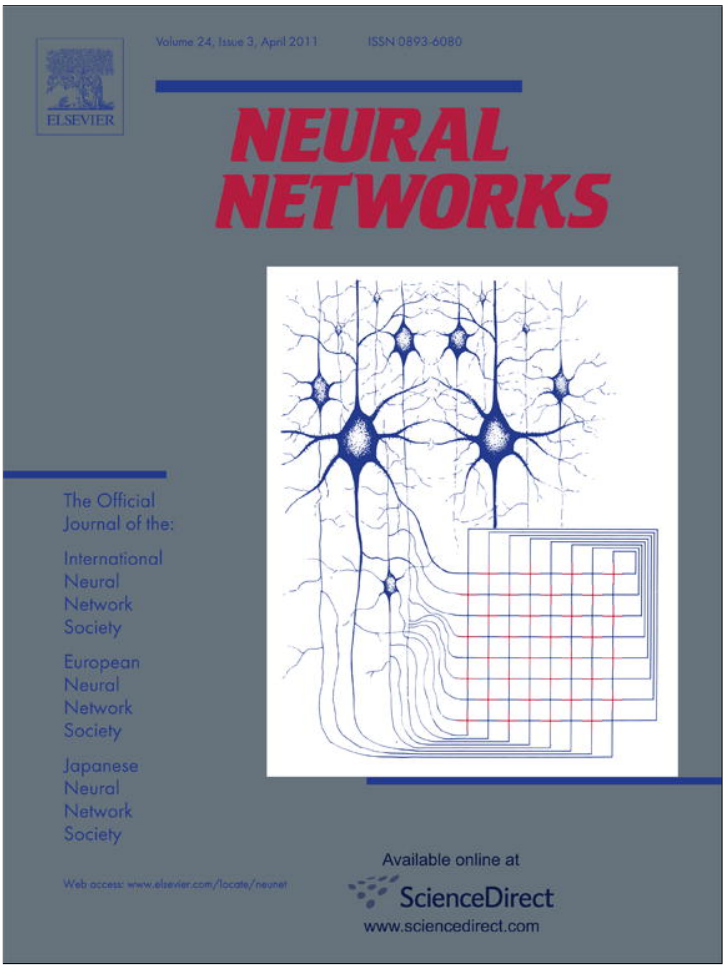

Neural Networks 24 (2011) 257–266

Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

A neuromorphic model of spatial lookahead planning

Richard Ivey, Daniel Bullock*, Stephen Grossberg

Department of Cognitive and Neural Systems, Center for Adaptive Systems and Center of Excellence for Learning in Education, Science, and Technology, Boston University, 677 Beacon Street, Boston, MA 02215, United States

ARTICLE INFO

Article history:
Received 26 March 2010
Received in revised form 1 November 2010
Accepted 3 November 2010

Keywords:
Spatial cognition
Gradient climbing
Lookahead planning
Working memory
Parietal cortex
Inhibition-of-return

ABSTRACT

In order to create spatial plans in a complex and changing world, organisms need to rapidly adapt to novel configurations of obstacles that impede simple routes to goal acquisition. Some animals can mentally create successful multistep spatial plans in new visuo-spatial layouts that preclude direct, one-segment routes to goal acquisition. *Lookahead* multistep plans can, moreover, be fully developed before an animal executes any step in the plan. What neural computations suffice to yield preparatory multistep lookahead plans during spatial cognition of an obstructed two-dimensional scene? To address this question, we introduce a novel neuromorphic system for spatial lookahead planning in which a feasible sequence of actions is prepared before movement begins. The proposed system combines neurobiologically plausible mechanisms of recurrent shunting competitive networks, visuo-spatial diffusion, and inhibition-of-return. These processes iteratively prepare a multistep trajectory to the desired goal state in the presence of obstacles. The planned trajectory can be stored using a primacy gradient in a sequential working memory and enacted by a competitive queuing process. The proposed planning system is compared with prior planning models. Simulation results demonstrate system robustness to environmental variations. Notably, the model copes with many configurations of obstacles that lead other visuo-spatial planning models into selecting undesirable or infeasible routes. Our proposal is inspired by mechanisms of spatial attention and planning in primates. Accordingly, our simulation results are compared with neurophysiological and behavioral findings from relevant studies of spatial lookahead behavior.

1. Introduction

The ability to create preparatory spatial plans in complex novel environments where perceptually indirect actions are necessary to obtain a goal is a critical competence for successful interaction with real-world environments. Spatial planning has been hypothesized to co-develop with spatial skills, in both phylogeny and ontogeny, and to support a broad range of human intellectual pursuits (Diamond, 1985; Matthews, 1996). There is a broad range of higher-level flexible spatial planning behaviors in humans and other primates (Buttelmann, Carpenter, Call, & Tomasello, 2008; Carder, Handley, & Perfect, 2004; Miller & Cohen, 2001; Ward & Allport, 1997). Many anatomically and functionally disparate spatial skills share the common conceptual objective of generating and executing a spatial trajectory to transfer one or more objects from their initial state to a desired goal configuration. We introduce a neuromorphic model for generation of such multistep, goal-directed lookahead trajectories in novel 2D visuo-spatial environments.

Neurodegenerative diseases (Cohen & Freedman, 2005; Ersche, Clark, London, Robbins, & Sahakian, 2006; Sahakian et al., 1995) or damage can result in dysexecutive syndrome, a spectrum of deficits characterized by behavioral impulsivity and myopia, in which behavior generation is dominated by immediate stimuli and characterized by response perseveration in situations where task demands change (Carder et al., 2004; Ciaramelli, 2007; Dias, Robbins, & Roberts, 1996; Walker, Mikheenko, Argyle, Robbins, & Roberts, 2006) or when perceptually indirect actions need to be taken to achieve a goal (Carder et al., 2004; Carder, Handley, & Perfect, 2008; Colvin, Dunbar, & Grafman, 2001). These data indicate that specific neural mechanisms are responsible for the generation of flexible goal-directed behaviors in novel visuo-spatial environments, such as those shown in Fig. 1.

Sustained, spatially tuned activations during planning intervals are recordable in posterior parietal cortex (PPC) neurons (e.g., Andersen, Snyder, Bradley, & Xing, 1997; Chafee & Goldman-Rakic, 2000). Moreover, the PPC has long been associated with attentional control and spatial awareness, including an actor's ability to relate visible spatial locations to self and self-initiated actions to the local spatial layout. Unilateral PPC lesions reliably produce hemifield neglect syndromes (e.g., Committeri et al., 2007; He et al., 2007), in which the actor loses the ability to process or attend to locations and to plan actions in an entire spatial hemifield.

* Corresponding author. Tel.: +1 617 353 9486; fax: +1 617 353 7755.
E-mail address: danb@bu.edu (D. Bullock).

0893-6080/$ – see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neunet.2010.11.002

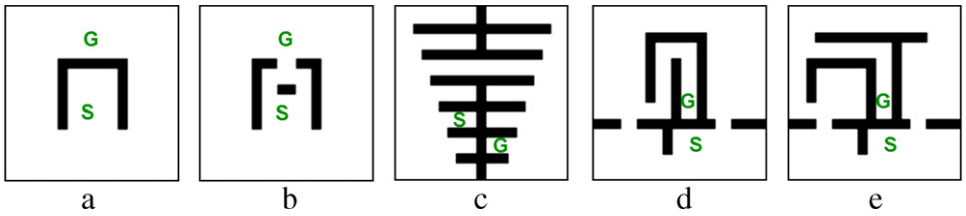

Fig. 1. Five examples of 2D spatial planning tasks that require indirect solutions and therefore nontrivial planning methods. *S* indicates the position of the start state, *G* the position of the goal. White is free space through which the trajectory may be generated. Black marks positions covered by obstacles.

