RESEARCH ARTICLE

Scale-Free Navigational Planning by Neuronal Traveling Waves

Azadeh Khajeh-Alijani, Robert Urbanczik, Walter Senn*

Department of Physiology, University of Bern, Bern, Switzerland

* senn@pyl.unibe.ch

Abstract

Spatial navigation and planning is assumed to involve a cognitive map for evaluating trajectories towards a goal. How such a map is realized in neuronal terms, however, remains elusive. Here we describe a simple and noise-robust neuronal implementation of a path finding algorithm in complex environments. We consider a neuronal map of the environment that supports a traveling wave spreading out from the goal location opposite to direction of the physical movement. At each position of the map, the smallest firing phase between adjacent neurons indicate the shortest direction towards the goal. In contrast to diffusion or single-wave-fronts, local phase differences build up in time at arbitrary distances from the goal, providing a minimal and robust directional information throughout the map. The time needed to reach the steady state represents an estimate of an agent's waiting time before it heads off to the goal. Given typical waiting times we estimate the minimal number of neurons involved in the cognitive map. In the context of the planning model, forward and backward spread of neuronal activity, oscillatory waves, and phase precession get a functional interpretation, allowing for speculations about the biological counterpart.

Introduction

Planning is a hallmark of higher cognitive functions. It has been particularly well studied as navigational planning that involves hippocampal-prefrontal cortex structures, and the neuronal processing involved in this case was suggested to be paradigmatic for planning in general [1]. Medial temporal lobe has been proposed to form a cognitive map [2, 3]. Yet, how navigational planning is possible in a noisy neuronal substrate remains an open question.

From a computational perspective, navigational planning amounts to finding the shortest route between two points. This can be formalized in terms of a path search problem in a graph specified by nodes and connections. There is a set of optimal algorithms solving this problem that go back to the classical breath-first search algorithm by Dijkstra [4]. In its backward version, this algorithm determines the distances from a target node backwards to successive neighbors throughout the graph until the start node is reached, and from there works stepwise forward to the target node again [5]. A bidirectional version of breath-first graph search

algorithm is also implemented by simultaneously triggered waves of activity at both the target and the start node that propagates through multiple networks by diffusive coupling [6].

In some form such a 2-step backspread–forwardtrack procedure is present in all of today's graph search algorithms, and it is difficult to imagine solutions of the planning problem which do not involve this core idea. Accordingly, various neuronal planning models have considered the backpropagation of activity from the goal across a topological map of the environment towards the start position [7–11]. However, these models suffer from an exponential decay of activity with distance from the goal. In technical solutions, evaluating small signals just requires high numerical range and precision. But in biological systems the large neuronal fluctuations prevent a reliable implementation across multiple spatial scales (Fig 1a). A recent proposal considers the spread of a single front of action potentials across a topographic map [12]. The direction from which the front reaches the start position first, indicates the shortest path. Yet, if independent noise is added by each neuronal processing step, information is again lost quickly. Here we suggest a phase-coding scheme that allows an agent to plan within a single network across many spatial scales, without requiring a hierarchical coding [13].

Results

We first outline the basic idea in the 1-dimensional case. The various models that assume a spread of directional information from the goal to putative starting points via diffusion of activity [7–11] inherently set a spatial scale for the planning that is limited by the length constant of the exponential activity decay. In fact, in the presence of a certain noise level and an upper bound of the activity at the goal position, the directional information will vanish after a few multiple of this length constant (Fig 1a).

To overcome the problem of exponential information decay (Fig 1a), we considered an encoding of the directional information in the local phase differences of a periodic traveling wave spreading across the map. All map neurons receive a synaptic drive that fires them periodically, with the goal-representing neuron being driven stronger and thus firing initially with a slightly higher frequency. Due to the recurrent nearest neighbor connectivity the surrounding neurons one after the other will adapt to this faster frequency and, eventually, the map neurons all fire with the same fast period, but with phase shifts that increase with distance from the goal. Hence, a local comparison of the phases allows for detecting the shortest path to the goal. The direction is determined by the neighboring neuron that fires earliest within a cycle, leading to a movement along decreasing phases (Fig 1b, top). Even at large goal distances, the same non-zero local phase difference between neighboring neurons is attained, although more time is needed for converging to the steady state the further away a neuron is from the goal (for a mathematical treatment see S1 Text).

Encoding goal, positions and obstacles

We next consider a 2-dimensional well-explored environment with possible obstacles. The planning layer is composed of spiking neurons each coding for a position in the environment. Excitatory synaptic connections between neurons exist if the places the neurons represent are adjacent to each other (Fig 2a). The readout layer associates 4 cardinal motion directions to each place in the planning layer. It reads out the local phase differences in the planning layer and translates the readings into a sequence of actions towards the goal.

For the planning layer, we consider an array of $N \times N$ synaptically coupled, single compartment Hodgkin-Huxley type neurons. The dynamics of the membrane potential V_{ij} of the

Fig 1. A comparison of the classical diffusion model and our scale-free traveling wave model in 1 dimension. For illustration, a linear environment with 20 discrete neurons and goal at position 1 is considered. In the diffusion model (a), activity (membrane voltage) is spread from the goal across the environment (top panel) with exponentially decaying gradient, and hence quickly fading directional information (bottom). In the traveling wave model (b), activation of the goal synaptically spreads through the environment such that the firing phase of adjacent neurons increases linearly with distance from the goal (top), resulting in a fixed and positive local phase difference and hence in directional information that does not decay in space and is not restricted to a specific scale ('scale-free'), bottom.

doi:10.1371/journal.pone.0127269.g001

planning neuron encoding position (i, j) is given by

$$C\frac{\mathrm{d}V_{ij}}{\mathrm{d}t} = -I_L^{ij} - I_{ion}^{ij} + \epsilon I_{syn}^{ij}(t) + I_{ext}^{ij}(t), \quad i, j = 1, 2, \cdots N. \tag{1}$$

Here, I_L^{ij} is a leak current, I_{ion}^{ij} defines the various intrinsic membrane currents responsible for the action potential generation, I_{syn}^{ij} is the synaptic current from the 4 neighboring neurons defining whether or not the place (i, j) can be reached from that neighbor (for details see Materials and Methods).

