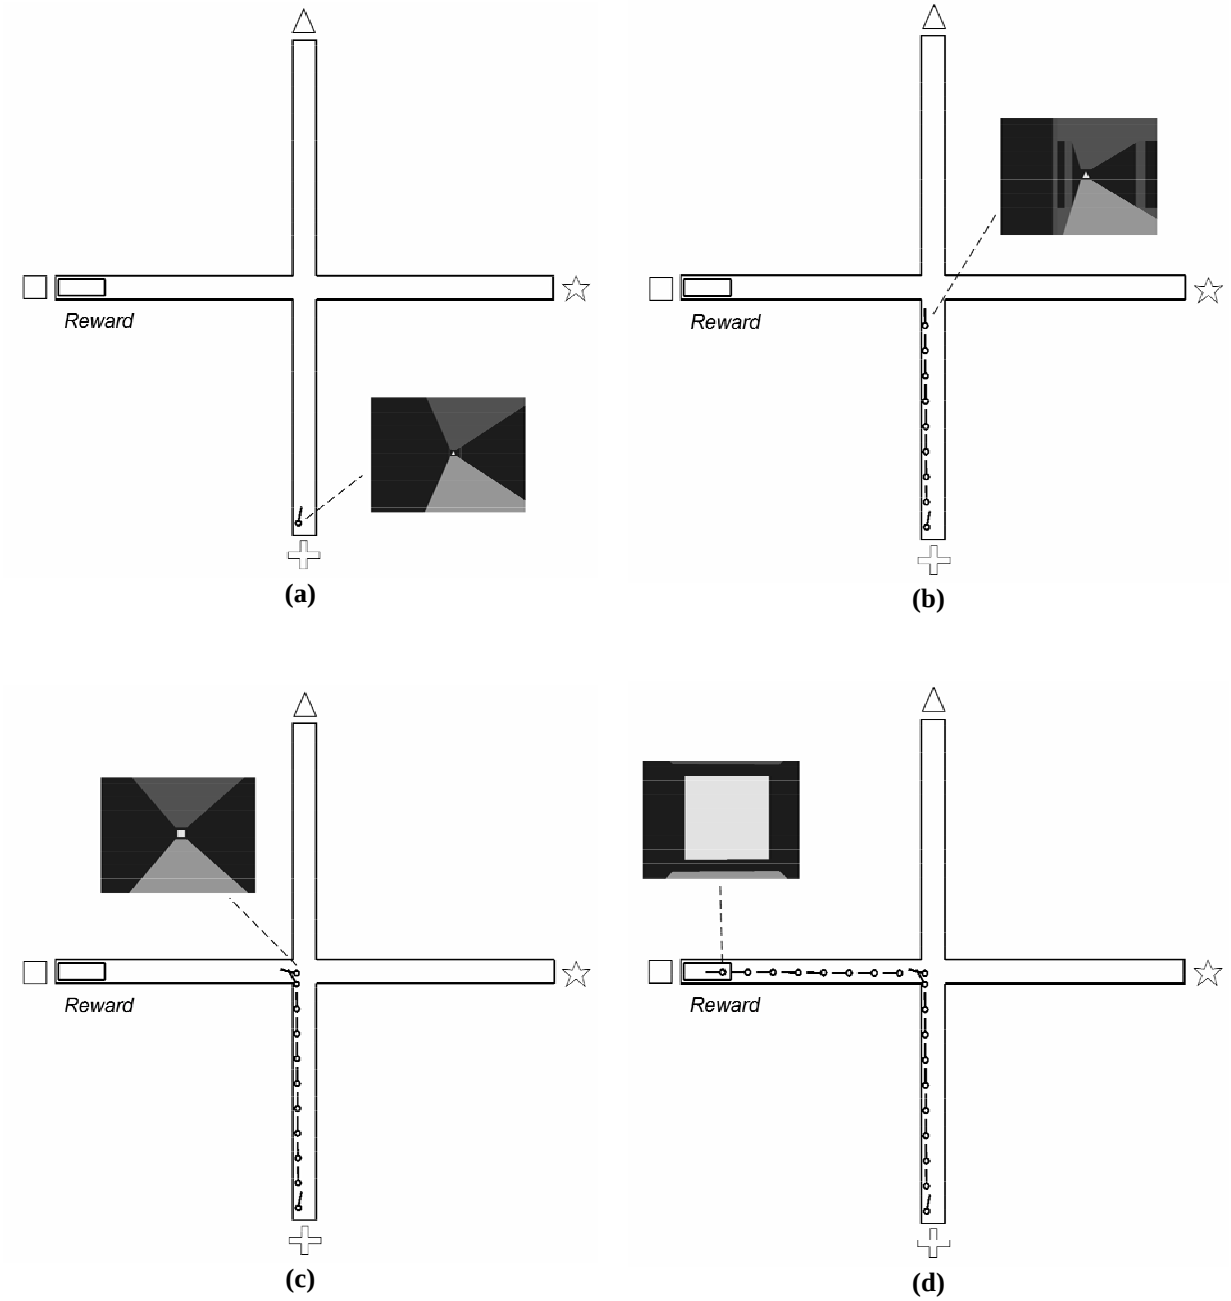

Figure 3. (a) Animat position and head direction facing the triangle cue at the start of the trial. Perspective-views of the 3D virtual reality scene at key locations within the maze are shown by a dashed line. **(b)** Animat position and head direction while approaching the triangle cue and nearing the choice point. **(c)** Animat position and head direction after a head orienting turn toward the square cue. **(d)** Animat position and head direction after reaching the goal location at the square cue.

As noted above, when the animat starts in a position that is shifted to the left side of the corridor, as in Figure 3a, motion signals in the left visual hemifield are stronger than those in the right hemifield. Left vs. right motion signals accumulate in the Visual Form and Motion System. When the left motion signal is sufficiently strong relative to the right motion cue and the form signal, a reactive head-orienting command is sent to the Motor Approach and Orienting System.

As the animat carries out these movements, it learns an invariant object category, or chunk, for the triangle visual cue. Top-down signals from the Visual Working Memory and Planning System (Figure 2) corresponding to the Triangle chunk learn the NET body approach and orienting movements computed by the Motor Approach and Orienting System. The triangle cue hereby learns to predict the Forward-Left body movement. The Forward-Left body movements are also stored in the Motor Working Memory and Planning System.