In addition to its general role in supporting working memory operations on, and representations of, task-relevant information (Funahashi, Chafee, & Goldman-Rakic, 1993; Levy & Goldman-Rakic, 1999; Smith et al., 1995), dorsolateral prefrontal cortex (dlPFC) has also been specifically implicated in the preparation and selection of multistep action plans across a range of spatial and nonspatial tasks (Funahashi, 2001; Tanji & Hoshi, 2008). Electrophysiological and functional imaging studies have found high recruitment, and substantial task-dependent selectivity, of dlPFC activity in tasks that require spatial lookahead planning (Averbeck, Chafee, Crowe, & Georgopoulos, 2003; Boussaoud & Wise, 1993; Miller & Cohen, 2001; Miller, Erickson, & Desimone, 1996; Mushiake, Saito, Sakamoto, Itoyama, & Tanji, 2006; Mushiake, Saito, Sakamoto, Sato, & Tanji, 2001; Saito, Mushiake, Sakamoto, Itoyama, & Tanji, 2005).

These neurobiological and behavioral data indicate that flexible lookahead spatial planning in higher primates utilizes a constellation of explicit processes that are distinct from lower-level conditioned behaviors. To qualitatively model such behavior, we introduce a novel neurodynamic model of spatial lookahead planning that integrates neural modeling concepts of attentional diffusion, transient inhibition-of-return, and competitive selection to enable mental construction of a feasible spatial trajectory from an initial state to a given goal state in the presence of complex, novel configurations of free space and obstacles. An earlier version of this model has been briefly presented in Ivey, Bullock, and Grossberg (2008). Consideration of interactions of the proposed model with cooperative neural processes, such as working memory storage of feasible trajectories and sequential execution, is deferred until the Discussion.

We constrain our attention to the class of planning models that can be formulated as continuous-time dynamical systems, due both to their desirable properties of analyzability and implementation in analog circuitry, and to their potential as candidate models of planning in brains. Our focus on neurodynamical systems contrasts with non-neurodynamic planning algorithms from the control theory and artificial intelligence literature (e.g., Dijkstra, 1959; Koenig & Likhachev, 2002; Stentz, 1995).

2. Model description

Prior dynamical system models of planning have typically dealt with uncluttered environments or relied on repeated learning over multiple attempts. Models that compute difference vectors from start to goal (e.g., Bullock, Grossberg, & Guenther, 1993) are well supported by neurobiological data from simple tasks (e.g., Bullock, Cisek, & Grossberg, 1998), but by themselves do not cope with the indirection required for route planning around obstacles. Conditioned chaining models (e.g. Butz, Sigaud, & Gerard, 2003; Capdepuy, Polani, & Nehaniv, 2007; Fu & Anderson, 2006; Sutton & Barto, 1998; Tolman, 1959) can succeed along familiar (prelearned) paths, but fail in novel or altered layouts, and may require many learning trials to reach acceptable performance (Roitblat, 1994). Attractor/repeller models (Browning, Grossberg, & Mingolla, 2009; Eichhorn, 2005; Elder, Grossberg, & Mingolla, 2009; Fajen & Warren, 2003) are robust for choosing paths around point or convex obstacles, but are insufficient with concave obstacles. These models do not claim lookahead planning as a competence, but are nevertheless possible candidate models and discussed here to note that they are not sufficient to address the present task.

The current model can be regarded as a cognitive preprocessor for the stages assumed in some of these simpler models. It uses goal-sourced attentional diffusion and gradient climbing, illustrated in Fig. 2, as key operations in prospective route planning. The attentional diffusion process is embodied by a 2D topographic map of the environment with excitatory input at the goal. Through diffusion dynamics, the activity spreads throughout the topographic map. The activity is blocked and redirected by environmental obstacles. Note that nearly all motile organisms, including primitive bacteria (Macnab & Koshland, 1972), can detect *external* gradients in the world formed by diffusion processes (e.g., odor gradients), and can climb up or down them as needed to achieve goals. Cognitive gradient climbing (CGC) models propose that at least some primates have discovered how to represent spatial gradients internally, and to exploit such internal representations for mental planning in novel layouts that include concave obstacles.

The adjective "cognitive" is chosen here, similar in spirit to its usage in Tolman's *cognitive maps* (Tolman, 1948), to distinguish between the two broad classes of gradient climbing by life forms. The strategy of seeking high densities of an externally sensed gradient, such as a chemical concentration gradient, requires no internal genesis or representation of the gradient since it already exists in the world. The strategy also requires no preparatory gradient climbing operation since such evaluation of the gradient occurs in-the-loop with physical movement in the world. In contrast, the CGC model explored here proposes that some animals have evolved the capacity to generate a mental representation of a gradient that does not inherently exist in the world, and to iterate a mental process that climbs that mentally generated gradient. In psychology, the external class of models would be called "behavioristic", the latter class "cognitive". Indeed, it was precisely the postulation by Lashley (1951), Tolman (1948) and others regarding the existence and manipulation of internal representations, notably spatial maps, that led to the reemergence of a cognitive psychology after the failure of Watsonian behaviorism. Note that a key aspect of external gradient climbing is that it is a memoryless process. In contrast, when a mental operation fundamentally manipulates internal states, as does the current model, it is properly called cognitive. The current model proposes a process whereby a sequence of forthcoming actions is prepared through internal mechanisms, a hallmark of the cognitive, as opposed to the behaviorist, tradition. Although the non-CGC models noted above are insufficient to account for key aspects of primate intelligence exhibited in complex novel environments, we propose that CGC co-exists with phylogenetically older planning mechanisms that suffice when task demands are simpler.

The current CGC model of spatial lookahead planning is built from neurobiologically plausible mechanisms whereby an entire plan of sequential actions may be mentally constructed and prepared for enactment before any action is taken in a complex novel

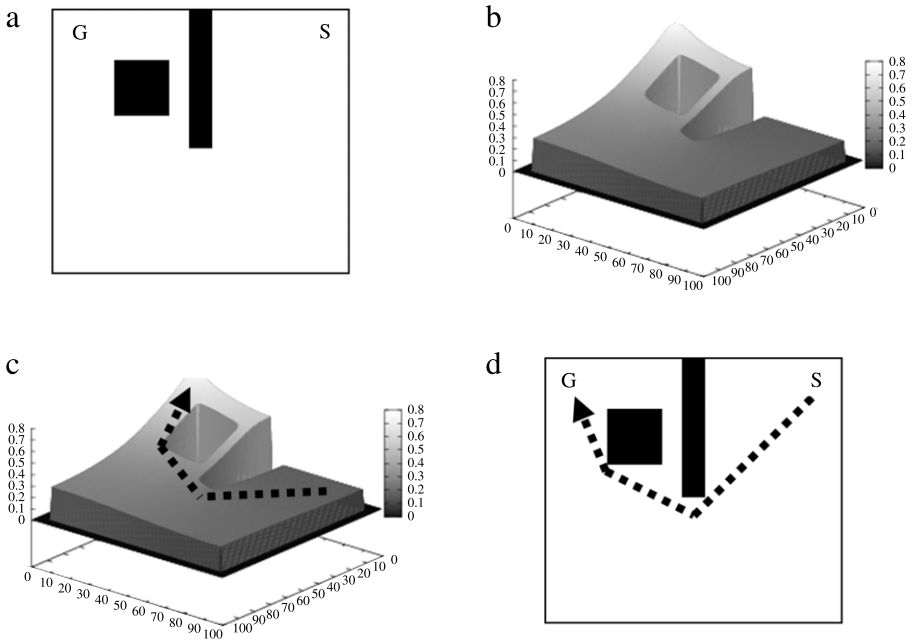

Fig. 2. General properties possessed by the class of gradient climbing models. (a) A spatial planning problem. Barriers are indicated by black blocks, while open space is white area. The goal and start positions are marked with *G* and *S*. (b) The model creates an activity gradient across neurons in a topographic map of the workspace, with a maximum at the goal. Topographic map sites corresponding to obstacles in the environment are inhibited or blocked during the gradient creation. (c) Through a model-specific mechanism, the gradient is ascended from the start position to form a trajectory (dotted black) to the goal position. (d) Same as (c), with the trajectory in black overlaid on the original map.

environment. The system uses principles of recurrent shunting competitive networks (Grossberg, 1973, 1980), filling-in of spatial attention within object or region boundaries (Fazl, Grossberg, & Mingolla, 2009), and inhibition-of-return mechanisms (Grossberg, 1978a, 1978b; Koch & Ullman, 1985; Posner & Cohen, 1984) to iteratively construct a continuous trajectory from a given initial state to a desired goal state in the presence of obstacles. Spatial attentional mechanisms in this and other models modulate space-dependent processing by exciting some locations while suppressing others. A companion article (Ivey, Bullock, & Grossberg, in preparation) shows how this continuous trajectory can be stored in working memory as a discrete series of target locations from which a series of eye or arm movements can execute the planned trajectory.