The external current $I_{ext}^{ij}(t)$ is produced by a population of N_E external neurons that stochastically fire with a certain Poisson rate v and have an excitatory connection strength J. By choosing these parameters, we can independently tune the mean μ_{ext} and the standard deviation σ_{ext} of the external current (see also Materials and Methods). This allows us to control the level of noise in the simulations. The goal neuron is driven by a slightly stronger external current than the other planning neurons ($\mu_{ext}^{goal} > \mu_{ext}$). In the absence of a recurrent synaptic coupling ($\epsilon = 0$), each neuron ij fires periodically with some jitter in the firing times. When the goal neuron has larger intrinsic oscillation frequency ($\mu_{ext}^{goal} > \mu_{ext}$), weak coupling ($\epsilon > 0$) shifts the firing phases of its neighbors and this progressively influences the firing phases of other neurons. This frequency difference propagates through the network and leads to a periodic traveling

Fig 2. Network architecture, traveling waves and planning time. (a) Planning and readout network. For each neuron in the planning layer, 4 actions can be assigned. Actions neurons associated to planning neuron (i, j) are W_{ij}, E_{ij}, N_{ij} and S_{ij} which receive synaptic input respectively from the left, right, north, and south neighbor of the neuron (i, j) and evoke a motion in the same directions (just one synaptic input is shown). Action neurons corresponding to the current place of the agent (here again (i, j)) are driven by an additional input ($I_{ext}^{a_{ij}}$). The first of the 4 action neurons that is fired by the passing traveling wave inhibits the other 3 action neurons. (b) Synaptically propagating waves of activity from the goal neuron at (1,1) across a planning layer of 20×20 neurons, for four different obstacle configurations. Colors code for firing phases at steady state relative to the goal neuron. (c) Time courses of the local phase difference for two sample neurons at positions (10,10) and (20,20), indicated by □ and ○ in the top left panel of b, with their local west-positioned neighbors. The time to reach the maximal local phase difference represents the planning time for these two start positions towards the goal (here 600 and 1050 ms, vertical lines), and subsequently the full path towards the goal can be read out.

wave [14–16]. For a neuron representing a transient obstacle, the external drive inhibits it such that it does not fire, even when receiving the weak lateral input.

Firing phase encodes distance to the goal

When turning on the external input to the goal and the other neurons, the planning network organizes itself into a periodic traveling wave spreading from the goal neuron through the network. After a transient period, the firing phases become ordered according to their spatial distance to the goal, i.e. spatially closer neurons to the goal fire earlier than those far away (Fig 2b). This situation repeats itself after a fixed time, the common period of the population. This periodic traveling wave is a stable phase-locked state in which all the neurons periodically fire with a common period but with a phase difference in the firing times [17, 18].

Fig 2b shows examples of stationary phases of a periodic traveling wave spreading from the goal neuron at position (1,1) across the planning layer with and without obstacles. The intrinsic frequency of the goal neuron is 18 Hz and that of the other neurons 17 Hz. The development of the local phase differences for two distal neurons, after injecting the small additional current into the goal neuron, shows a delayed propagation of the directional information with a swift increase and convergence to the steady state after 600 and, respectively, 1050 ms (Fig 2c). The steady state of the traveling wave is said to be reached when the increase of the local phase difference at two successive times is less than 10% of the current value. The time from turning on the external neurons until the steady state at the start position is referred to as planning time. Thereafter, the spatial ordering of the firings is established between start and goal, so that firing phases can be read out across the entire path without requiring additional planning. Remarkably, for readout it is necessary that a non-zero local phase difference is reached independently

of the distance from the goal (see Materials and Methods). This feature assures a noise-robust readout of the directional information at any point in the map without information loss with increasing distance from the goal.

Reading firing phases and local phase differences

Associated to each position, four action neurons representing the cardinal directions (W, E, N, S) receive inputs from the planning layer and an additional input representing the agent current position. In the example of Fig 2a where the agent is assumed to be at position (i, j), the four action neurons receive the common subthreshold input $I_{ext}^{a_{ij}}$.

When the traveling wave spreading from the goal through the planning layer arrives at the closest neighbor from (i, j) towards the goal (the left neuron at $(i, j − 1)$), the synaptic input from the planning layer to the readout layer fires action neuron W_{ij} first (Fig 2a). The other three action neurons (E_{ij}, N_{ij} and S_{ij}) at that position are inhibited by the spiking of the first action neuron. The agent will move towards the new position $(i, j − 1)$ where the next action can be read out. Note that once the periodic traveling wave reached the steady state at the animal's position, it did so for all positions towards the goal, and no planning time is required anymore before reading out the action. Hence, at the new place, the next action can be read out from the next period of the wave.

Background noise is overcome by longer readout times

To check for the noise robustness of our architecture we modulated the noisiness of the external input current. This was achieved by changing the number of external neurons N_E firing with a constant Poisson rate, and the synaptic strength J with which they drive the planning neurons. For instance, a mean $\mu_{ext} = 12$ mV/ms and standard deviation $\sigma_{ext} = 0.7$ mV/ms of I_{ext}^{ij} is obtained by a total afferent Poisson rate of 16′000 Hz, a synaptic time constant of 2 ms, and a synaptic strength of $J = 0.375$ mV/ms (cf. the voltage trace in Fig 3b, top panel, and Materials and Methods). For such realistic noise the network still displays close to periodic traveling waves and a shortest path can reliably be found (Fig 3a). The external current was driving the goal neurons with a periodicity of roughly 18 Hz and the other planning neurons with 17 Hz. For large values of σ_{ext}, stochastic traveling waves are generated in which the firing times of neurons may become disordered [19] and the readout mechanism fails to find a shortest path, although the goal itself is still found (Fig 3b). In a real brain, the neurons may be subject to a time-dependent common modulation that makes the firing irregular, although it remains correlated (Fig 3c). Furthermore, there may also be a stochastic bias in the external current that drive the individual neurons. This may transiently revert the order of firing among neighboring neurons, but due to the recurrent connectivity, differences in the drive can be corrected and a short path to the goal can still be found (Fig 3d). The bias was produced by randomly varying the input firing rates such that the input currents I_{ext}^{ij} had a temporal mean μ_{ext}^{ij} that itself was varied.

