Neural Networks 23 (2010) 1139–1154

Contents lists available at ScienceDirect

Neural Networks

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

A computational model of a global neuronal workspace with stochastic connections

Dustin Connor*, Murray Shanahan

Department of Computing, Imperial College London, 180 Queen's Gate, London SW7 2AZ, England, UK

ARTICLE INFO

Article history:
Received 8 October 2008
Received in revised form 9 July 2010
Accepted 12 July 2010

Keywords:
Global workspace theory
Global neuronal workspace
Reverberating networks
Cortical competition

ABSTRACT

This paper presents a model of long-range cortical communication by means of a global neuronal workspace similar to that proposed by Dehaene and Naccache (2001). The model resembles that of Shanahan (2008), which was based on reverberating circuits of one-to-one connections, but uses a stochastic wiring regime in place of the highly regular scheme used there. The paper offers a systematic analysis of the influence of certain parameters on the dynamics of networks built according to this regime. Armed with a fuller understanding of the origins of the observed behaviour in the model, the qualitative behaviour of the previous model is replicated, and further unexplored choices are examined for ongoing work.

1. Introduction

Though most cognitive neuroscientists agree that the human brain integrates information from many different sources across functional modalities, it is still unclear how this is accomplished from an architectural and dynamical point of view. One theory which attempts to account for some of this information integration is Global Workspace Theory. Originally put forth and developed by Baars (1988, 2002), Dehaene and Naccache (2001), Dehaene, Kerszberg, and Changeux (1998) and Dehaene, Sergent, and Changeux (2003) used neural imaging data and neuropsychological testing to further postulate that the architecture presented by this theory is instantiated in the brain as a Global Neuronal Workspace (GNW).

Highlighted in the GNW is the role of neurons with long-range cortico-cortical connections, allowing communication between otherwise topologically and functionally distinct areas. Dehaene et al. (2003) hypothesized that access to and mobilization of these neurons accounted for several seemingly unconnected observations in neuropsychological testing data, such as correlation across distant regions and sustained activity in certain areas of the cortex. They then built a neural network model of cortical columns competing for access to this GNW and tested it according to a simulated paradigm borrowed from neuropsychology, the attentional blink. The predictions of the GNW continue to be tested and confirmed at present (e.g. Marzouki & Grainger, 2007; Van den Bussche, Segers, & Reynvoet, 2008).

While modelling access alone to the nominal GNW provided interesting results as well as further avenues of investigation, Shanahan (2008) extended the modelling of this structure to include the dynamics and interactions of the set of excitatory neurons with long-range connections to and from different areas, that is to say the hypothesized workspace neurons themselves. This work emphasized the central role of two particular aspects of the dynamics of the workspace neurons: broadcast and reverberation.

Broadcast is the ability of this set of neurons to disseminate patterns of activation across many otherwise disconnected groups, preserving the information inherent in their spatiotemporal structure. By utilizing the network of long-range connections afforded by a GNW, distant and functionally separate groups of neurons can coordinate and communicate.

Reverberation refers to the tendency of such a structure to sustain a pattern of activation over an extended period of time, usually through recurrent connections. This is not a new concept—synaptic reverberation has been thought to underlie persistent neural activity for 30 years and has been studied as a potentially critical aspect of working memory (e.g. Durstewitz & Seamans, 2006; Wang, 2001), and decision-making (Lo & Wang, 2006; Wang, 2002), amongst other things. Many of these studies specifically focus on the investigation and emergence of bistability between reverberation and decay of activation in neural dynamics (e.g. Bentley & Salinas, 2004; Kalitzin, van Dijk, & Spekreijse, 2000), but the overall effect of a sustained neural pattern of activation is shared as a predicted quality of a GNW.

Bringing the ideas of broadcast and reverberation together in a GNW framework, we present a model to investigate the role of the set of neurons with long-range connections within otherwise unconnected groups of neurons, in other words, the dynamics of a GNW-like network. The main aim of the presented work is to

* Corresponding author. Tel.: +44 779 621 3476.
E-mail address: dustin.connor02@imperial.ac.uk (D. Connor).

0893-6080/$ – see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neunet.2010.07.005

confirm the expected GNW behaviour using a less constrained model with fewer parameters, and a potentially more biologically plausible setup than previous models.

While the currently presented model removes some parameters and constraints from previous work, there are still distinct choices which could be questioned. Our goal is to present this work as a useful step towards exploring a space of possible instantiations of the GNW concept, in which this particular stochastic setup is one point, the model presented in Shanahan (2008) is another, and the information integration across many different sources instantiated in the human brain may be another.