As shown in Fig. 3, the core neural CGC planner proposed here is a dynamical system that consists of neuronal elements organized into three topographically organized maps that process spatial information in a world-centered, or *allocentric*, coordinate system. Maps 1 through 3 are governed by Eqs. (1) through (14). The known goal position and start position cause excitatory inputs to corresponding positions in Maps 1 and 2. The positions covered by known obstacles send further inputs to Map 1. These inputs need not be continually sensed throughout the planning process. For example, the goal could be visually occluded, provided that its position is maintained in working memory.

Map 1 dynamics define an *anisotropic*, or directionally dependent, diffusion process that radiates activity from the goal position throughout the represented environment (see Eq. (1)). This diffusion is blocked and redirected by map positions corresponding to known or sensed obstacles in the environment. Such a boundary-gated surface filling-in process is well established in models of how the brain sees (e.g., Grossberg, 1994; Grossberg & Todorovic, 1988). An interaction between surface filling-in in prestriate visual cortex and spatial attention in parietal cortex has been predicted to enable a spotlight of spatial attention to spread until it fits itself to object form and region shape (Fazl et al., 2009). This interaction is

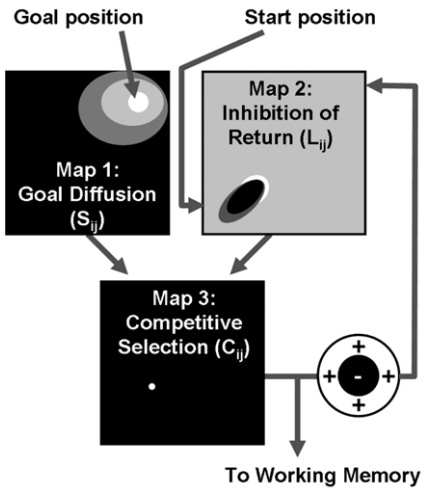

Fig. 3. Model Maps. Computations in Maps 1 through 3, each covering the same environment, are defined by Eqs. (1), (6) and (10). See the text for details.

called a *surface-shroud resonance*. Our model of anisotropic diffusion is inspired by the surface-shroud resonance concept, but simplifies it to use only boundary-gated surface diffusion whose effect is interpreted to be the corresponding spread of spatial attention in posterior parietal cortex (PPC).

Through this attentional diffusion process, a gradient of activity is created across map positions. Ascending the gradient from an initial start position that is input to Map 2 generates a feasible path (i.e., one that circumvents obstacles) to the goal position. In contrast to earlier surface filling-in models, however, the present model diffuses from a single position (that is, from the goal) rather than from many source positions along a boundary contour. This, along with the specific system dynamics and parameters, enable convergence to a gradient formed in open space rather than serving a filling-in style function. Dynamic information processing by a

Table 1
Network parameters used in all simulations.

Parameter	Value	Parameter	Value
D	1.0	T	80 000.0
E	30 000.0	U	2.0
F	60 000.0	V	0.0
K	30.0	Z	0.4
M	0.8	I_{start}	400.0
N	800.0	I_g^+	10 000.0
Q	1.6	I_g^-	7000.0
R	0.1		

recurrent competitive dynamical system in Maps 2 and 3 (see Eqs. (6) and (10)) realizes the iterative gradient climbing operation. This is accomplished by the following neural mechanisms when they interact together. First, model equations in Map 2 are defined by neural membrane, or shunting, equations, whose activities have finite upper and lower bounds of activity. Despite these fixed activity bounds, cell activities retain their sensitivity to analog differences in input strength by balancing recurrent shunting excitation and inhibition within on-center off-surround networks that self-normalize network activity. In the current model, signal functions are chosen to enable the network to make a winner-take-all choice whereby all network activity is concentrated into the cell, or cell population, that receives the largest total input at any time (Grossberg, 1973).

In addition to these classical properties, the current model introduces a novel form of transient inhibition-of-return that is achieved by an off-center on-surround projection from local regions in Map 3 to the corresponding local regions in Map 2 (see Fig. 3). This projection inhibits the current choice location while it enhances nearby locations as possible next locations in an emerging lookahead trajectory. In particular, these enhanced locations from Map 2 multiplicatively gate goal diffusion activities that are received in Map 3 as inputs from Map 1. These gated signals are then fed into the competitive selection process in Map 3, and an emerging lookahead trajectory is iteratively created from the start position to the goal position as the network automatically cycles through this feedback loop.

Once the spatiotemporal pattern that constitutes the lookahead trajectory is created, it can be sampled, segmented, and stored through time in a sequential working memory, in preparation for subsequent plan execution. Such a storage process may take place between PPC and the prefrontal cortex (PFC). The specifics of such operations, downstream from Map 3, are outside the scope of the current article, but can be made consistent with well-established prior neural models, such as the Item and Order, or Competitive Queuing, working memory models that were introduced by Grossberg (1978a, 1978b) and subsequently applied and developed to explain various kinds of data (e.g., Boardman & Bullock, 1991; Bullock & Rhodes, 2003; Houghton, 1990; Page & Norris, 1998). More recently, Item and Order models have been further developed to predict how working memory and learned planning operations may be realized by the laminar circuits of PFC (e.g., Grossberg & Pearson, 2008; Silver, Bullock, Grossberg, Histed, & Miller, 2009).

The model equations and parameters are defined next. The parameters used in all the simulations are given in Table 1. In experiments not reported on in detail here, the model generated qualitatively similar trajectories after parameter values were modified by 5%.

Map 1. Diffusion in Map 1 enables formation of an activity gradient across a 2D space. Diffusion is greatly reduced at any positions that are covered by obstacles. In open spaces, the gradient declines with distance from the goal, which serves as the diffusion source. In all three 2D Maps, subscripts i and j index the two spatial map dimensions. Different maps share the same positional indices (i, j). The boundary-gated diffusion (Cohen & Grossberg, 1984; Grossberg & Todorovic, 1988) activity S_{ij} at position (i, j) in Map 1 is governed by:

$$\frac{dS_{ij}}{dt} = -DS_{ij} + (1 - S_{ij})I_{ij} + \sum_{pq \in N_{ij}} P_{pqij}(S_{pq} - S_{ij}) - S_{ij}H_{ij}, \tag{1}$$

where D is a decay parameter (see Table 1). I_{ij} is an excitatory input at the goal position such that

$$I_{ij} = \begin{cases} I_g^+ & \text{if } (i, j) \text{ corresponds to the goal location} \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

and H_{ij} is an inhibitory input neighboring the cells of the goal location, where

$$H_{ij} = \begin{cases} I_g^- & \text{if } (i, j) \text{ is an immediate neighbor of the goal location} \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

Term H_{ij} creates a sharper gradient at the goal position and prevents saturation of the gradient near the goal. N_{ij} is the set of cells that are immediate neighbors of cell (i, j). The permeability P_{pqij} of the diffusion from neighbors (p, q) to (i, j) is given by:

$$P_{pqij} = \frac{E}{(1 + FB_{pqij})}, \tag{4}$$

where

$$B_{pqij} = \begin{cases} 1 & \text{if there is an obstacle boundary between} \\ & \text{cells } (p, q) \text{ and } (i, j) \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

E and F are scalar parameters that specify the permeability of the diffusion process. The large diffusion permeability value $E = 30\,000$ in unobstructed space is consistent with prior neural modeling work (Fazl et al., 2009). Although the permeability is large, leading to a fast diffusion, the continuously active driving input at the goal position given by (2) enables formation of a robust gradient rather than a spatially homogeneous response.