To quantify the degradation with noise we introduced a measure of the planning performance (PP) by calculating the ratio between the shortest path connecting start and goal, and the average path lengths that have been chosen, $PP = \frac{\text{shortest path length}}{\langle \text{chosen path lengths} \rangle}$. The planning performance—here evaluated for start-goal distance of 10 steps—only slowly degrades with the noise level (Fig 4a). The performance in the presence of noise can be improved by extending the readout time for each action selection, i.e. the time for accumulating evidence about the direction to take. The longer this time, the more cycles can be evaluated to check which of the four action neurons associated to the current position typically fires first. With an oscillation
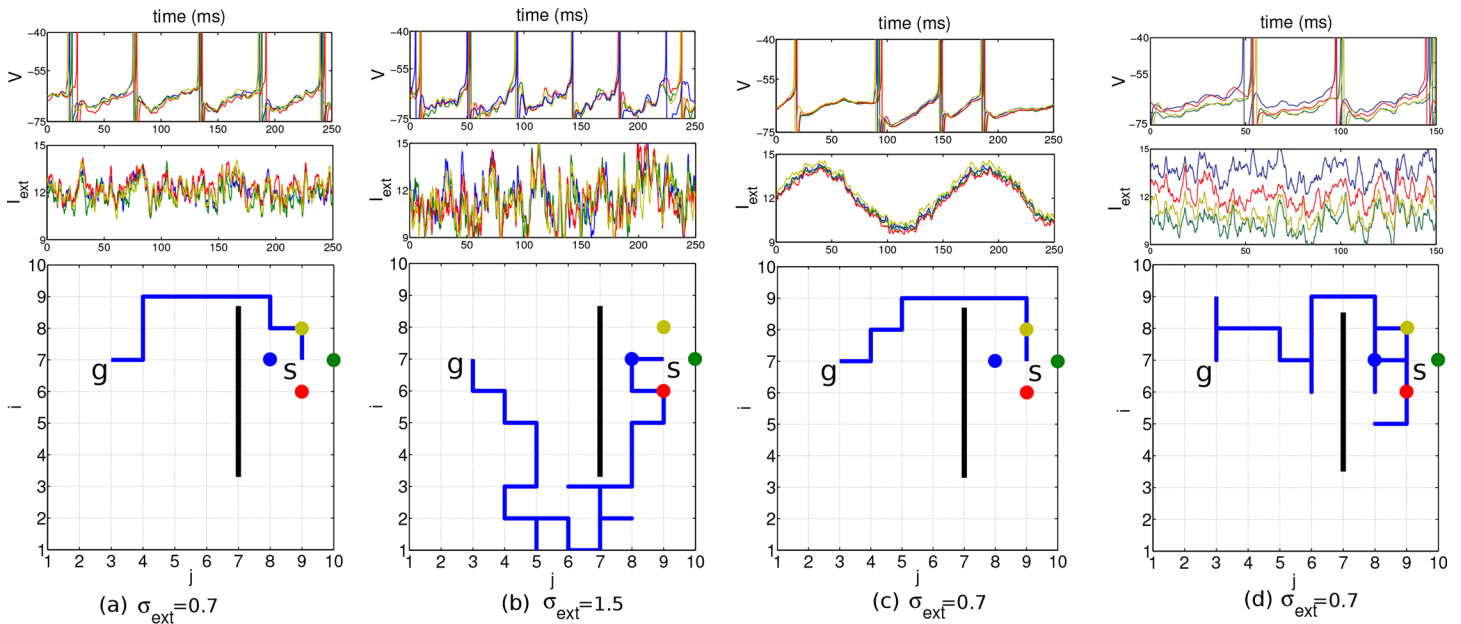
Fig 3. Results in a network of 10 × 10 planning neurons with different noise scenarios. Start position at (7,9), goal position at (7,3) and obstacle indicated by the black bar. Top row: voltage of the four nearest neighbor neurons at the start position (identity color coded). Middle row: external input current these neurons receive. Bottom row: 2D environment with the chosen path (blue line) from the start to the goal. Standard deviations σ_{ext} of the input currents indicated below. a, b: The mean of the input current generated by the background Poisson spiking neuron was constant and identical for all neurons (μ_{ext} = 12 mV/ms), except for the goal neuron (marked with g, $\mu_{ext}^{goal} = 12.5$). c: A common sinusoidal fluctuation in the Poisson firing rate of the background neurons does not disturb the relative timing among neighboring neurons. d: A randomly chosen bias in the mean μ_{ext}^{ij} of the individual input currents with standard deviation 3 does not prevent the agent from finding a short path to the goal. In all simulations, the planning times were 600 ms, the readout times 250 ms, and the coupling strength was ϵ = 0.15.

doi:10.1371/journal.pone.0127269.g003

Fig 4. Effect of noise on planning performance and readout time in the network used in Fig 3. (a) Planning performance, shown for 3 different readout times of 60, 180 and 240 ms (corresponding to 1, 3 and 4 readout cycles, bottom to top), declines with increasing noise (average across 10 chosen paths, error bars represent standard deviations of mean). (b) Readout time used at each position such that a shortest 10-step path is found, evaluated for the different noise levels. Parameters, network- and task configuration as used in Fig 3.

doi:10.1371/journal.pone.0127269.g004

frequency of 17 Hz and a readout time of 240 ms, for instance, one obtains 4 readout cycles with an inter-cycle interval of roughly 60 ms, and a considerable improvement in performance as compared to evaluating only a single cycle (Fig 4a). Given a noise level σ_{ext} we determined the shortest readout time such that the goal is always found on a shortest path in 10 random start-goal configurations of distance 10. The required readout time increases roughly linearly with the noise level (Fig 4b).

Planning performance increases with frequency

A more subtle way to assure a high planning performance, beside increasing the readout time, is to increase the oscillation frequency of the planning neurons by injecting stronger external input currents. Fig 5a shows a comparison of the performance versus noise curve for a slow and a fast periodic traveling wave, and for a solitary wave. The performance is best for the fast traveling wave at a given noise level. Two reasons contribute to this effect. First, a high oscillation frequency requires stronger external input currents, and the neuron is shifted from the noise-driven regime into the drift-dominated regime where spike timing becomes more precise [20]. Second, the steady state with its phase-locking pattern represent an attractor of the phase dynamics and this attractor becomes more stable with higher frequency and hence cleans up the noise [21].

When the frequency goes to zero, the periodic wave degenerates to a single solitary wave which propagates once through the network. For this solitary wave, both benefits of a fast periodic wave disappear. First, because the spiking must be triggered by a few neighboring neurons, the planning neurons need to be in a subthreshold but depolarized regime, where they are also sensitive to noise. Second, the timing pattern of a single wave cannot profit from the phase attractor property of a periodic wave. Correspondingly, the directional information in the
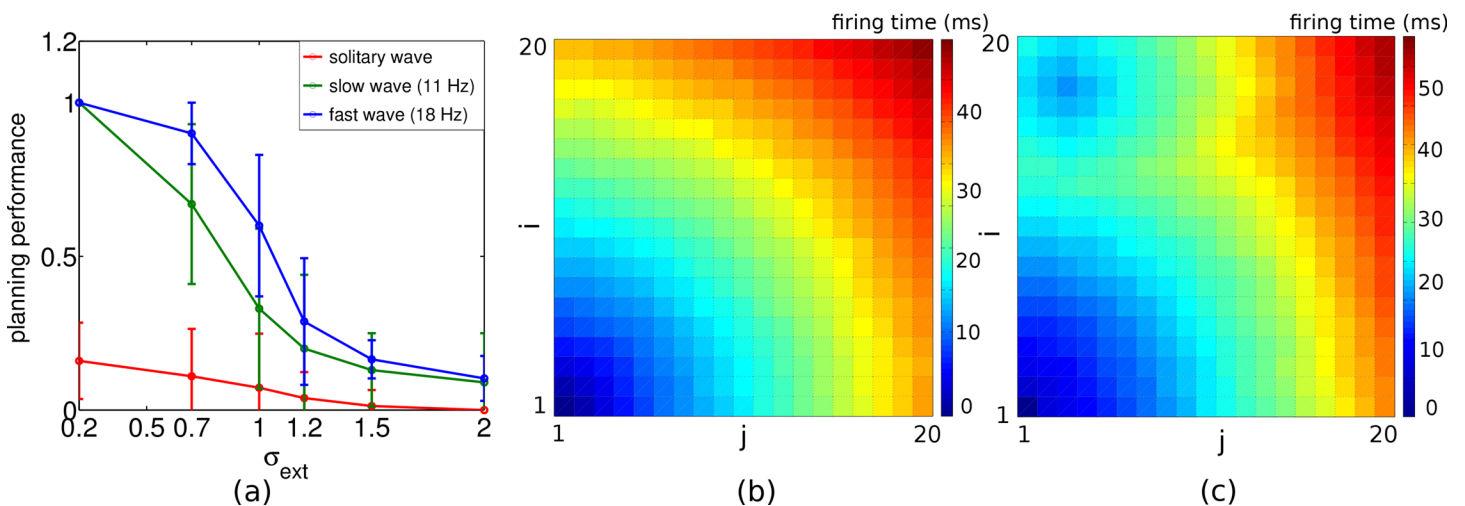