In this paper we first discuss the model as a whole, and the construction of a GNW-like structure within that model, paying special attention to the reasoning behind many of the choices made during this construction. We then single one parameter of the model out for the in-depth study, the scaling factor for synaptic weights, and analyse the changes in model behaviour due to varying this parameter over a broad range of values. Furthermore, we describe the implementation of this GNW in a wider model and reproduce results from similar work as a proof-of-concept exercise. Finally, we discuss future work with this type of model and the exploration of the space of possible instantiations of the GNW as a whole.

2. The model

The model consists of a single set of 1600 spiking neurons representing the workspace (1280 excitatory, 320 inhibitory) and several separate groups of neurons attached to the workspace representing functionally distinct neuronal clusters, or specialists, each with 2048 spiking neurons (1638 excitatory, 410 inhibitory). The ratio of roughly 4:1 excitatory to inhibitory neurons matches earlier findings regarding such ratios in the mammalian cortex (Binzegger, Douglas, & Martin, 2004; Gabbott & Somogyi, 1986). The specialist areas should be considered as distant from each other, the only connections between them being mediated by the workspace (Fig. 1). The specialists are reciprocally connected to a range of workspace neurons—the spatially nearest set of neurons in the workspace. Neurons within a single specialist area are also densely connected to each other but are not connected to neurons in a different specialist area except via the workspace. For some experiments (see Section 4), a competition is built in to two neighbouring specialists, such that they each react to the same input state, but with different responses. All such cases are accompanied by an additional layer of diffuse lateral inhibition between the two specialists.

The workspace neurons are organized as two one-dimensional rings: one with excitatory neurons, connected to each other in self-excitatory cycles of long-range connections in order to best reflect the underlying concepts of reverberation (self-excitatory cycles) and broadcast (long-range connections), and one with inhibitory neurons, connected diffusely to and from the excitatory neurons (see Section 2.1 and Figs. 2 and 4). It is important to note that the ring organization represents a spatial embedding, not the actual connectivity. It is from this spatial embedding that a distance metric is defined, and thus neurons are connected according to long- or short-range connectivity rules (see Sections 2.1 and 4.1 for more detail).

A state of activation can start and be sustained in the workspace while propagating quickly to different localities via the excitatory cycles, eventually to be picked up by the specialists. One or more of these specialists may respond to the state of workspace activation and will attempt to influence the workspace in turn.

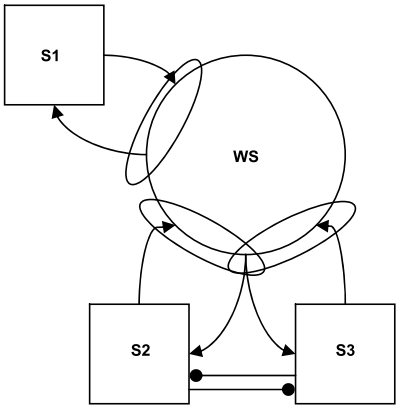

Fig. 1. Overall schematic of the model. Connections within the workspace area (WS) allow for broadcast and reverberation of a state, while connections within the specialist areas (S1-3) allow them to learn a spatiotemporal pattern in order to have an effect on the state in the workspace. Connections between the workspace and the specialist areas are highly focal and have only a small spread of conduction delays, allowing for the maintenance of the spatiotemporal properties of a state of activation between the different areas. S2 and S3 are arranged as competing neighbours, with lateral inhibition. Circled sections of the workspace denote the range of workspace neurons to which each specialist is connected.

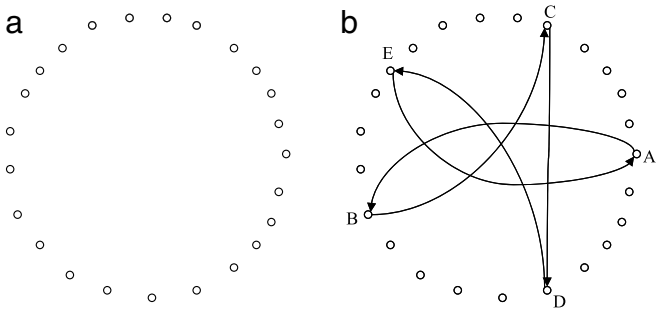

Fig. 2. (a) The neuronal workspace ring arrangement before being connected. (b) An example of a cycle of connections ($A \rightarrow B \rightarrow C \rightarrow D \rightarrow E \rightarrow A$). Each neuron would be connected in one and only one cycle of this type. The distance parameter is typically chosen so that this occurs after an average of four or five neurons, usually sufficient to have a representative in any part of the workspace where a specialist area might be attached.