Map 2. Map 2 allows both inhibition of the position already used at the prior step in a constructed trajectory and priming of eligible (adjacent) next steps in the planning space. The activity of L_{ij} at position (i, j) in Map 2 is governed by:

$$\frac{dL_{ij}}{dt} = -KL_{ij} + (1 - L_{ij})(J_{ij} + G_{ij}^{(s)} * \Psi) - L_{ij}(G_{ij}^{(c)} * \Psi), \tag{6}$$

where K is a decay parameter and J_{ij} is an excitatory input at the goal position such that

$$J_{ij} = \begin{cases} I_{\text{start}} & \text{if } (i, j) \text{ is the start position} \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

and I_{start} is the parameterized input at the start position (see Table 1). The inhibitory center kernel $G_{ij}^{(c)}$ and the excitatory surround kernel $G_{ij}^{(s)}$ are convolved $(*)$ with Ψ in an off-center on-surround network. Ψ is the thresholded and weighted output from Map 3 (see Eqs. (10) and (12)). In particular, $G_{ij}^{(c)}$ and $G_{ij}^{(s)}$ comprise a 2D off-center, on-surround receptive field given by:

$$G_{ij}^{(c)} = \begin{cases} N & (i, j) = \text{center} \\ 0 & \text{elsewhere} \end{cases} \tag{8}$$

$$G_{ij}^{(s)} = \begin{cases} Ne^{-\frac{(x-i)^2+(y-j)^2}{2Q^2}} & \text{if } x = i-2, i-1, i, i+1, i+2 \\ & \text{and } y = j-2, j-1, j, j+1, j+2 \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

where scalar parameter N modulates receptive field strength.

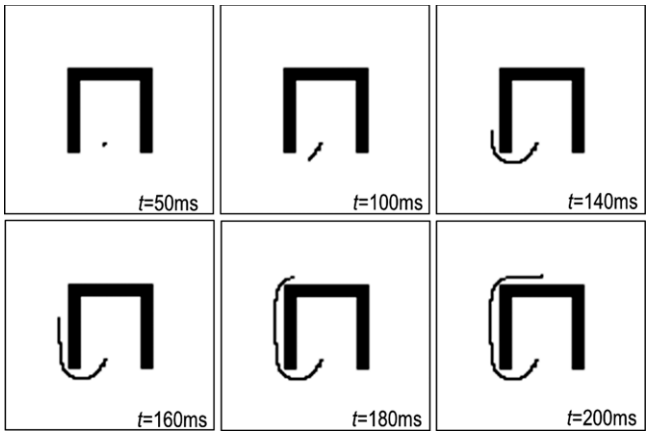

Fig. 4. Centroid positions (shown as a black curve), computed from Map 3 outputs, trace out a spatial trajectory that runs from the start position to the goal position while avoiding the obstacle (shown in black).

Map 3. The external input to each position (i, j) in Map 3 is a scaled product $S_{ij}L_{ij}$ of the activities S_{ij} and L_{ij} at corresponding positions in Map 1 (gradient) and Map 2 (eligible continuations). Map 3's recurrent self-excitation and lateral-inhibition define a competition that leads to peak activation at the site that receives the maximal input $S_{ij}L_{ij}$, which corresponds to the best-rated eligible continuation. Map 3 activity C_{ij} at position (i, j) is governed by:

$$\frac{dC_{ij}}{dt} = -RC_{ij} + (1 - C_{ij})(TS_{ij}L_{ij} + Ug(C_{ij})) - (C_{ij} + V)\left(Z\sum_{kl\neq ij} g(C_{ij})\right). \quad (10)$$

In (10), R is a decay parameter and parameter T scales the input $S_{ij}L_{ij}$. The signal function

$$g(x) = x^5 \quad (11)$$

governs recurrent on-center off-surround feedback within Map 3. This recurrent feedback ensures maximal activation at the position with the maximal current input $S_{ij}L_{ij}$. It would lead to winner-take-all (WTA) behavior (Grossberg, 1973) if the system were allowed to go to equilibrium in the absence of any changes to inputs from the other maps. However, during gradient climbing, these inputs do change, Map 3 does not go to equilibrium, and so there is an evolving distribution of activity that spans a traveling peak and nearby spatial positions. Map 3 output is governed by:

$$\Psi_{ij} = (20[C_{ij} - M]^+)^2. \quad (12)$$

Since activity must be strong enough to surpass the threshold M in (12) to produce a response, but it is also bounded by 1 by equation dynamics in (10), a scaling parameter and faster-than-linear power function is used to increase the contrast of the output signal. Map 3 activations project to Map 2, as in Eq. (6). To extract a unique position to define the traveling peak at successive times during gradient climbing, a thresholded centroiding operation is used. Let

$$P_i = \frac{1}{\sum_{ij}\Psi_{ij}}\sum_{ij} i\Psi_{ij} \quad (13)$$

and

$$P_j = \frac{1}{\sum_{ij}\Psi_{ij}}\sum_{ij} j\Psi_{ij}. \quad (14)$$

The position (P_i, P_j) defines the model output through time. The threshold M in Eq. (12) suppresses positions (i, j) with low activity and weights the output position heavily toward the most active cell positions, which trace a continuous trajectory from the start position to the goal position (see Figs. 4 and 5).

3. Planning simulation

We first illustrate system behavior using a canonical spatial lookahead planning task shown in Fig. 1(a), consisting of a concave obstacle around which the system must plan a path. The point goal position used in the simulation is indicated in the figure by G, and the simulated start position by S. The black concave object is an obstacle. Simulations were performed by numerical integration using a software implementation of the Euler method applied to model Eqs. (1)–(14) and parameters listed in Table 1, with each time step equal to 0.005 ms, as modeled by the simulation. All times reported in the figures and text are thus produced by summing time steps using the standard numerical integration procedure.

For Task 1a, Figs. 6 through 8 present snapshots of the activity through time within the model Maps 1 through 3. Activity amplitude at each map position is coded in the snapshots with white signifying maximal activity. Fig. 6 snapshots sample and depict the time course of diffusion, in which the Map 1 activity S_{ij} (Eq. (1)) flows around the concave obstacle. In the final panel of Fig. 6 (for $t = 200$ ms) the intensities have been discretized into bands as an aid to make the gradient more visible to readers. The gradient created by goal-sourced diffusion in Map 1 represents the direction that should be moved to acquire the goal.

In contrast to the other gradient climbing models (Glasius, Komoda, & Gielen, 1994, 1995, 1996; Yang & Meng, 2001), the current model defines a complete dynamical system, with no non-dynamical model components, to generate a spatial trajectory before action is taken. The spatiotemporal behavior of Maps 2 and 3 (Eqs. (6) and (10)) is depicted by successive snapshots in Figs. 7 and 8. Fig. 4 shows as a black trajectory the cumulative (integrated) weighted centroids of activity in Map 3 that are defined by (13) and (14). This trajectory is overlaid on the Task 1a scenario, in which the concave obstacle appears in black. By $t = 200$ ms, a complete trajectory from the start position to the goal position was iteratively constructed.

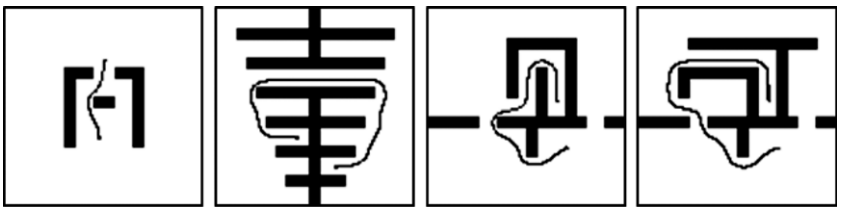

Fig. 5. Selected trajectories at the time of goal acquisition for Tasks 1b–1e (see Fig. 1) are shown in black. The tasks' obstacles are also overlaid in black. While the entire trajectory is shown at once here, the trajectories were created as a continuous sequence from the start position to the goal position. The centroid defined by Eqs. (13) and (14) varies continuously through space. The closest cell to the centroid is marked as black. Thus, small fluctuations near a cell boundary may result in multiple adjacent cells being marked.