After the animat turns left, invariant preprocessing and learned ART categorization within the Visual Form and Motion System encode a 3D representation of the square cue. This 3D representation is stored in the Visual Working Memory and Planning System (Figure 2), while the NET body displacement in the Motor Approach and Orienting System is reset to prepare for the next movement. Then the cycle of computing the NET head and body displacements begin again, as the animat navigates toward the square cue.

The square cue is at the rewarded location. When the animat reaches this location, it receives a reward, such as food. The active hunger drive representation is then associated with the currently active plan chunks stored in both the Visual Working Memory and Planning System and the Motor Working Memory and Planning System (Figure 2). In particular, the visual Triangle-Square list chunk is learned and associated with the hunger drive representation. In addition, signals from the Triangle-Square chunk learn the NET body orienting and approach movement computed by the Motor Approach and Orienting System, and thereby learns to predict the Forward body movement that brings the animat to the square cue after it turns left in the West arm of the maze. Figure 4 summarizes this sequence of events.

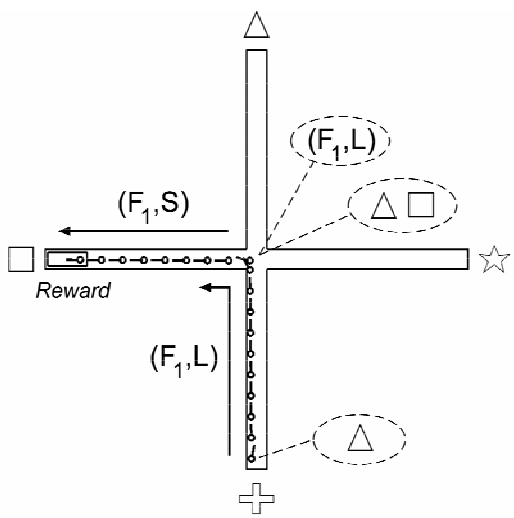

Figure 4. An initial maze trial in which the animat is under reactive visual guidance is shown in this diagram. An approach movement toward the triangle cue is interrupted by motion signals to the left. After a reactive head orienting movement, the square cue comes into view. After approaching the square cue, the rewarded location is reached and adaptive weights are adjusted to strengthen the association between the forward-left-forward sequence and the current motivational state. The arrows and symbols (F_1,L) and (F_1,S), along with the triangle and triangle-square symbols in the dotted ellipses, summarize that a forward-left movement sequence with a forward distance of F_1 is associated with the Triangle list category, and a forward-straight movement also with a forward distance of F_1 is associated with the Triangle-Square list category.

One additional point should be made: All animat behaviors are motivated by some Drive Representation (Figure 2). During initial exploratory activities, an exploratory drive is active. As learning occurs, the exploratory drive is supplanted by the consummatory motivational sources that correspond to the reward; e.g., the hunger drive when the animat is rewarded by food. These processes are now described in greater detail.

2.4. Visual Form and Motion System.

The Visual Form and Motion System processes signals from the What Parvo cortical processing stream and the Where Magno cortical processing stream (Figure 5). This separation of functionality endows the animat with three major capabilities. First, the animat can utilize target object recognition and cognitive-emotional conditioning circuitry to learn, choose, and execute motivationally–compatible movements within an overall plan. Second, the animat can use form information to localize visual references, or beacons, to measure its progress over varying terrain. Finally, the animat can process motion boundaries generated during movement toward a choice point within a maze. As the animat nears a choice point, its field–of–view and the intensity of boundary-derived motion signals increase, which can trigger a reactive head–orienting movement. The visual system also drives several important control signals within the model, as described below.

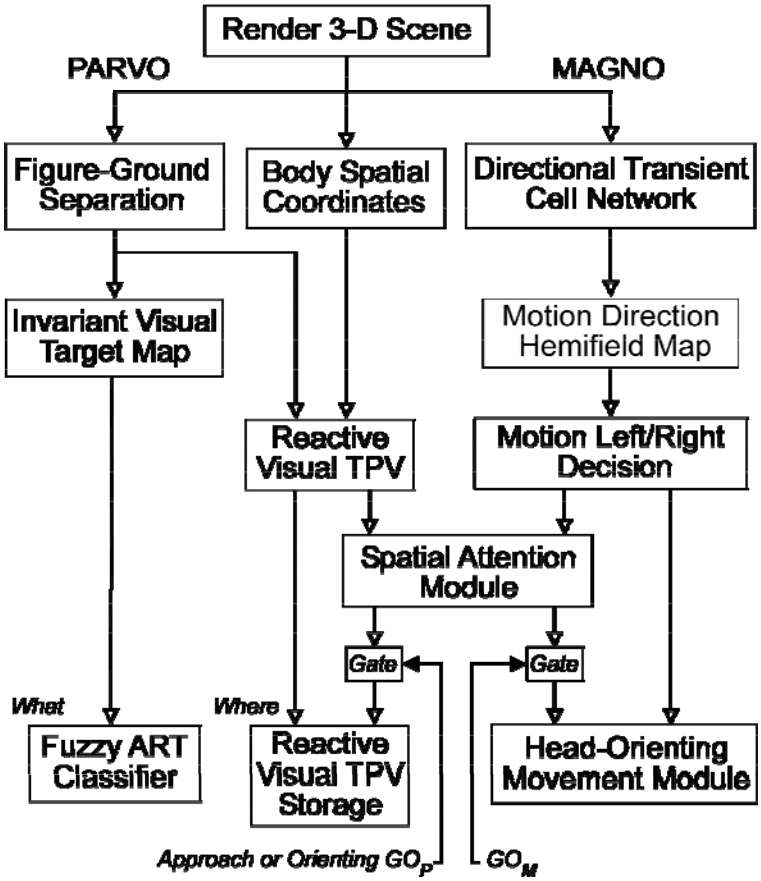

Figure 5. The Visual Form and Motion System flow diagram depicts the stages of visual processing in the model. See text for details.

The visual environment is simulated in a virtual reality environment by rendering 3D chromatic scenes as 2D "snapshots" at regular intervals during head-orienting and body-approach movements. As indicated in Figure 1, the visual environment is simplified in SOVEREIGN, which focuses on the various learning and navigational aspects of sequential goal-oriented navigation. A visual target object is separated from the background by a two-stage Figure-Ground Separation module that is within the Parvo stream (Figure 5, left stream). At present, the first processing stage is accomplished in a simple way by using visual targets that are yellow (Figure 1), or have the grayscale corresponding to yellow, and are thereby selected from the background. The second processing stage selects object boundaries via convolution with a 2D Laplacian-of-Gaussian filter. Future model developments will include more sophisticated neural models for 3D vision and figure-ground separation (Cao and Grossberg, 2005; Fang and Grossberg, 2007; Grossberg and Yazdanbakhsh, 2005; Kelly and Grossberg, 2000).