Fig 5. Planning performance decreases with decreasing oscillation frequency and is worst for a solitary wave. (a) Planning performance as a function of the noise level for a 18 Hz (blue) and a 11 Hz (green) intrinsic oscillation frequency with 4 readout cycles to select a single action. For the solitary wave, planning performance was measured after averaging the path lengths across 4 sweeps. Error bars from 10 realizations. (b, c) Color coded spike times relative to the spike time of the goal position (1,1) for (b) the 18 Hz periodic traveling wave after reaching steady state and (c) the solitary wave, both with noise level $\sigma_{ext} = 0.2$. For the solitary wave, the spike times do not faithfully represent distance from goal, and hence the action selection mechanism may yield a path to a non-goal position (the blue island in (c)).

doi:10.1371/journal.pone.0127269.g005

spike timing order for a single solitary wave quickly degrades with increasing noise level (lowest curve in Fig 5a).

To further compare the directional information in a periodic traveling wave and a solitary wave we were considering the spike timing map relative to the goal neuron for these two cases. As expected, the spike timings for the solitary wave does not fully reflect the distances to the goal, as e.g. read off from the isolated blue island in panel c.

Planning time increases with distance and frequency

The characteristic feature of encoding the directional information in local phase differences is that, in the steady state, this information does not decay throughout the whole network (Figs 1b and 2c). However, the time to reach a steady state in the local phase differences (i.e. the planning time) increases with the distance from the goal, measured along the shortest path. While, at positions close to the goal, local phase differences are already at steady state, further away, neighboring neurons are still firing synchronously (Fig 6a). Fig 6b shows the time needed for the full phase difference to spread across a 2-dimensional network without obstacles. In the case of obstacles, the planning time increases linearly with the length of the shortest path to the goal.

Importantly, planning time also increases with the oscillation frequency (Fig 6c). This is because for increasing frequencies the excitatory postsynaptic potentials start to fall into the refractory period of the previous postsynaptic spike, and they are therefore less efficient in advancing the phase of the postsynaptic neuron [22, 23]. Ultimately, disproportionally more cycles are needed to reach steady state. One might think of increasing the synaptic strength at higher frequencies to speed up the convergence. Then the steady state is indeed reached earlier, but the stronger coupling (ϵ) of the neurons reduces the final phase differences [17]. Hence, in the presence of noise, planning performance again decreases (data not shown). Both, the coupling strength and the oscillation frequency, can be chosen to optimize the trade-off between better planning performance (Fig 5a) and longer planning time (Fig 6c). As we have optimized this trade-off, the relation between planning time and network size (measured along the diagonal, Fig 6b) yields a prediction of network size involved in behaviorally estimated planning times.


Fig 6. Determinants of the planning time. (a) Snapshots of the local phase differences along diagonal positions from the goal. The intrinsic oscillation frequency was 17 Hz for the non-goal positions and 18 Hz for the goal position. Noise level was σ_{ext} = 0. (b) Planning time (i.e. the time to reach roughly 90% of the final local phase difference) increases linearly with the distance from the goal. (c) Planning time also increases with the frequency of the intrinsic oscillation.

doi:10.1371/journal.pone.0127269.g006

Planning in a complex and changing environment

To test our network with more challenging problems we considered classical path finding tasks that have also been suggested to rate animal intelligence [24]. Whether our agent finds the shortest path through a narrow hole in an obstacle depends on the size of the hole, and on the level of noise present in the planning network. While for small noise the slippage can be found without problems (Fig 7a), increasing the fluctuations in the external synaptic input to the planning neurons precludes the finding of the shortcut (Fig 7a, 7b). We next wondered whether the network can deal with a moving goal that changes its position while the agent is on its way. This is in fact possible without pausing to wait until the network relaxes to the new steady state. Once in the steady state, a continuous displacement of the goal leads to a continuous adaptation of the firing phases of the individual planning neurons, and a direct path to the moving target is found on the fly without delay (Fig 7c). Such faithful online modifications of the optimal path would be difficult to explain if the direction field were represented by asymmetric connections that are subject to synaptic plasticity obeying its own dynamics (e.g. by anti-STDP, see [12]). Finally, we challenged the network by a complex maze where a slight
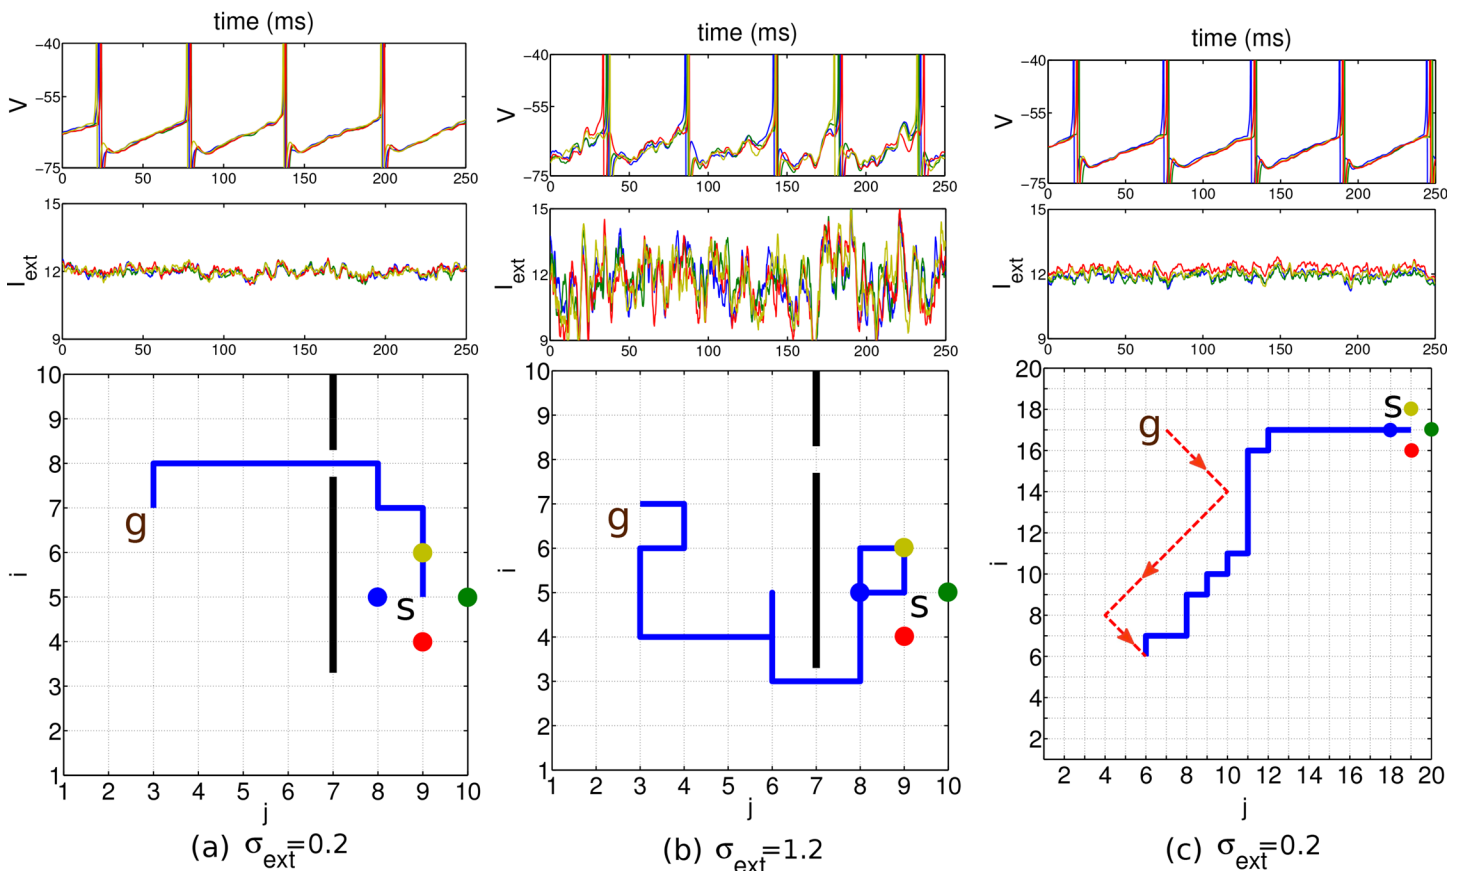