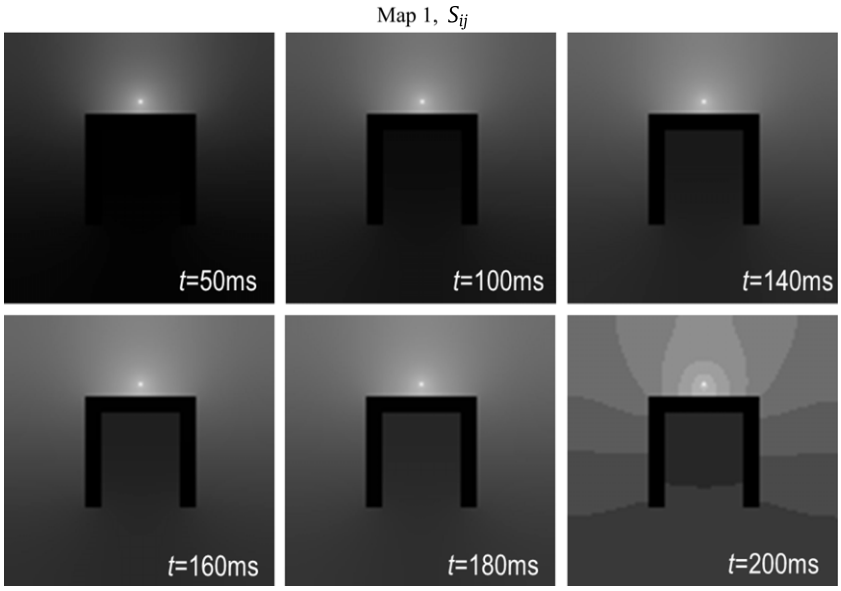

Fig. 6. Diffusion of activity S_{ij} (Eq. (1)) in Map 1 through time for the spatial planning task depicted in Fig. 1(a). The activity gradient has highest intensity at the goal position and progressively lower intensities stretching to the start position. To facilitate viewing, in the final panel, for $t = 200$ ms, the analog intensity values, output by Map 1 have been discretized into bands to create a gradient contour map.

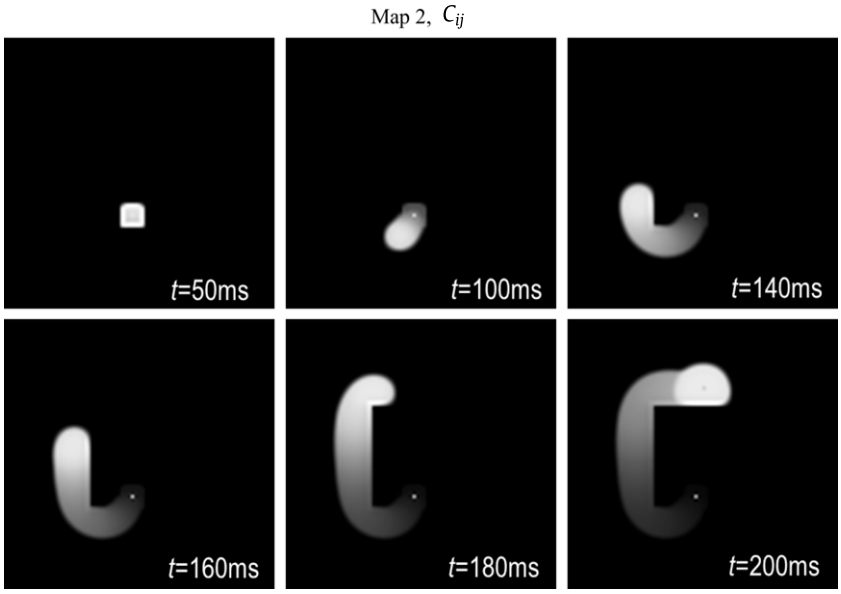

Fig. 7. Evolution through time of activations L_{ij} (Eq. (6)) across sites in Map 2, showing gradient climbing from start to goal.

The four additional tasks depicted in Fig. 1(b)–(e) were used to assess the generality and robustness of the model. A successful, approximately minimum-distance, trajectory was generated for each of these tasks on the first attempt without error (Fig. 5) using the gradient constructed by Map 1 (Fig. 9). We predict that this result generalizes to other vision-based planning tasks similar to those shown here, namely tasks that humans can quickly evaluate on relatively brief visual inspection if the key information types (goal, start, and obstacles) are correctly visualized by the viewer. More complex problems, such as navigating complicated mazes (Crowe, Averbeck, Chafee, & Georgopoulos, 2005), likely invoke further strategies, such as hierarchical search, and are beyond the scope of this model.

4. Discussion

Due to the variety and complexity of planning behaviors and deficits, as well as the need to work with primates capable of performing complex spatial tasks, neural models of flexible spatial planning of the type typically attributed to prefrontal–parietal interactions have been scarce in comparison to models of visual, auditory, and pattern recognition systems. Many descriptions of flexible spatial planning behavior are qualitative rather than quantitative. Experimental studies tend to focus on novel environmental configurations to avoid overtraining, but this prevents sufficient data collection for statistical comparisons. We have therefore focused on modeling the recurring qualitative theme of flexible successful behavior in the face of novel environmental variation without repeated performance attempts.

Reactive gradient climbing based on sensed gradients in novel environments is a common practice among animals, some of which lack complex nervous systems. The current model postulates a multi-level process that embodies a cognitive gradient climbing (CGC) strategy. Activity spreads throughout a topographic map from the goal position, with obstacles inhibiting or blocking the

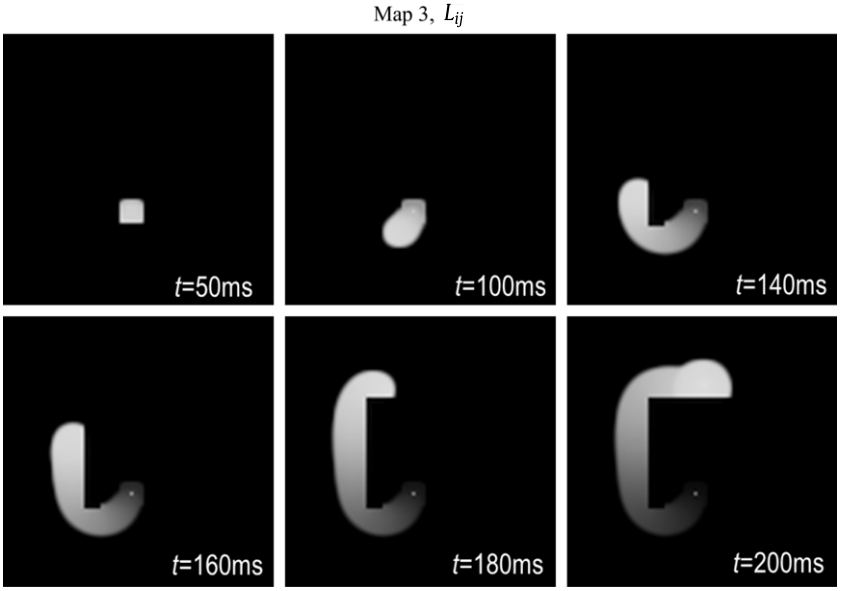

Fig. 8. Evolution through time of activations C_{ij} (Eq. (10)) across sites in Map 3, showing competitively sharpened winning positions on the path from start to goal. The weighted centroids of the thresholded evolving activation trace out a continuous spatial trajectory that comprises the plan (see Fig. 4).