When an object falls within the visual field and it is separated from its background, the coordinates of its centroid, in the 2D image plane, are computed (cf., Russell, 1932) and passed to the Reactive Visual TPV module (Figure 5). The Reactive Visual TPV module converts the centroid from image plane coordinates to head-centered spatial coordinates by using the perspective transformation (Schilling, 1990). The Body Spatial Coordinates module computes the angle between the head and body, before combining this information with the target coordinates in the Reactive Visual TPV module to compute the body-centered distance and angle coordinates of the visual target. The Reactive Visual TPV module updates the Reactive Visual TPV Storage module until the volitional Approach and Orienting GO signal (GO_P) releases a head-orienting or body-approach movement (Figure 5). The head-orienting movement brings the visual target to the center of gaze. Such a transformation into body–centered coordinates can be learned by using a more elaborate network (Greve et. al., 1993; Grossberg et. al., 1993; Guenther et al., 1994).

The left path of the Parvo stream in Figure 5 is devoted to learning a size-invariant and position-invariant object category representation of a visual target within the Invariant Visual Target Map. In order to accomplish this, the figure-ground-separated visual target undergoes a log-polar transformation followed by Gaussian coarse-coding (Baloch and Waxman, 1991; Bradski and Grossberg, 1995). The log-polar transform computes a representation of the visual target object that is size-invariant and position-invariant. This invariant map representation of the target is then transformed into an object category, leading to further compression and invariance under modest changes in object shape, by using unsupervised incremental learning by a Fuzzy ART classifier (Carpenter et al., 1991). The Fuzzy ART classifier and Reactive Visual TPV Storage modules comprise What and Where cortical representations of visual target objects.

The Fuzzy ART classifier can be generalized in a future version of SOVEREIGN to enable learning of 3D target objects from one of multiple views. This requires additional processing stages to learn individual view categories which can be associatively linked to a view-invariant object category (Baloch and Waxman, 1991; Bradski and Grossberg, 1995; Fazl, Grossberg, and Mingolla, 2007).

2.5 Motor Approach and Orienting System.

As noted above, the Motor Approach and Orienting System directs body-approach and head-orienting movements (Figure 2). Cumulative estimates of each approach-orienting movement that is processed within the Motor Approach and Orienting System are stored in the Motor Working

Memory and Planning System (Figure 2). This section summarizes how these estimates are computed.

The Motor Approach and Orienting System flow diagram is shown in Figure 6. Target position information originates from one of two sources. First, it can be received from the body-centered distance and angle coordinates of the visual target object computed in the Reactive Visual TPV module (Figure 5). Second, it can be received from learned top–down signals from the processing stage that computes Motivated WHAT and WHERE Decisions (Figure 6). These decisions comprise responses from the animat's learned experience which are compatible with the current motivational state. An approximate measure of head-orienting and body displacement is computed by the NET module (Figure 6). Target position information flows from the Reactive Visual TPV to the Reactive Visual TPV Storage module. The NET activity is subtracted from the Stored TPV via learned weights to compute the Reactive Difference Vector, or DV, which represents the angle and distance to move. The learned weights from the NET activity are necessary to calibrate the DV activity. Similarly, learned top-down commands from the Motivated WHAT and WHERE Decisions activate the Planned DV, where NET movement signals are subtracted, yielding a planned angle and distance to move. Calibration of planned commands is accomplished entirely by the top-down adaptive weights. The Reactive DV and Planned DV are the first motor control stages which can elicit head-orienting and body-approach movements.

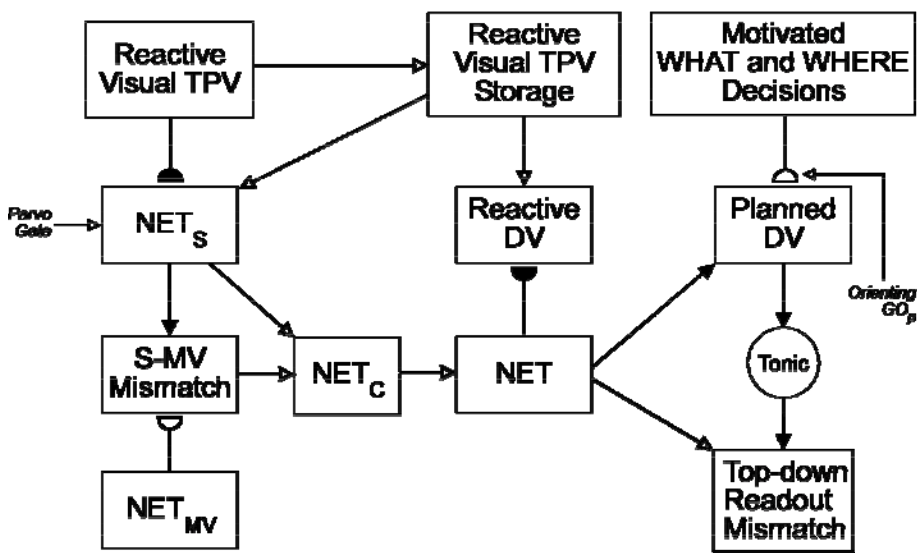

Figure 6. The Motor Approach and Orienting System flow diagram depicts the control hierarchy which generates motor outflow commands. See text for details.

The NET estimates of head-orienting and body-approach displacement requires multiple stages of processing to be computed (Figure 6). NET estimates during navigation replace the outflow present position estimates that are computed during hand/arm movements. The NET_S module (Figure 6) calculates this displacement using target positions computed by the Visual Form and Motion System (Figure 5). Body–centered spatial coordinates are denoted by an "s" subscript. The NET_S field activity encodes the net body movement toward a target in spatial coordinates by subtracting the Reactive Visual TPV activity from the Reactive Visual TPV Storage module activity.

Initially, target position information flows from the Reactive Visual TPV to the Reactive Visual TPV Storage module and a short burst of learning zeros the difference at the NET_S module. As the animat moves toward a target, updates to the Reactive Visual TPV Storage module cease and the Reactive Visual TPV decreases, thereby allowing the NET_S module activity to grow. Vestibular and proprioceptive feedback signals are integrated into distances by the NET_{MV} module (Figure 6).