Fig 7. Finding shortcut and adaptive planning. (a, b) Turning up the noise level in the external input (indicated below each column) may be a way to prevent the detection of a shortest path. Other parameters as indicated in the caption of Fig 3. (c) Adaptive planning for a moving goal. A goal at initial goal position (17,7) is moved along the red line after planning time is over and the agent heads from the start towards this goal position. The goal moves one step at each readout cycle. The moving target can reliably be traced (blue line). The planning time is 1 s and the readout time is 250 ms. $\mu_{ext}^{goal} = 12.5$, $\mu_{ext} = 12$.
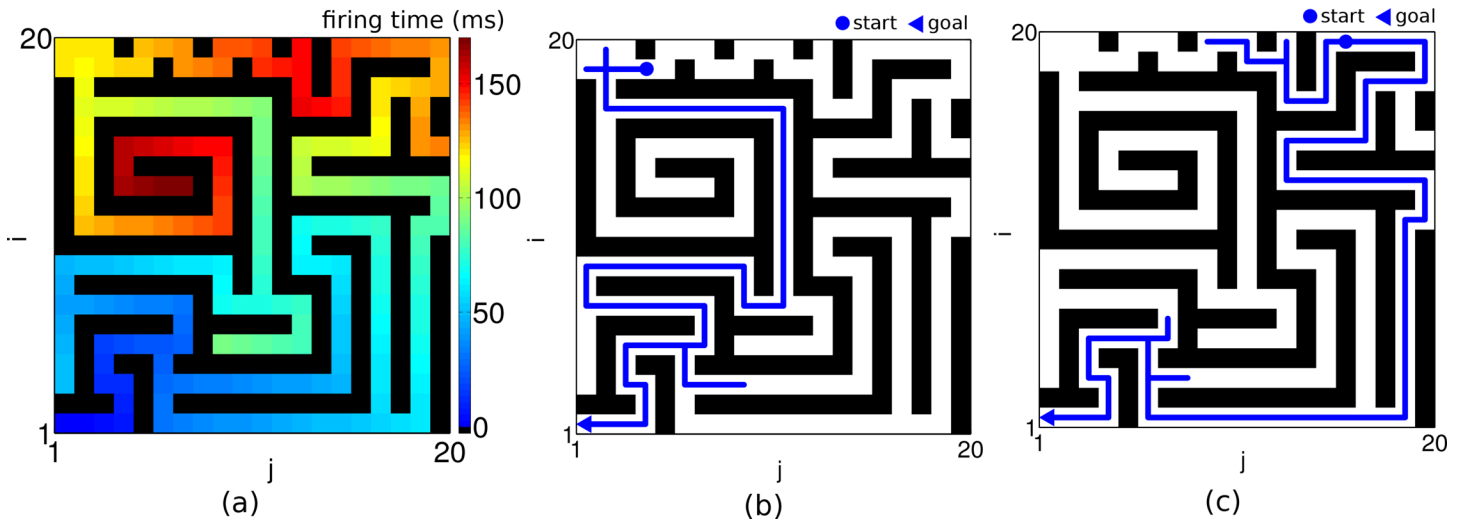
doi:10.1371/journal.pone.0127269.g007

Fig 8. Planning in complex environments. (a) Propagated wave of activity from the goal at (1,1) through a 20 × 20 network with obstacles (black bars) is demonstrated in space-space color coded plot of firing times relative to the goal neuron. The intrinsic frequency of the goal neuron is 18 Hz. (b, c) Two examples of navigation path (blue lines) in the network from different start positions (•) to the same goal (◄) such that a slight shift of the starting position implies an entirely different shortest path. The input to the planning layer is μ_{ext}^{goal} = 12.5 and μ_{ext} = 12 and the noise level is σ_{ext} = 0.7. The planning time takes 1 s and readout time is 150 ms.

doi:10.1371/journal.pone.0127269.g008

shift of the starting position implies an entirely different shortest path (Fig 8). Planning times, readout times and performance were as in an open environment.

Discussion

We have shown how directional information towards a goal can spread without information loss in a topographical network when the information is encoded in local phase differences of periodically firing neurons. In response to an additional drive of a selected goal neuron the network activity robustly self-organizes with a positive non-zero local phase difference even at large distances from the goal neuron. At any position in the network the direction of the shortest path to the goal is towards the neighbor with the smallest firing phase. The lower bound of these phase differences across the network guarantees a scale-free planning such that in the steady state, even arbitrarily far from the goal, a minimal directional information can be read out (S1 Text).

The backward spread of directional information from a goal location [7–12, 25] or both the backward spread from the goal and forward spread from the start [6] have earlier been suggested as strategies for navigational planning. However, these previous models suffer from the problem of a strong spatial information decay that causes them to break down in the presence of realistic neuronal noise. Based on the suggested phase coding, a scale-free and reliable planning now becomes possible within a single network. We have shown that the attractor property of the phase locking state makes the coding scheme robust against background noise generated from the stochastical firing of external neurons. Higher noise levels can be countered by additional readout cycles to be evaluated per action decision.