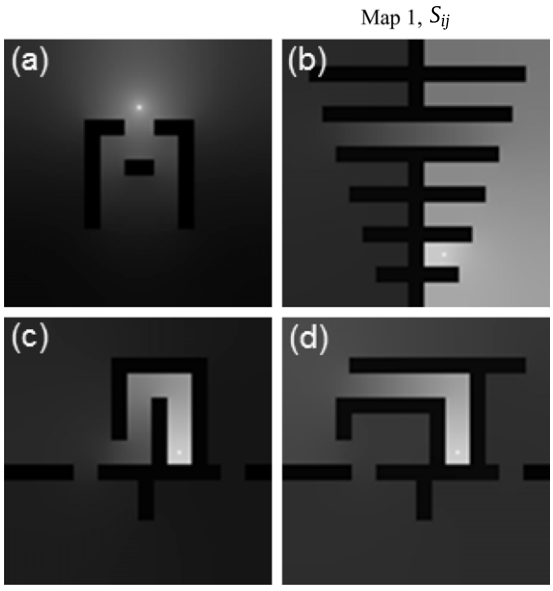

Fig. 9. Diffusion gradient in Map 1 at the time of goal acquisition for Tasks 1b–1e (see Fig. 1). (a) Task 1b at $t = 45$ ms. (b) Task 1c at $t = 800$ ms. (c) Task 1d at $t = 435$ ms. (d) Task 1e at $t = 1300$ ms.

spreading activity. To prepare a trajectory from a given start position to the goal, a local gradient is iteratively climbed from the start position. Prior models (Glasius et al., 1994, 1995, 1996; Yang & Meng, 2001) have used activation with local inhibition to implement the spreading activity. The current model instead uses a boundary-gated diffusion to create the gradient, which redirects energy rather than inhibiting it. Prior models were defined algorithmically (Yang & Meng, 2001) or did not pre-plan whole trajectories because the core iteration required external feedback of the current manipulator position during execution to choose next steps (Glasius et al., 1996). Prior models (Glasius et al., 1996; Yang & Meng, 2001) also predicted excessively long reaction times and a high effective signal-to-noise ratio or common noise-induced planning errors due to shallow gradients generated in typical cases (Fig. 11).

As noted in the Introduction, prior dynamical system models of planning have typically dealt with either relatively simple

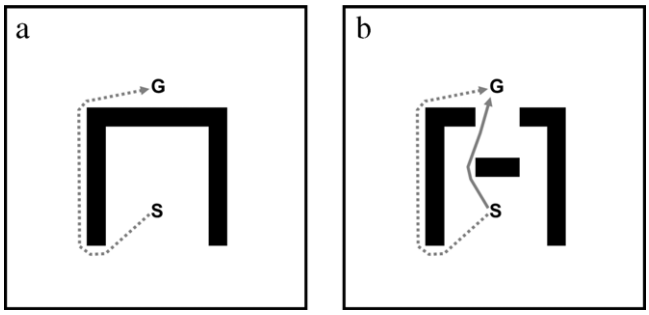

Fig. 10. Planning environments with trajectories that highlight a key problem for the class of attractor–repeller (A/R) models. (a) An A/R model's parameters might be tuned such that the repelling force of the obstacles enables escape from the concave area and successful goal acquisition (dotted trajectory). (b) The same environment with a gap created by displacing a section of the obstacle toward the start position. An A/R model that successfully escaped the concave area in the left problem would on the right problem compute a stronger repulsive force from the displaced section and would therefore take the longer route (dotted line) rather than the shorter route (solid line).

environments, by computing single direction vectors, or have relied on repeated learning over multiple attempts to discover adequate sequences. Attractor/Repeller (A/R), or potential field, models hypothesize that goals act as attractors and obstacles act as repellers during visually reactive navigation. These models have been shown to match trajectories generated by humans navigating in 3D in the presence of simple point or convex obstacles (Eichhorn, 2005; Elder et al., 2009; Fajen & Warren, 2003; Huang, Fajen, Fink, & Warren, 2006), but break down with concave obstacles or other complex obstacle configurations (Fig. 10). While many A/R models use point obstacles only, modifications have been proposed to extend A/R models to non-point obstacles (Huang et al., 2006). Our current model is used to describe lookahead planning in a two-dimensional space. Despite these different goals of the two classes of models, it is of interest to ask whether A/R models embody a competence for lookahead planning as well as for visually reactive navigation. A key issue with using A/R models for lookahead planning is revealed by the thought experiment depicted in Fig. 10, in which example trajectories are shown. An A/R model with parameters adequately tuned may successfully escape the concave obstacle shown in the map on the left. However, that same system

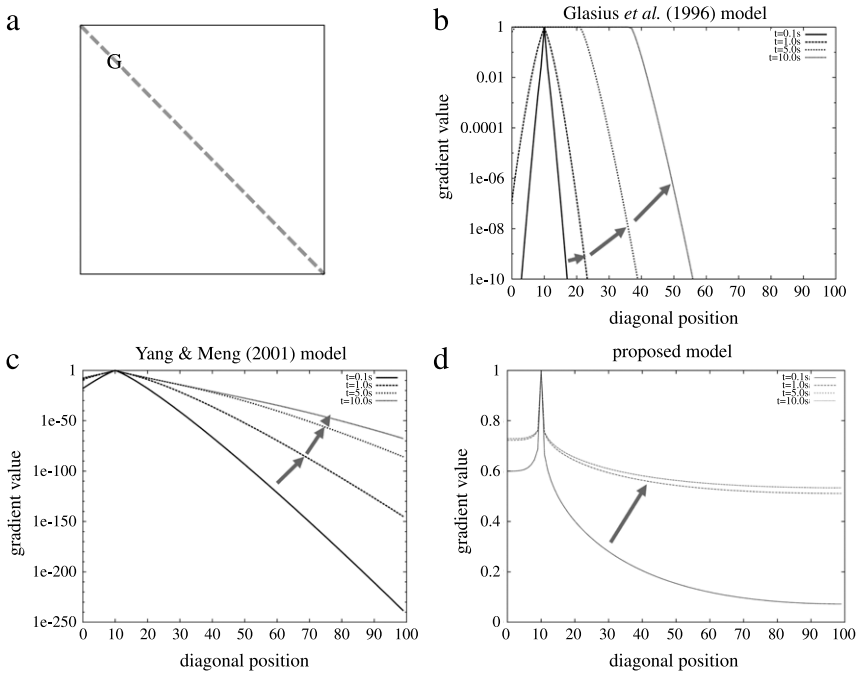

Fig. 11. Comparison of three gradient cell activation values of three gradient climbing models in free space as a function of time. (a) An obstacle-free 100 × 100 cell environment. The goal is at position (10, 10). Each model's mechanism for creating goal gradients was numerically simulated in this environment. The cell activity values for the positions along the diagonal dashed line are plotted in (b), (c), and (d) at $t = 0.1$, 1.0, 5.0 and 10.0 s. Note the logarithmic axes on (b) and (c), and the linear axis on (d). (b) Gradient values using the model by Glasius et al. (1996). By $t = 10$ s, the network has responded along less than half of the diagonal with appreciable gradient energy, but in doing so, the positions near the gradient have become saturated. A subsequent gradient climbing mechanism could not respond at positions with activity values that have not yet responded or have fully saturated at the upper bound. (c) Gradient values through time for the model by Yang and Meng (2001). The network converges more quickly than (b) and does not saturate at the upper bound, but has very small activation values. (d) In contrast to prior models, the proposed model does not saturate and quickly converges to final gradient values.

is confounded by the map on the right where a section of the obstacle is displaced toward the start position. The class of A/R models would compute a greater repulsive force and would be more likely to select a downward movement rather than the desirable shorter path.

Our model successfully selects desirable trajectories in tasks that pose problems for the A/R class of models. In particular, Tasks 1c–1e (Fig. 1) illustrate further counterexamples to A/R approaches. Our model belongs to a class of continuous-time CGC models, which use internally generated representations of gradients to account for planning in novel layouts that include concave obstacles (Glasius et al., 1994, 1995, 1996; Yang & Meng, 2001). These models are related to discrete-time Dynamic Wave Expansion models (Lebedev, Steil, & Ritter, 2005) and are extensions of resistive grid, or Laplacian, planning models (Bugmann, Taylor, & Denham, 1995; Connolly, Burns, & Weiss, 1990).