Learning at the output of the NET_{MV} module calibrates the vestibulo-motor signals relative to the visual signals at the S-MV Mismatch module. This adaptive process uses a slow learning rate while visual signals are available from the Visual Form and Motion System (Figure 5). The resulting activity at the S-MV Mismatch module serves as a correction factor which can account for the animat's progress either without visual feedback (e.g., in the dark) or over uneven (e.g., slippery) terrain. When the NET_S and NET_{MV} module activity are identical, then the correction factor is zero and the S-MV Mismatch module activity is also zero.

The NET module combines signals from the NET_S and S-MV Mismatch modules into a robust sensory-motor representation of body displacement. The NET_S module is only active when Parvo signals are present in the Invariant Visual Target Map module (Figure 5). Learned weights from the NET module inhibit activity of the Reactive DV (Figure 6). When the animat has reached the target under visual guidance, the Reactive Visual TPV reaches zero and these adaptive weights are updated, thereby inhibiting the Reactive DV, and stopping the movement. On future trials, the Reactive DV module can be driven to zero by a calibrated level of activity in the NET module, regardless of whether visual input is available.

While under reactive control, visual target coordinates flow into the Reactive Visual TPV Storage module and activate the Reactive DV module, which initiates head and body movements. The Approach or Orienting GO_P control signals are activated when the Reactive or Planned DV command is released under volitional control. The activation of the Approach or Orienting GO_P allows the DV signals to initiate a head-orienting or body-approach movement. Updates to the Reactive Visual TPV Storage module (Figure 6) continue until the Approach or Orienting GO_P is activated. However, under planned control, Motivated WHAT and WHERE Decisions (Figure 6) learn to read out planned head-orienting and body-approach movements. Top-down commands are computed in the Planned DV module, which can initiate head and body movements in response to motivationally-compatible plan items. As then plan unfolds, NET increases until the Planned DV approaches zero, thus terminating a planned movement. The Top-down Readout Mismatch module compares the activity of the learned top-down command and the NET module. A sufficiently large discrepancy between these fields can elicit a control signal to select a different top-down signal from the Motivated WHAT and WHERE Decisions. For instance, the control signal is released when a planned response is interrupted by a strong motion signal which activates the Head-Orienting Movement module (Figure 5) and the animat turns away from the planned response direction.

3. A Key End-to-end SOVEREIGN Simulation

In the Motivated Choice Experiment, the animat learns the route to two different goal locations under two different motivational states. The simulation summary contains the following types of information: (1) Explanation of the experimental setup; (2) movement trajectories; (3) a step-by-step description of model dynamics; (4) snapshots of visual input at key moments; (5) multi-trial learning; and (6) summary of the model properties demonstrated.

3.1 Motivated Choice Experiment.

In this classic example of spatial learning, the animat learns the route to two different goal locations under two different motivational states. Specifically, the Forward-Left-Forward sequence when hungry leads to a food reward, whereas the Forward-Right-Forward sequence leads to a water reward when thirsty. Upon reaching the end of the goal arm, the animat is rewarded and long-term memory weights are updated using a slow learning rate. For the first five training trials of this sequence, the animat has a high hunger drive and is rewarded with food at the end of each trial. For the next five trials, the animat is thirsty and is rewarded with water at the end of each trial. The eleventh and final trial retests the response under the hunger drive to demonstrate that learning under each drive is preserved.

The diagram shown in Figure 7a shows the position, actions and local views seen by the animat during this training trial. Similarly, the diagram in Figure 7b shows the Visual and Motor Working Memory and Planning System plan chunks which are stored during this experimental trial. This simulation is similar to that presented in earlier examples. However, a summary is offered here for completeness. The animat starts this trial shifted to the left within the corridor, thereby increasing Magno signals in the left hemifield. The hungry animat learns to categorize the triangle cue and updates the Visual Working Memory and Planning System. It approaches the triangle cue under reactive control. As it nears the choice point, Magno signals trigger a head-orienting movement to the left bringing the square cue into view. The Triangle chunk is associated with the exploratory drive and can now sample the Forward-Left movement. The animat then learns to categorize the square cue. The Visual and Motor Working Memory and Planning System are updated and the animat approaches the square cue under reactive control. Upon reaching the food reward, all active chunks are associated with the hunger drive and the Forward-Left and Triangle-Square chunks sample the Forward-Straight movement command. The initial training trial is complete.

The diagram in Figure 7c shows the learned plan chunks and their associated motor responses which are gradually strengthened during training. After several trials, the hungry animat starts this trial centered in the corridor, yielding Magno signals which are balanced between left and right. After learning to categorize the triangle cue and updating the Visual Working Memory and Planning System, the Triangle chunk can read out the command to go Forward-Left via top-down signals. The planned command overrides reactive signals and the animat moves forward and turns left. Both the approach speed and Parvo gain are increased because the previously rewarded plan element has been reactivated. After the turn is complete, the Triangle chunk is again associated with the exploratory drive. This learning is triggered by the exploratory learning signal in the absence of explicit reward, and is activated after a head turn is completed. After learning to categorize the square cue, the Visual and Motor Working Memory and Planning System are updated and the Forward-Straight command is directly read out via top-down Motivated WHAT and WHERE Decision signals. After approaching under planned control, the animat is rewarded with food and the active chunks are associated with the hunger drive. Finally, the Forward-Left and Triangle-Square chunks can sample the Forward-Straight movement command.