Spatial planning has been suggested to involve neurons in the medial prefrontal cortex (mPFC) that could be responsible for goal encoding [26]. In fact, impaired planning performance after mPFC lesioning [27], and human activity mapping during detour tasks, confirm the engagement of this area in navigating towards a goal [28, 29]. Wherever the goal encoding

neurons reside, our theory predicts that they should display a well-timed firing pattern as part of the direction encoding network. Crucially, this network does not need to uniformly tile the whole free space of an environment. Instead, the nodes of the planning network typically represent critical spots where directional decisions need to be taken, as in the vicinity of an obstacle or at crossing or bifurcation points of paths. If the synaptic strengths within the planning network are constant, the algorithm will find the shortest path within the graph of these nodes. While the shortest graph-theoretic path may not match with the shortest Euclidean path in the 2-dimensional environment, varying the synaptic strengths may convey additional distance information. For instance, a stronger synaptic strength will advance the firing of the postsynaptic neuron, and by virtue of moving towards earlier spikes, this implicitly tells that the path across this postsynaptic neuron is shorter.

Although we are considering a planning network with nearest neighbor connectivity, this network may still represent only a sparse selection of a larger network with many more neurons around each of our planning cells. When changing the environment, another sub-network with neurons that are physically close but have still only sparse overlap may be used for planning in this new environment. In a similar way as place cells in the hippocampus are remapped in a novel environment [30], the same planning neurons may be used in different spatial contexts.

Phase-coding of directional information has several hallmarks. First, the model predicts that the waiting time for an agent put into a well-explored maze until it heads off towards the home location increases with the distance, i.e. the length of the shortest path, to the home location. This waiting time corresponds to our planning time necessary for the network to settle in a steady state. Second, phase coding inherently assumes (relative) periodic firings that leads to activity oscillations, although it does not imply the existence of a global metronome (see Fig 3c and [31]). This is reminiscent to oscillatory activity during navigational planning observed in humans [32] and rats [33]. Third, to make a single decision about a new direction, oscillations must transiently build up in our planning network before the decision is taken. Such oscillations could be related to hippocampal-cortical oscillations observed when an animal must take a directional decision in front of the junction in a Y-maze [34]. Fourth, knowing the planning times of an agent that hypothetically uses our algorithm, we can predict the network size involved in this planning. If planning time is estimated from the waiting time at the start position to be 2 s, for instance, the steady state after this time is reached at roughly diagonal position 35 (Fig 6b), and this yields a squared network size of roughly 1200 neurons (or small neuronal populations), offering 1200 decision spots for the navigation in an environment. Multiple copies of such a network could provide redundancy and could therefore help to reduce the readout time, but not the planning time itself. So the limit on the network size imposed by the planning time implies a limit on the spatial resolution of the internal map.

In terms of human cognition the suggested breadth-first search algorithm may underly the pop-up effect of just 'seeing' the shortest path when navigating in a relatively simple environment or looking at a map thereof. In complex planning problems, however, the shortest path ceases to simply pop-up when the spatial resolution of the available network or networks becomes insufficient. We then need to resort to heuristic strategies such as defining intermediate goals to decompose the complex planning problem into a sequence of simpler tasks, such as known for the transition from parallel to sequential search and sequence representation [35, 36]. In this sense, on a behavioral level, our model predicts that planning time should at some point start to increase nonlinearly with task complexity and that this should go in hand with an increasing likelihood that the chosen path is suboptimal.

The planning model finally may be extended by a network of position cells that provides the required positional information. This would endow the model with different interesting

features. Such a position network may be used for mental navigation in terms of forward replay to new or old goals. In turn, reverse replay of past sequences in such a position network could be used for learning the backward connections in the planning network that point from the goal to the start, e.g. using classical spike-timing dependent plasticity [37]. Replay activity is in fact observed in hippocampal recordings [38]. Finally, because the movement is directed towards adjacent planning cells that fire earlier and earlier in phase, a place cell that receives input from the corresponding planning neuron would similarly advance its firing phase during the movement across its place field. Yet, as far as our planning network that operates by a phase code remains hypothetical, any re-interpretation of the experimentally observed phase precession (see e.g. [1]) in terms of a directional input from a planning network must also remain speculative.

Methods

Model equations and parameters

We modeled the voltage dynamics of the planning neuron at position (i, j) according to

$$C\frac{dV_{ij}}{dt} = -I_L^{ij} - I_{Na}^{ij} - I_K^{ij} - I_M^{ij} + I_{ext}^{ij}(t) + \epsilon I_{syn}^{ij} \ , \tag{2}$$

with leak current $I_L^{ij} = g_L(V_{ij} - E_L)$, sodium current $I_{Na}^{ij} = g_{Na}m^3h(V_{ij} - E_{Na})$, potassium current $I_K^{ij} = g_K n^4(V_{ij} - E_k)$, and outward potassium current with low threshold $I_M^{ij} = g_M q(V_{ij} - E_k)$. The subscript ij for the gating variables m, h, n and q is ignored to lighten the notation. The conductance is $C = 1$ (unitless), and the leak conductance is $g_L = 0.2$ (in units of ms^{-1}). All parameters taken from [17, 39], see also Supporting Information (S2 Text). The dynamics of the ion currents, in particularly the after-hyperpolarizing current I_M, implies that an action potential elicited by some planning neuron cannot fire a neighboring neuron that just fired before, and this guarantees a forward spread of each single activity wave throughout the network without reverberations.

The synaptic current $I_{syn}^{ij}(t)$ is obtained as a sum across the neighboring neurons (if existing),

$$I_{syn}^{ij}(t) = g_{syn} \sum_{k=i\pm1,\ l=j\pm1} s_{kl}(t)(E_e - V_{ij}) \ .$$

The conductance is set to $g_{syn} = 1$ if the place (i, j) can be reached and $= 0$ if it represents a fixed obstacle. The reversal potential is $E_e = 0$. The synaptic gating variable s_{kl} of the presynaptic neuron (k, l) describes the release probability as a function of the presynaptic potential, see [40] and Supporting Information (S2 Text). The external current $I_{ext}^{ij}(t)$ in Eq 2 that drives the planning neurons (i, j) is produced by N^{ij} afferents that are stochastically selected with connection probability $c = 0.8$ from a pool of N_E external neurons, each stochastically firing with a Poisson rate of v Hz. Denoting the spike times of the n'th afferent to neuron ij by t_n^{sp}, the synaptic strength by J and the synaptic time constant by τ_s ($= 2$ ms), the external current follows the dynamics

$$\tau_s \frac{dI_{ext}^{ij}}{dt} = -I_{ext}^{ij} + J\tau_s \sum_n^{N^{ij}} \sum_{t_n^{sp}} \delta(t - t_n^{sp}) \ .$$

This stochastic current can be characterized by its mean and standard deviation of the form $\mu_{ext}^{ij} = Jv\tau_s c N_E$ and $\sigma_{ext} = J\sqrt{\frac{1}{2} v\tau_s c N_E}$, respectively [41], both in units of mV/ms. A value of $\mu_{ext}^{ij} = 12.5$ and $\sigma_{ext} = 0.7$ corresponds to an average of $N_E cv = 73'500$ input spikes per seconds with a synaptic strength of $J = 0.081$ (in units of mV/ms), and for $\sigma_{ext} = 1.5$ we had $N_E cv =$

16′000 and $J = 0.375$. The case $\sigma_{ext} = 0$ was simulated by a DC current for I_{ext}^{ij} with mean μ_{ext}^{ij} (Figs 1 and 2).