Glasius et al. (1994, 1995, 1996) presented a CGC model using spreading excitation through a 2D topographic map and inhibitory barriers. In Glasius et al. (1996) a contrast enhancement process selects the next local target. This local target position is passed to the manipulator for execution. Feedback from the manipulator on its current position causes a new local target to be selected. Because their planning model is specified using in-the-loop feedback from an external manipulator, it does not address lookahead planning *per se*, and is not consistent with data showing multiple steps of preparatory activity prior to any execution (Mushiake et al., 2006; Saito et al., 2005). It is also desirable to exclude current manipulator position from the model definition in order to eliminate difficulties in handling latencies in feedback and so that the forthcoming plan can be compared with competing alternatives or otherwise evaluated prior to execution. Finally, their model dynamics predict long reaction times, even for unobstructed trajectories (Fig. 11). Humans are much faster.

The Yang and Meng (2001) model uses a similar gradient production mechanism to that of Glasius et al. but diverges on the mechanism to climb the gradient. Obstacle objects are treated as inhibitory regions with short-range connections. When activity reaches the provided start position, a process begins to extract the spatial trajectory from the gradient algorithmically, using an iterated MAX operator to select among the local neighbors. Yang and Meng do not propose neural mechanisms that could correctly control such an iterated MAX operation. Our simulations of the Yang and Meng (2001) model revealed that the spreading activation level became very small (less than 10^{-20}) a short distance away from the goal (Fig. 11). In a noisy system, such as a biological neural network or an analog hardware implementation, noise would completely overwhelm gradient information in their model. Our model also displayed robustness to parameter changes.

The problems of the Yang and Meng (2001) model and the Glasius et al. (1996) model are specific to their formulations and do not exclude the gradient concept as a framework for modeling flexible spatial planning behavior.

This new model is realized as a novel dynamical system that generates spatial lookahead plans in the presence of complex obstacles. While neurobiological data on the specific method by which spatial plans are generated are sparse, the proposed model is consistent with prior neural modeling concepts. In particular, boundary-gated diffusive filling-in is a common component of many neural models of vision (Grossberg & Todorovic, 1988; Grossberg & Yazdanbakhsh, 2005; Kelly & Grossberg, 2000). Spreading due to a surface-shroud resonance is predicted to occur between prestriate cortical areas, such as V4 and parietal cortex (Fazl et al., 2009). Frontal projections of spreading parietal activity may also spread and/or be converted into a series of discrete movement commands that are stored in working memory. Top-down projections from working memory to PPC

could provide supporting information on goal and start positions and remembered obstacle configurations. The model proposes neurally plausible mechanisms by which topographic spatial maps may interact to control lookahead visuo-spatial planning behavior.

Acknowledgements

Daniel Bullock, Stephen Grossberg, and Richard Ivey were supported in part by CELEST, an NSF Science of Learning Center (SBE-0354378). Stephen Grossberg and Richard Ivey were supported in part by the SyNAPSE program of DARPA (HR0011-09-C-0001). Richard Ivey was additionally supported in part by the ACES program of the National Science Foundation (DGE-0221680).

References

Andersen, R., Snyder, L., Bradley, D., & Xing, J. (1997). Multimodal representation of space in the posterior parietal cortex and its use in planning movements. *Annual Review of Neuroscience*, *20*(1), 303–330.

Averbeck, B., Chafee, M., Crowe, D., & Georgopoulos, A. (2003). Neural activity in prefrontal cortex during copying geometrical shapes. I: single cells encode shape, sequence, and metric parameters. *Experimental Brain Research*, *150*(2), 127–141.

Boardman, I., & Bullock, D. (1991). A neural network model of serial order recall from short-term memory. In *Proc. international joint conference on neural networks, II* (pp. 879–884).

Boussaoud, D., & Wise, S. (1993). Primate frontal cortex: effects of stimulus and movement. *Experimental Brain Research*, *95*(1), 28–40.

Browning, A., Grossberg, S., & Mingolla, E. (2009). Cortical dynamics of navigation and steering in natural scenes: motion-based object segmentation, heading, and obstacle avoidance. *Neural Networks*, *22*(10), 1383–1398.

Bugmann, G., Taylor, J., & Denham, M. (1995). Route finding by neural nets. In J. Taylor (Ed.), *Neural networks* (pp. 217–230). Henley-on-Thames: Alfred Waller Ltd.

Bullock, D., Cisek, P., & Grossberg, S. (1998). Cortical networks for control of voluntary arm movements under variable force conditions. *Cerebral Cortex*, *8*(1), 48–62.

Bullock, D., Grossberg, S., & Guenther, F. (1993). A self-organizing neural model of motor equivalent reaching and tool use by a multijoint arm. *Journal of Cognitive Neuroscience*, *5*(4), 408–435.

Bullock, D., & Rhodes, B. (2003). Competitive queuing for serial planning and performance. In M. Arbib (Ed.), *The handbook of brain theory and neural networks* (pp. 241–248). Cambridge, MA: MIT Press.

Buttelmann, D., Carpenter, M., Call, J., & Tomasello, M. (2008). Rational tool use and tool choice in human infants and great apes. *Child Development*, *79*(3), 609–626.

Butz, M., Sigaud, O., & Gerard, P. (2003). Internal models and anticipations in adaptive learning systems. *Lecture Notes in Computer Science*, *2684*, 86–109.

Capdepuy, P., Polani, D., & Nehaniv, C. (2007). Grounding action selection in event-based anticipation. *Lecture Notes in Computer Science*, *4648*, 253–262.

Carder, H., Handley, S., & Perfect, T. (2004). Deconstructing the Tower of London: alternative moves and conflict resolution as predictors of task performance. *The Quarterly Journal of Experimental Psychology, Section A*, *57*(8), 1459–1483.

Carder, H., Handley, S., & Perfect, T. (2008). Counterintuitive and alternative moves choice in the water jug task. *Brain and Cognition*, *66*(1), 11–20.

Chafee, M., & Goldman-Rakic, P. (2000). Inactivation of parietal and prefrontal cortex reveals interdependence of neural activity during memory-guided saccades. *Journal of Neurophysiology*, *83*(3), 1550–1566.

Ciaramelli, E. (2007). The role of ventromedial prefrontal cortex in navigation: a case of impaired way finding and rehabilitation. *Neuropsychologia*, *46*(7), 2099–2105.

Cohen, S., & Freedman, M. (2005). Cognitive and behavioral changes in the Parkinson-plus syndromes. In W. Weiner, A. Lang, & K. Anderson (Eds.), *Behavioral neurology of movement disorders* (pp. 166–186). Hagerstown, MD: Lippincott Williams & Wilkins.

Cohen, M., & Grossberg, S. (1984). Neural dynamics of brightness perception: features, boundaries, diffusion, and resonance. *Perception and Psychophysics*, *36*, 428–456.

Colvin, M., Dunbar, K., & Grafman, J. (2001). The effects of frontal lobe lesions on goal achievement in the water jug task. *Journal of Cognitive Neuroscience*, *13*(8), 1129–1147.

Committeri, G., Pitzalis, S., Galati, G., Patria, F., Pelle, G., Sabatini, U., et al. (2007). Neural bases of personal and extrapersonal neglect in humans. *Brain*, *130*(2), 431–441.

Connolly, C., Burns, J., & Weiss, R. (1990). Path planning using Laplace's equation. In *Proceedings of the IEEE international conference on robotics and automation* (pp. 2102–2106). Cincinnati: IEEE.

Crowe, D., Averbeck, B., Chafee, M., & Georgopoulos, A. (2005). Dynamics of parietal neural activity during spatial cognitive processing. *Neuron*, *47*(6), 885–891.

Diamond, A. (1985). Development of the ability to use recall to guide action, as indicated by infants' performance on A-not-B. *Child Development*, *56*, 868–883.

Dias, R., Robbins, T., & Roberts, A. (1996). Primate analogue of the Wisconsin card sorting test: effects of excitoxic lesions of the prefrontal cortex in the marmoset. *Behavioral Neuroscience*, *110*(5), 872–886.

Dijkstra, E. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, *1*(1), 269–271.

Eichhorn, M. (2005). A reactive obstacle avoidance system for an autonomous underwater vehicle. In *IFAC world congress* (pp. 3–8). Prague: IFAC.