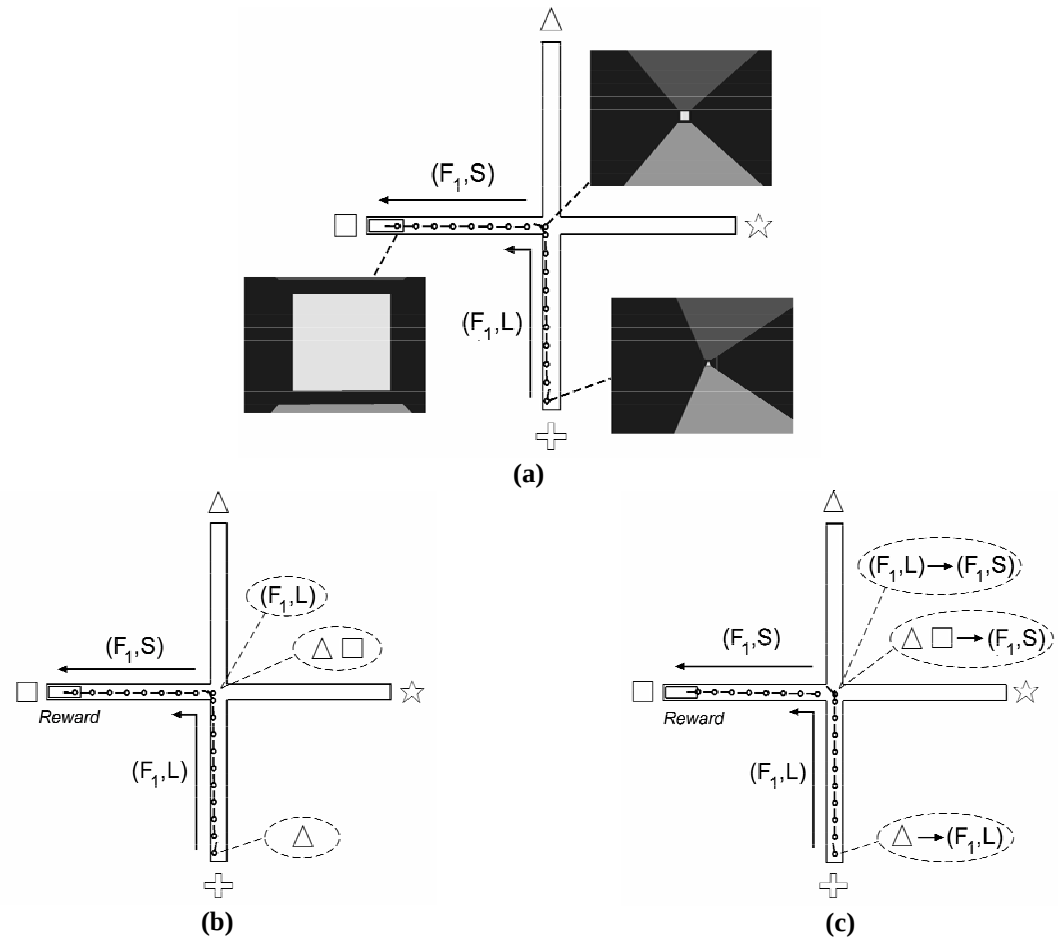

Figure 7. (a) Perspective views are shown for selected points during maze exploration toward the goal location in the left arm. **(b)** Each ellipse graphically depicts the short-term memory chunks represented in both the Visual and Motor Working Memory and Planning System during exploratory learning. **(c)** Each ellipse graphically depicts the short-term memory chunks and associated motor responses in both the Visual and Motor Working Memory and Planning System after exploratory learning.

The diagram in Figure 8a shows the position, actions and local views seen by the animat during the second phase of the maze experiment. Similarly, the diagram in Figure 8b shows the Visual and Motor Working Memory and Planning System plan chunks which are stored during this experimental trial. After several learning trials, the thirsty animat starts this trial centered in the corridor, yielding Magno signals which are balanced between the left and right sides of the visual field. After learning to categorize the triangle cue and updating the Visual Working Memory and Planning System, the Triangle chunk is able to read out the command to go Forward-Right via top-down Motivated WHAT and WHERE Decision signals. The planned command overrides reactive signals and the animat moves forward and turns right. Both the approach speed and Parvo gain are increased because the previously rewarded plan element has been reactivated. After the turn is complete, the Triangle chunk is again associated with the exploratory drive. After learning to categorize the star cue, the Visual and Motor Working Memory and Planning System are updated and the Forward-Straight command is directly read out via top-down signals. After approaching under planned control, the animat is rewarded with water and the active chunks are associated with the thirst drive. Finally, the Forward-Right and Triangle-Star chunks can sample the Forward-Straight movement command.

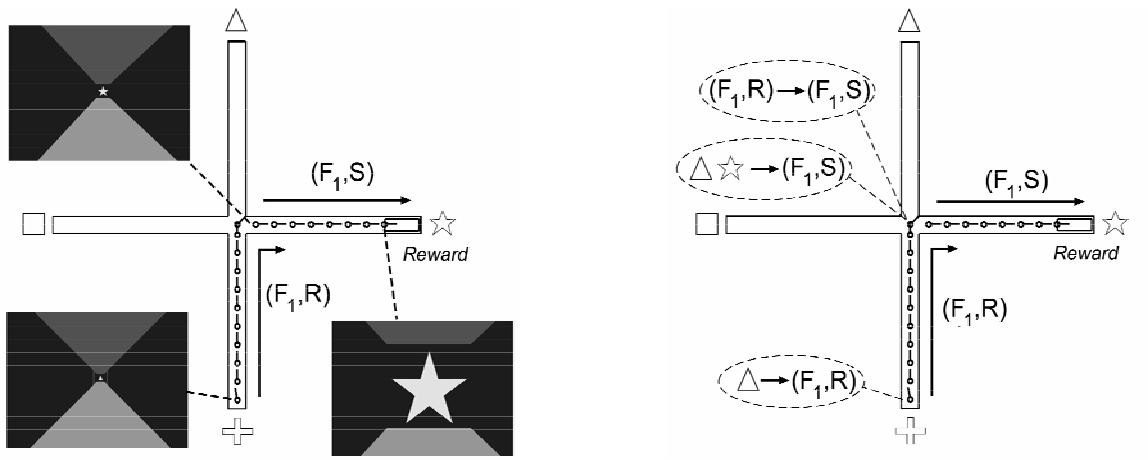

(a) **(b)**

Figure 8. (a) Perspective views are shown for selected points during maze exploration under planned control. **(b)** Each ellipse graphically depicts the short-term memory chunks and associated motor responses in both the Visual and Motor Working Memory and Planning System after exploratory learning.

Although the long-term memory (LTM) weights rapidly learn the appropriate responses, several trials are required for the weights to achieve asymptotic levels. The total level of conditioned reinforcement (CR) and incentive motivational learning is shown in the six panels of Figure 9. In the upper panel of Figure 9a, the CR adaptive weight from the Visual Working Memory and Planning System (Figure 2), denoted "SENSORY", to the hunger Drive Representation module is plotted as a function of training trial. Recall that the first five training trials comprise the Forward-Left-Forward movement sequence to receive a food reward while hungry. The next five training trials comprise the Forward-Right-Forward movement sequence to receive a water reward while thirsty. The eleventh, and final test trial demonstrates no loss of learning of the first sequence while hungry. In the middle panel of Figure 9a, the same CR data is shown for the Motor Working Memory and Planning System (Figure 2) to the hunger drive. In the bottom panel Figure 9a, the total learned CR versus training trial is shown. The CR levels increase and remain high as a consequence of the learning law employed by the system.