For each neuron encoding a position in the environment, 4 action neurons are assigned in the action layer (Fig 2a). These action neurons are described by the leaky integrate-and-fire model with dynamics of the membrane potential

$$\frac{\mathrm{d}V_{a_{ij}}}{\mathrm{d}t} = -\frac{V_{a_{ij}}}{\tau_m} + I_{ext}^{a_{ij}} + I_{PR}^{a_{ij}}(t) + I_{RR}^{a_{ij}}(t) \ . \tag{3}$$

The index a stands for the cardinal directions W, E, N, and S that are represented by an action neuron at each position (i, j). Moreover, $\tau_m = 20$ ms is the membrane time constant, $I_{ext}^{a_{ij}}$ is a constant external input and $I_{PR}^{a_{ij}}$ and $I_{RR}^{a_{ij}}$ are synaptic inputs to action neurons a_{ij} from the planning (P) and readout (R) layer, respectively. An action neuron emits a spike whenever its potential reaches a threshold potential $V_{thr} = -50$ and is then instantaneously reset to $V_{reset} = -65$. The external input (e.g. from hippocampal place cells) to the action neurons is set to $I_{ext}^{a_{ij}} = 10$ if the agent is at position (i, j) and 0 else.

The action neuron a_{ij} receives synaptic input $I_{PR}^{a_{ij}}$ from the planning layer if its corresponding neuron emits a spike (see Fig 2a). More precisely, $I_{PR}^{a_{ij}}(t) = g_e\, s_{\alpha(i,j,a)}(t)$ with index function $\alpha(i, j, W) = (i, j-1)$, $\alpha(i, j, E) = (i, j+1)$, $\alpha(i, j, N) = (i-1, j)$, $\alpha(i, j, S) = (i+1, j)$. Here, $g_e = 0.8$ is the excitatory synaptic conductance and $s_{kl}(t)$ is again the synaptic gating variable driven by the membrane potential of the presynaptic neuron kl. To insure that at most the first of the four action neurons allocated to the same position is spiking we consider the mutual inhibition among these neurons of the form

$$I_{RR}^{a_{ij}} = g_{inh}\sum_{t_{a_{ij}}<t}\mathrm{PSP}(t - t_{a_{ij}}) \ ,$$

with postsynaptic potential $\mathrm{PSP}(t) = (t/\tau^2)\exp(-t/\tau)\Theta(t)$ characterized by $\tau = 2$ ms, $g_{inh} = -20$ and step function $\Theta(t) = 1$ for $t > 0$ and $\Theta(t) = 0$ else. The sum is taken over all spikes emitted by presynaptic neuron a_{ij} at times $t_{a_{ij}}$.

Scale-free planning: Non-decaying local phase differences

Here we give an intuitive account of why the local phase difference in the steady state is strictly positive throughout a network of arbitrary size. Let us assume that the local phase differences at some position were zero, i.e. the oscillator fire synchronously with its neighbors. In this case no interaction is possible and the oscillators fire with their common intrinsic frequency. But this contradicts the fact that the intrinsic frequency of the goal is higher than others, and that in the steady-state of a periodic traveling wave all oscillators fire with the same frequency. Next, to reach a common oscillation frequency, the difference between the driving currents of the goal and the remaining neurons, $I_{ext}^{goal} > I_{ext}$, needs to be compensated by different firing phases. As the individual synaptic currents have a strictly positive initial slope, we conclude that the firing phase difference between the current position and its neighbors can not be smaller than a fixed positive value. For an analytical explanation see Supporting Information (S1 Text).

Supporting Information

S1 Text. Non-fading directional information in the planning network: Mathematical proofs.
(PDF)

S2 Text. Details on the HH-type model.
(PDF)

Author Contributions

Conceived and designed the experiments: WS RU AKA. Performed the experiments: AKA. Analyzed the data: AKA WS RU. Contributed reagents/materials/analysis tools: AKA WS RU. Wrote the paper: WS AKA RU.

References

1. Buzsaki G, Moser EI. Memory, navigation and theta rhythm in the hippocampal-entorhinal system. Nat Neurosci. 2013 Feb; 16(2):130–138. doi: 10.1038/nn.3304 PMID: 23354386

2. O'Keefe J, Nadel L. The hippocampus as a cognitive map. Oxford: Clarendon Press; 1978.

3. Redish AD, Touretzky DS. Cognitive maps beyond the hippocampus. Hippocampus. 1997; 7:15–35. doi: 10.1002/(SICI)1098-1063(1997)7:1%3C15::AID-HIPO3%3E3.0.CO;2-6 PMID: 9138665

4. Dijkstra EW. A note on two problems in connexions with graphs. Numerische Mathematik. 1959; 1:269–271. doi: 10.1007/BF01386390

5. LaValle SM. Planning algorithms. Cambridge University Press; 2006.

6. Gorchetchnikov A, Hasselmo ME. A biophysical implementation of a bidirectional graph search algorithm to solve multiple goal navigation tasks. Connection Science. 2005; 17(1–2):145–166. doi: 10.1080/09540090500140925

7. Glasius R, Komoda A, Gielen SC. A biologically inspired neural net for trajectory formation and obstacle avoidance. Biol Cybern. 1996 Jun; 74(6):511–520. doi: 10.1007/BF00209422 PMID: 8672558

8. Gaussier P, Revel A, Banquet JP, Babeau V. From view cells and place cells to cognitive map learning: processing stages of the hippocampal system. Biol Cybern. 2002 Jan; 86(1):15–28. doi: 10.1007/s004220100269 PMID: 11918209

9. Hasselmo ME. A model of prefrontal cortical mechanisms for goal-directed behavior. J Cogn Neurosci. 2005;p. 1115–1129. doi: 10.1162/0898929054475190 PMID: 16102240

10. Martinet LE, Sheynikhovich D, Benchenane K, Arleo A. Spatial learning and action planning in a prefrontal cortical network model. PLoS Comput Biol. 2011; 7:e1002045. doi: 10.1371/journal.pcbi.1002045 PMID: 21625569

11. Hirel J, Gaussier P, Quoy M, Banquet JP, Save E, Poucet B. The hippocampo-cortical loop: spatio-temporal learning and goal-oriented planning in navigation. Neural Netw. 2013 Jul; 43:8–21. doi: 10.1016/j.neunet.2013.01.023 PMID: 23500496