Elder, D., Grossberg, S., & Mingolla, E. (2009). A neural model of visually guided steering, obstacle avoidance, and route selection. *Journal of Experimental Psychology: Human Perception and Performance*, *35*, 1501–1531.

Ersche, K., Clark, L., London, M., Robbins, T., & Sahakian, B. (2006). Clinical research profile of executive and memory function associated with amphetamine and opiate dependence. *Neuropsychopharmacology*, *31*, 1036–1047.

Fajen, B., & Warren, W. (2003). Behavioral dynamics of steering, obstacle avoidance, and route selection. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(2), 343–362.

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), 1–48.

Fu, W., & Anderson, J. (2006). From recurrent choice to skill learning: a reinforcement-learning model. *Journal of Experimental Psychology: General*, *135*(2), 184–206.

Funahashi, S. (2001). Neuronal mechanisms of executive control by the prefrontal cortex. *Neuroscience Research*, *39*(2), 147–165.

Funahashi, S., Chafee, M., & Goldman-Rakic, P. (1993). Prefrontal neuronal activity in rhesus monkeys performing a delayed anti-saccade task. *Nature*, *365*(6448), 753–756.

Glasius, R., Komoda, A., & Gielen, S. (1994). Population coding in a neural net for trajectory formation. *Network: Computation in Neural Systems*, *5*(4), 549–563.

Glasius, R., Komoda, A., & Gielen, S. (1995). Neural network dynamics for path planning and obstacle avoidance. *Neural Networks*, *8*(1), 125–133.

Glasius, R., Komoda, A., & Gielen, S. (1996). A biologically inspired neural net for trajectory formation and obstacle avoidance. *Biological Cybernetics*, *74*(6), 511–520.

Grossberg, S. (1973). Contour enhancement, short term memory, and constancies in reverberating neural networks. *Studies in Applied Mathematics*, *52*(3), 213–257.

Grossberg, S. (1978a). A theory of human memory: self-organization and performance of sensory-motor codes, maps, and plans. *Progress in Theoretical Biology*, *5*, 233–374.

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). Biological competition: decision rules, pattern formation, and oscillations. *Proceedings of the National Academy of Sciences*, *77*(4), 2338–2342.

Grossberg, S. (1994). 3-D vision and figure-ground separation by visual cortex. *Perception and Psychophysics*, *55*, 48–120.

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*(3), 677–732.

Grossberg, S., & Todorovic, D. (1988). Neural dynamics of 1-D and 2-D brightness perception: a unified model of classical and recent phenomena. *Perception and Psychophysics*, *43*, 241–277.

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

He, B., Snyder, A., Vincent, J., Epstein, A., Shulman, G., & Corbetta, M. (2007). Breakdown of functional connectivity in frontoparietal networks underlies behavioral deficits in spatial neglect. *Neuron*, *53*(6), 905–918.

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.

Huang, W., Fajen, B., Fink, J., & Warren, W. (2006). Visual navigation and obstacle avoidance using a steering potential function. *Robotics and Autonomous Systems*, *54*(4), 288–299.

Ivey, R., Bullock, D., & Grossberg, S. (2008). A neuromorphic model of spatial lookahead planning. In *International conference on cognitive and neural systems. Session 2, poster 2* (pp. 1–23).

Ivey, R., Bullock, D., & Grossberg, S. (2010). Lookahead planning of sequential actions: from boundary-gated flow of spatial attention to working memory and action (in preparation).

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

Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: towards the underlying neural circuitry. *Human Neurobiology*, *4*(4), 219–227.

Koenig, S., & Likhachev, M. (2002). *D** lite. In *Proceedings of the eighteenth national conference on artificial intelligence* (pp. 476–483).

Lashley, K. (1951). The problem of serial order in behavior. In L. Jeffress (Ed.), *Cerebral mechanisms in behavior* (pp. 112–146). New York: Wiley.

Lebedev, D., Steil, J., & Ritter, H. (2005). The dynamic wave expansion neural network model for robot motion planning in time-varying environments. *Neural Networks*, *18*(3), 267–285.

Levy, R., & Goldman-Rakic, P. (1999). Association of storage and processing functions in the dorsolateral prefrontal cortex of the nonhuman primate. *Journal of Neuroscience*, *19*(12), 5149–5158.

Macnab, R., & Koshland, D. (1972). The gradient-sensing mechanism in bacterial chemotaxis. *Proceedings of the National Academy of Sciences*, *69*(9), 2509–2512.
Matthews, A. (1996). Development of preterm and full-term infant ability on A-not-B, recall memory, transparent barrier detour, and means-end tasks. *Child Development*, *67*(6), 2658–2676.
Miller, E., & Cohen, J. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*(1), 167–202.
Miller, E., Erickson, C., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience*, *16*(16), 5154–5167.
Mushiake, H., Saito, N., Sakamoto, K., Itoyama, Y., & Tanji, J. (2006). Activity in the lateral prefrontal cortex reflects multiple steps of future events in action plans. *Neuron*, *50*(4), 631–641.
Mushiake, H., Saito, N., Sakamoto, K., Sato, Y., & Tanji, J. (2001). Visually based path-planning by Japanese monkeys. *Cognitive Brain Research*, *11*(1), 165–169.
Page, M., & Norris, D. (1998). The primacy model: a new model of immediate serial recall. *Psychological Review*, *105*(4), 761–781.
Posner, M. I., & Cohen, Y. (1984). Components of visual orienting. In H. Bouma, & T. U. Bouwhuis (Eds.), *Attention and performance X* (pp. 531–556). Hillsdale, NJ: Erlbaum.
Roitblat, H. (1994). Mechanism and process in animal behavior: models of animals, animals as models. In D. Cliff, P. Husband, J. Meyer, & S. Wilson (Eds.), *From animals to animats 3: proceedings of the third international conference on simulation of adaptive behavior* (pp. 12–21). Cambridge, MA: MIT Press.
Sahakian, B., Elliott, R., Low, N., Mehta, M., Clark, R., & Pozniak, A. (1995). Neuropsychological deficits in tests of executive function in asymptomatic and symptomatic HIV-1 seropositive men. *Psychological Medicine*, *25*(6), 1233–1246.
Saito, N., Mushiake, H., Sakamoto, K., Itoyama, Y., & Tanji, J. (2005). Representation of immediate and final behavioral goals in the monkey prefrontal cortex during an instructed delay period. *Cerebral Cortex*, *15*(10), 1535–1546.
Silver, M., Bullock, D., Grossberg, S., Histed, M., & Miller, E. (2009). Interactions between rank-selective spatial working memory and the supplementary eye fields during the dynamic control of eye movements. *Program no. 575.11. 2009 neuroscience meeting planner*. Chicago, IL: Society for Neuroscience. Online.
Smith, E., Jonides, J., Koeppe, R., Awh, E., Schumacher, E., & Minoshima, S. (1995). Spatial versus object working memory: PET investigations. *Journal of Cognitive Neuroscience*, *7*(3), 37–356.
Stentz, A. (1995). The focussed D^* algorithm for real-time replanning. In *Proceedings of the international joint conference on artificial intelligence* (pp. 476–483).
Sutton, R., & Barto, A. (1998). *Reinforcement learning: an introduction*. Cambridge, MA: MIT Press.
Tanji, J., & Hoshi, E. (2008). Role of the lateral prefrontal cortex in executive behavioral control. *Physiological Review*, *88*(1), 37–57.
Tolman, E. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*(4), 189–208.
Tolman, E. (1959). *Principles of purposive behavior*: *Vol. 2*. New York: McGraw-Hill.
Walker, S., Mikheenko, Y., Argyle, L., Robbins, T., & Roberts, A. (2006). Selective prefrontal serotonin depletion impairs acquisition of a detour-reaching task. *European Journal of Neuroscience*, *23*(11), 3119–3123.
Ward, G., & Allport, A. (1997). Planning and problem solving using the five disc Tower of London task. *The Quarterly Journal of Experimental Psychology, Section A*, *50*(1), 49–78.
Yang, S., & Meng, M. (2001). Neural network approaches to dynamic collision-free trajectory generation. *IEEE Transactions on Systems, Man and Cybernetics, Part B*, *31*(3), 302–318.