The learned CR to the thirst input of the Drive Representation module is similar to the hunger drive input, discussed above. In the upper panel of Figure 9b, the learned CR from the Visual Working Memory and Planning System (Figure 2), denoted "SENSORY", to the thirst input of the Drive Representation module is plotted as a function of training trial. In the middle panel, the same CR data is shown for the Motor Working Memory and Planning System (Figure 2) to the thirst drive. Finally, in the bottom panel, the total learned CR versus training trial is shown. The CR levels only increase after the fifth trial, when the thirst drive is active, and the animat receives a water reward for reaching the goal.

The exploratory drive performs an important role in spatial learning. The exploratory drive input to the Drive Representation module is inhibited by consummatory drive activity. However, it is never completely inactive, thereby allowing some conditioned reinforcement (CR) and incentive motivational learning to it. In the upper panel of Figure 9c, the learned CR from the Visual Working Memory and Planning System (Figure 2), denoted "SENSORY", to the exploratory input of the Drive Representation module is plotted as a function of training trial. In the middle panel the same CR data is shown for the Motor Working Memory and Planning System (Figure 2) to the exploratory drive. Finally, in the bottom panel, the total learned CR versus training trial is shown. During the first five trials, the Star chunk and Forward-Left motor

sequence were conditioned to the exploratory drive. This resulted in low CR levels for this first phase of training and a rapid increase in the CR levels after the fifth trial. The effect seen is that, during reactivation of the Star chunk, learned CR contributes to an increase in the Drive Representation module activity. This increase, in turn, is sampled by the Visual and Motor Working Memory and Planning System modules.

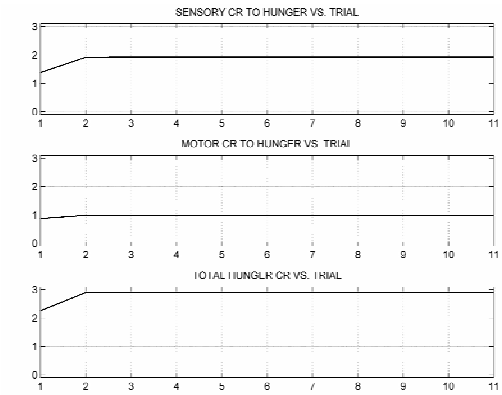

(a)

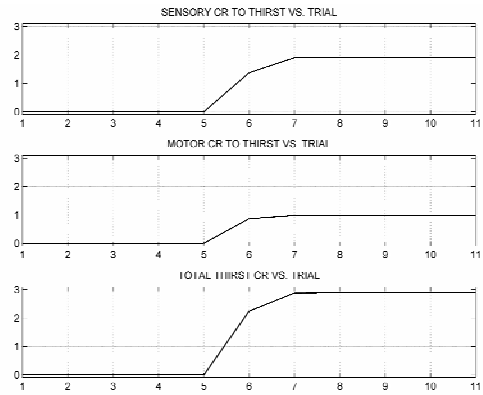

(b)

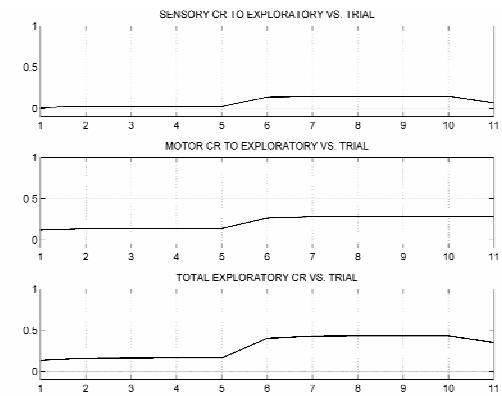

(c)

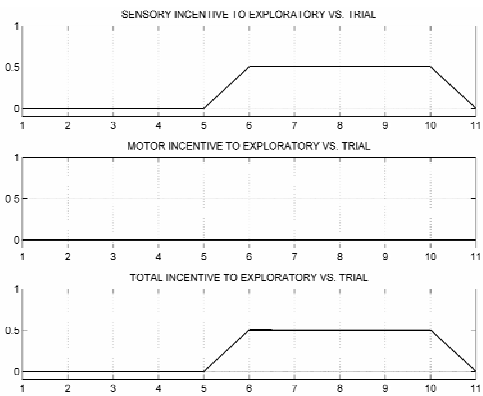

(d)

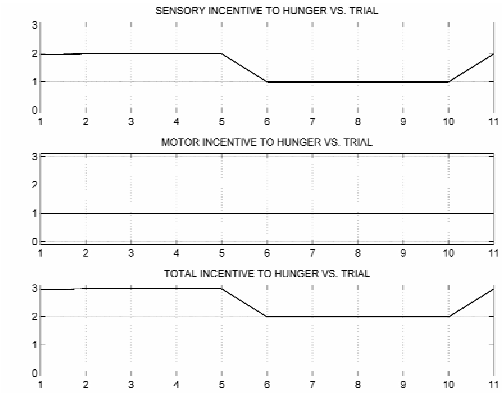

(e)

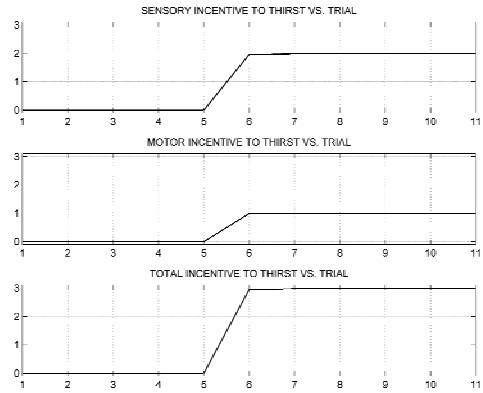

(f)

Figure 9. This series of plots show the acquisition of learned conditioned reinforcement and incentive motivation over a series of 11 experimental trials navigating to goal locations under either the hunger or thirst drives (see text for details). In each panel, the Visual Memory System (sensory), Motor Memory System (motor) and total learned associations are plotted vs. trial. **(a)** Acquisition of sensory and motor conditioned reinforcement to the hunger drive vs. trial. **(b)** Acquisition of sensory and motor conditioned reinforcement to the thirst drive vs. trial. **(c)** Acquisition of sensory and motor conditioned reinforcement to the exploratory drive vs. trial. **(d)** Acquisition of sensory and motor incentive motivation from the Sensory-Drive Heterarchy module exploratory channel vs. trial. **(e)** Acquisition of sensory and motor incentive motivation from the Sensory-Drive Heterarchy module hunger channel vs. trial. **(f)** Acquisition of vision and motor incentive motivation from the Sensory-Drive Heterarchy module thirst channel vs. trial.