12. Ponulak F, Hopfield JJ. Rapid, parallel path planning by propagating wavefronts of spiking neural activity. Front Comput Neurosci. 2013; 7. doi: 10.3389/fncom.2013.00098 PMID: 23882213

13. Dehaene S, Changeux JP. A hierarchical neuronal network for planning behavior. Proc Natl Acad Sci USA. 1997 Nov; 94:13293–13298. doi: 10.1073/pnas.94.24.13293 PMID: 9371839

14. Ermentrout GB, Kopell N. Frequency plateaus in a chain of weakly coupled oscillators, I. SIAM J Math Anal. 1984; 15:215–237. doi: 10.1137/0515019

15. Kopell N, Ermentrout GB. Symmetry and phase-locking in chains of weakly coupled oscillators. Comm Pure Appl Math. 1986; 39:623–660. doi: 10.1002/cpa.3160390504

16. Ermentrout GB, Kleinfeld D. Traveling electrical waves in cortex: insights from phase dynamics and speculation on a computational role. Neuron. 2001 Jan; 29(1):33–44. doi: 10.1016/S0896-6273(01)00178-7 PMID: 11182079

17. Ermentrout GB. The analysis of synaptically generated traveling waves. J Comput Neurosci. 1997; 5:191–208. doi: 10.1023/A:1008822117809

18. Bressloff PC, Coombes S. Travelling waves in chains of pulse-coupled integrate-and-fire oscillators with distributed delays. Physica D. 1999; 130:232–254. doi: 10.1016/S0167-2789(99)00013-5

19. Sompolinsky H, Golomb D, Kleinfeld D. Cooperative dynamics in visual processing. Phys Rev, A. 1991 Jun; 43(12):6990–7011. doi: 10.1103/PhysRevA.43.6990 PMID: 9905051

20. Rauch A, La Camera G, Luscher HR, Senn W, Fusi S. Neocortical pyramidal cells respond as integrate-and-fire neurons to in vivo-like input currents. J Neurophysiol. 2003 Sep; 90(3):1598–1612. doi: 10.1152/jn.00293.2003 PMID: 12750422

21. Goel P, Ermentrout B. Synchrony, stability, and firing patterns in pulse-coupled oscillators. Physica D: Nonlinear Phenomena. 2002; 163(3):191–216. doi: 10.1016/S0167-2789(01)00374-8

22. Lewis TJ, Rinzel J. Dynamics of spiking neurons connected by both inhibitory and electrical coupling. J Comput Neurosci. 2003; 14(3):283–309. doi: 10.1023/A:1023265027714 PMID: 12766429

23. Hansel D, Mato G, Meunier C. Synchrony in excitatory neural networks. Neural Comput. 1995 Mar; 7 (2):307–337. doi: 10.1162/neco.1995.7.2.307 PMID: 8974733

24. Hebb DO, Williams K. A method of rating animal intelligence. J Gen Psychol. 1946; 34:59–65. doi: 10.1080/00221309.1946.10544520 PMID: 21015350

25. Toussaint M. A sensorimotor map: Modulating lateral interactions for anticipation and planning. Neural Computation. 2006; 18(5):1132–1155. doi: 10.1162/neco.2006.18.5.1132 PMID: 16595060

26. Poucet B, Lenck-Santini PP, Hok V, Save E, Banquet JP, Gaussier P, et al. Spatial navigation and hippocampal place cell firing: the problem of goal encoding. Rev Neurosci. 2004; 15:89–107. doi: 10.1515/REVNEURO.2004.15.2.89 PMID: 15202682

27. Granon S, Poucet B. Medial prefrontal lesions in the rat and spatial navigation: evidence for impaired planning. Behav Neurosci. 1995 Jun; 109(3):474–484. doi: 10.1037/0735-7044.109.3.474 PMID: 7662158

28. Spiers HJ, Maguire EA. A navigational guidance system in the human brain. Hippocampus. 2007; 17 (8):618–626. doi: 10.1002/hipo.20298 PMID: 17492693

29. Viard A, Doeller CF, Hartley T, Bird CM, Burgess N. Anterior hippocampus and goal-directed spatial decision making. J Neurosci. 2011 Mar; 31(12):4613–4621. doi: 10.1523/JNEUROSCI.4640-10.2011 PMID: 21430161

30. Kubie JL, Muller RU. Multiple representations in the hippocampus. Hippocampus. 1991 Jul; 1(3):240–242. doi: 10.1002/hipo.450010305 PMID: 1669297

31. Nikolic D, Fries P, Singer W. Gamma oscillations: precise temporal coordination without a metronome. Trends Cogn Sci (Regul Ed). 2013 Feb; 17(2):54–55. doi: 10.1016/j.tics.2012.12.003

32. Kahana MJ, Sekuler R, Caplan JB, Kirschen M, Madsen JR. Human theta oscillations exhibit task dependence during virtual maze navigation. Nature. 1999 Jun; 399(6738):781–784. doi: 10.1038/21645 PMID: 10391243

33. Siapas AG, Lubenov EV, Wilson MA. Prefrontal phase locking to hippocampal theta oscillations. Neuron. 2005; 46:141–151. doi: 10.1016/j.neuron.2005.02.028 PMID: 15820700

34. Benchenane K, Peyrache A, Khamassi M, Tierney PL, Gioanni Y, Battaglia FP, et al. Coherent theta oscillations and reorganization of spike timing in the hippocampal- prefrontal network upon learning. Neuron. 2010; 66(6):921–936. doi: 10.1016/j.neuron.2010.05.013 PMID: 20620877

35. Nakayama K, Silverman GH. Serial and parallel processing of visual feature conjunctions. Nature. 1986; 320(6059):264–265. doi: 10.1038/320264a0 PMID: 3960106

36. Rhodes BJ, Bullock D, Verwey WB, Averbeck BB, Page MP. Learning and production of movement sequences: behavioral, neurophysiological, and modeling perspectives. Hum Mov Sci. 2004 Nov; 23 (5):699–746. doi: 10.1016/j.humov.2004.10.008 PMID: 15589629

37. Clopath C, Gerstner W. Voltage and Spike Timing Interact in STDP—A Unified Model. Front Synaptic Neurosci. 2010; 2:25. doi: 10.3389/fnsyn.2010.00025 PMID: 21423511

38. Diba K, Buzsaki G. Forward and reverse hippocampal place-cell sequences during ripples. Nat Neurosci. 2007; 10:1241–1242. doi: 10.1038/nn1961 PMID: 17828259

39. Traub RD, Jefferys JG, Miles R. Analysis of the propagation of disinhibition-induced after-discharges along the guinea-pig hippocampal slice in vitro. J Physiol (Lond). 1993;p. 267–287. doi: 10.1113/jphysiol.1993.sp019946

40. Golomb D, Amitai Y. Propagating neuronal discharges in neocortical slices: computational and experimental study. J Neurophysiol. 1997; 78:1199–1211. PMID: 9310412

41. Amit DJ, Brunel N. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cereb Cortex. 1997; 7(3):237–252. doi: 10.1093/cercor/7.3.237 PMID: 9143444