Incentive motivational signals aid in the selection of Visual and Motor Memory System output channels. These signals are important for enabling a chunk to read out the correct motor command under different motivational states. The total learned incentive motivation from the exploratory channel of the Sensory-Drive Heterarchy module is shown as a function of trial in Figure 9d. The upper, middle and lower panels are as discussed above. The acquisition of learned incentive motivation to the hunger and thirst drives is shown in Figures 9e and 9f. The upper, middle and lower panels are as above. The initially large incentive motivation signals under the hunger drive decrease later in the experiment when learning under the thirst drive. This results from the ongoing exploratory learning to the prepotent output channel which weakens its prior conditioning to the hunger drive. In summary, this experimental setup demonstrates the capacity for the SOVEREIGN model to learn different routes to a goal location under different motivational states.

4. General Discussion and Conclusions

The model has several notable strengths relative to other available models, including the following ones: First, it provides an end-to-end model that includes on-line vision, visual recognition learning and categorization, working memory storage of sequences of visual and motor categories, learning of sequential cognitive and motor plans, cognitive-emotional interactions whereby reinforcement learning can select plans that can attain a currently valued goal, and balancing of visually reactive exploratory vs. planned movement decisions, based upon the relative salience of bottom-up and top-down information through time. Second, unlike various other models (e.g., Barto and Sutton, 1981; Dayan, 1987; Schmajuk, 1990), no explicit spatial goal gradient, proportional to the spatial distance from the goal, is required to guide goal-oriented sequential behavior in SOVEREIGN. Third, list chunks provide a compact context-sensitive code for learning plans to navigate a large number of different routes. Fourth, reliable, single–trial learning of a maze can occur if the animat happens to find the goal location during an exploratory trial. Fifth, the animat can respond to the same sequence of visual or motor events in different ways to achieve different goals when different drives are prepotent.

One weakness in the current version of SOVEREIGN is that its navigational behaviors are all route-based. The model does not yet include mechanisms of spatial navigation (O'Keefe and Dostrovsky, 1971; O'Keefe and Nadel, 1978) such as the role of hippocampal place fields, head-direction cells, and the theta rhythm (e.g., Burgess et al., 1995). Such a development would require an understanding of how place fields form, notably the role of entorhinal grid cells in their formation (e.g., Hafting, Fyhn, Molden, Moser, and Moser, 2005), which other modeling research is currently investigating (e.g., Fuhs and Touretzky, 2006; Gorchetchnikov and Grossberg, 2007).

References

Baloch, A. A. and Waxman, A. M. (1991). Visual learning, adaptive expectations, and behavioral conditioning of the mobile robot MAVIN. *Neural Networks, 4*, 271–302.

Barto, A. G. and Sutton, R. S. (1981). Landmark learning: An illustration of associative search. *Biological Cybernetics, 42*, 1–8.

Berzhanskaya, J., Grossberg, S. and Mingolla, E. (2007). Laminar cortical dynamics of visual form and motion interactions during coherent object motion perception. *Spatial Vision*, in press.

Bradski, G., Carpenter, G.A., and Grossberg, S. (1994). STORE working memory networks for storage and recall of arbitrary temporal sequences. *Biological Cybernetics, 71,* 469-480.

Bradski, G. and Grossberg, S. (1995). Fast learning VIEWNET architectures for recognizing 3-D objects from multiple 2-D views. *Neural Networks, 8,* 1053-1080.

Bullock, D., Cisek, P. and Grossberg, S. (1998). Cortical networks for control of voluntary arm movements under variable force conditions. *Cerebral Cortex, 8,* 48-62.

Bullock, D. and 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. and Grossberg, S. (1991). Adaptive neural networks for control of movement trajectories invariant under speed and force rescaling. *Human Movement Science, 10,* 3-53.

Burgess, N., Recce, M., and O'Keefe, J. (1995). Hippocampus – Spatial models. In *The Handbook of Brain Theory*. Bradford Books / MIT Press, Cambridge, MA.

Cao, Y. and Grossberg, S. (2005). A laminar cortical model of stereopsis and 3D surface perception: Closure and da Vinci stereopsis. *Spatial Vision, 18,* 515-578.

Carpenter, G.A., Grossberg, S., Markuzon, N., Reynolds, J.H., and 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., and Rosen, D. B. (1991). Fuzzy ART: Fast stable learning and categorization of analog patterns by an adaptive resonance system. *Neural Networks, 4,* 759–771.

Cohen, M.A. and Grossberg, S. (1986). Neural dynamics of speech and language coding: Developmental programs, perceptual grouping, and competition for short-term memory. *Human Neurobiology, 5,* 1-22.

Cohen, M. A. and Grossberg, S. (1987). Masking fields: A massively parallel neural architecture for learning, recognizing, and predicting multiple groupings of patterned data. In Grossberg, S. (Ed.), *Neural Networks and Natural Intelligence*, Chapter 7, pp. 317–367. The MIT Press, Cambridge, MA.

Contreras-Vidal, J.L., Grossberg, S., and Bullock, D. (1997). A neural model of cerebellar learning for arm movement control: Cortico-spino-cerebellar dynamics. *Learning and Memory, 3,* 475-502.

Dayan, P. (1987). Navigating through temporal difference. In *Neural Information Processing Systems, 3,* pp. 464–470. Morgan Kaufmann Publishers, San Mateo, CA.

Fang, L. and Grossberg, S. (2007). From stereogram to surface: How the brain sees the world in depth. Technical Report CAS/CNS TR-06-013, Boston University. Submitted for publication.

Fazl, A., Grossberg, S., and Mingolla, E. (2007). View-invariant object category learning, recognition and search: How spatial and object attention are coordinated using surface-based

attentional shrouds. Technical Report CAS/CNS TR-07-011, Boston University. Submitted to *Cognitive Psychology*, April, 2007.

Fiala, J.C., Bullock, D., Grossberg, S. (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.

Fuhs, M.C. and Touretzky, D.S. (2006). A spin glass model of path integration in rat medial entorhinal cortex. *Journal of Neuroscience, 26*, 4266 – 4276.

Gaudiano P. and Grossberg S. (1991). Vector associative maps: Unsupervised real-time error-based learning and control of movement trajectories. *Neural Networks, 4*, 147-183.

Gnadt, W. (2007). SOVEREIGN: An autonomous neural system for incrementally learning to navigate towards a rewarded goal. *Ph.D. thesis*, Boston University.

Gnadt, W. and Grossberg, S. (2005a). SOVEREIGN: *A Self-Organizing, Vision, Expectation, Recognition, Emotion, Intelligent, Goal-oriented Navigation system.* In Douglas Blank and Lisa Meeden (Eds.), Developmental Robotics: Papers from the 2005 Spring Symposium, March 21-23, pp. 106-110. American Association for Artificial Intelligence, Menlo Park, California.

Gnadt, W. and Grossberg S. (2005b). *SOVEREIGN: A Self-Organizing, Vision, Expectation, Recognition, Emotion, Intelligent, Goal-oriented Navigation system.* Dynamical Neuroscience XIII: Computational Cognitive Neuroscience satellite symposium of the annual Society for Neuroscience meeting, November 10-11, 2005.

Gnadt, W. and Grossberg S. (2006). *SOVEREIGN: A Self-Organizing, Vision, Expectation, Recognition, Emotion, Intelligent, Goal-oriented Navigation system.* Tenth International Conference on Cognitive and Neural Systems, Boston, Massachusetts, May 2006.

Goodale, M. A. and Milner, D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences, 15*, 20–25.

Gorchetchnikov, A. and 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.

Greve, D., Grossberg, S., Guenther, F., and Bullock, D. (1993). Neural representations for sensory–motor control, I: Head–centered 3–D target positions from opponent eye commands. *Acta Psychologica, 82*, 115–138.

Grossberg, S. (1980). How does a brain build a cognitive code? *Psychological Review, 87*, 1-51.

Grossberg, S. (1984a). Some psychophysiological and pharmacological correlates of a developmental, cognitive and motivational theory. In R. Karrer, J. Cohen, and P. Tueting (Eds.), *Brain and information: Event related potentials*. New York: New York Academy of Sciences, pp.58-142.

Grossberg, S. (1984b). Some normal and abnormal behavioral syndromes due to transmitter gating of opponent processes. *Biological Psychiatry, 19*, 1075–1118.

Grossberg, S. (2000a). The complementary brain: Unifying brain dynamics and modularity. *Trends in Cognitive Sciences, 4*, 233-246.

Grossberg, S. (2000b). The imbalanced brain: From normal behavior to schizophrenia. *Biological Psychiatry, 48*, 81-98.

Grossberg, S., Guenther, F., Bullock, D., and Greve, D. (1993). Neural representations for sensory–motor control, II: Learning a head–centered visuomotor representation of 3–D target position. *Neural Networks, 6*, 43–67.

Grossberg, S. and 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. and Merrill, J.W.L. (l996). The hippocampus and cerebellum in adaptively timed learning, recognition, and movement. *Journal of Cognitive Neuroscience, 8,* 257-277.

Grossberg, S., Mingolla, E. and Viswanathan, L. (2001). Neural dynamics of motion integration and segmentation within and across apertures. *Vision Research, 41,* 2521-2553.

Grossberg, S. and Myers, C.W. (2000). The resonant dynamics of speech perception: Interword integration and duration-dependent backward effects. *Psychological Review, 107,* 735-767.

Grossberg, S., & Pearson, L. (2007). Laminar cortical dynamics of cognitive and motor working memory, sequence learning and performance: Towards a unified theory of how the cerebral cortex works. Submitted for publication.

Grossberg, S. and Seidman, D. (2006). Neural dynamics of autistic behaviors: Cognitive, emotional, and timing substrates. *Psychological Review, 113,* 3, 483-525.

Grossberg, S. and Yazdanbakhsh, A. (2005). Laminar cortical dynamics of 3D surface perception: Stratification, transparency, and neon color spreading. *Vision Research, 45***,** 1725-1743.

Guenther, F. H., Bullock, D., Greve, D., and 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.

Hafting, T., Fyhn, M., Molden, S., Moser, M.H., and Moser, E.I. (2005). Microstructure of the spatial map in the entorhinal cortex, *Nature, 436,* 801 – 806.

Kelly, F.J. and Grossberg, S. (2000). Neural dynamics of 3-D surface perception: Figure-ground separation and lightness perception. *Perception and Psychophysics, 62***,** 1596-1619.

Leonard, B. and McNaughton, B. L. (1990). Spatial representation in the rat: Conceptual, behavioral, and neurophysiological perspectives. In Kesner, R. P. and Olton, D. S. (Eds.), *Neurobiology of Comparative Cognition,* Chapter 13, pp. 363–422. Lawrence Erlbaum Associates, Hillsdale, NJ.

Mishkin, M., Ungerleider, L. G., and Macko, K. A. (1983). Object vision and spatial vision: Two cortical pathways. *Trends in Neurosciences, 6,* 414–417.

Munn, N. L. (1950). *Handbook of Psychological Research on the Rat.* Houghton Mifflin Co., Boston, MA.

O'Keefe, J.M. and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research, 34,* 171 – 175.

O'Keefe, J. and Nadel, L. (1978). *The Hippocampus as a Cognitive Map.* Oxford University Press, New York, NY.

Pribe, C., Grossberg, S., and Cohen, M.A. (1997). Neural control of interlimb oscillations, II: Biped and quadruped gaits and bifurcations. *Biological Cybernetics, 77,* 141-152.

Raizada, R. and Grossberg, S. (2003). Towards a theory of the laminar architecture of cerebral cortex: Computational clues from the visual system. *Cerebral Cortex, 13,* 100-113.

Russell, J. T. (1932). Depth discrimination in the rat. *Journal of Genetic Psychology, 40,* 136–159.

Schilling, R. J. (1990). *Fundamentals of Robotics: Analysis and Control.* Prentice Hall, Englewood Cliffs, NJ.

Schmajuk, N. A. (1990). Role of the hippocampus in temporal and spatial navigation: An adaptive neural network. *Behavioral Brain Research, 39,* 205–229.

Ungerleider, L. G. and Mishkin, M. (1982). Two cortical visual systems: Separation of appearance and location of objects. In Ingle, D. L., Goodale, M. A., and Mansfield, R. J. W. (Eds.), *Analysis of Visual Behavior,* pp. 549–586. The MIT Press, Cambridge, MA.