Diffuse inhibition in the workspace increases as total workspace activation increases, keeping total activation in the workspace from continually growing. Only a new state driven directly by a specialist will be capable of taking over the workspace to be broadcast. After this competition, there will only be trace remnants of the previous reverberating state.

For a spatiotemporal pattern to be transmitted from specialist to workspace to the other specialists, it is important that the spatiotemporal properties of a state of activation be preserved. To this end, connections between different neuronal groups (e.g. from specialist area 1 to the workspace) are focal with only a small spread in conduction delay. In addition, the cycles of excitatory connections in the workspace have no connections between cycles, and a workspace neuron belongs to one and only one cycle. This ensures that a state of activation reaching the workspace does not subsequently smear across other possible states.

The workspace could theoretically be connected to any number of specialists, and a more complete model might have neighbouring specialists attached all around the workspace. The dynamics studied here, however, requires only the use of three specialist areas, and this is chosen to minimize run times of the model, thus allowing for more in-depth study of the parameters involved.

```
α ← 192; x ← 0; y ← 0; z ← 0
A ← All Neurons; B ← Ø; C ← Ø; D ← Ø
While A ≠ Ø
    While A ≠ Ø
        x ← rand from A
        A ← A/{x}
        C ← C ∪ {x}
        B ← B ∪ (A ∩ {n : dist(n, x) ≤ α})      // dist(n,x) is the shortest
        A ← A/{n : dist(n, x) ≤ α}              // distance around the ring
        If y ≠ 0                                // between neurons n and x
            Make connection from y to x
        Else
            z ← x
        End
        y ← x
    End
    If x = z
        D ← D ∪ {x}                             // Check if x is an orphan
        C ← C/{x}
    Else
        Make connection from x to z
    End                                         // Reset neurons for next
    A ← A ∪ B                                   // cycle
    B ← Ø; x ← 0; y ← 0; z ← 0
End
```

Fig. 3. Psuedo-code to describe the construction of connections between the excitatory neurons of the workspace. *Key*: $\alpha \rightarrow$ short-distance parameter, $x \rightarrow$ label of the current neuron, $y \rightarrow$ label of the previous neuron, $z \rightarrow$ label of the first neuron in this cycle, $A \rightarrow$ set of all neurons available for connection, $B \rightarrow$ set of all neurons unavailable for connection just for this cycle, $C \rightarrow$ set of all neurons already connected in a cycle, $D \rightarrow$ set of all 'orphan' neurons, neurons with no possible long-range connections in the current setup.

2.1. Workspace structure and connectivity

The key to understanding the construction of the workspace area of neurons lies in grasping the principles and rules by which the self-excitatory cycles of long-range connections are created.

The neurons contained in the workspace are distinguished in this model from neurons in the specialist groups as those neurons which have long-range, reverberation-inducing connections, appropriate for sustained broadcast. All connections within the workspace are constrained to be of long range by ensuring that the distance between any two connected workspace neurons is greater than a predefined parameter, α. The distance between neurons is measured by the number of neurons which lie between them on the ring via the shortest path.[1]

Next, the workspace neurons should promote reverberation via these long-range connections. To this end, self-excitatory cycles, loops of neuronal connectivity, are created such that neuron $A \rightarrow B \rightarrow C \rightarrow D \rightarrow E \rightarrow A$ (Fig. 2(b)). Keeping with the minimum distance parameter, each of these neurons is in a distant part of the workspace compared to the other neurons in that cycle. The neurons of a given cycle act as representatives of the activity of the cycle for that specific area of the workspace, and hence for any specialists attached to that area of the workspace. If A is active, it will in turn activate B, C, and eventually A again, and this will be sustained until stopped by inhibition. This occurs either in the form of slow decay with time, or an abrupt and rapid decay of the state caused by a competing state entering the workspace.

It is also assumed for this model that the *only* role of this communication infrastructure is to broadcast and sustain activation. This means a state of activity of neurons passing through the workspace does not change, with some allowances made for a reasonable level of decay in the state. To accomplish this, the cycles of connections cannot overlap—each neuron in the workspace must belong to one and only one cycle, and connections between cycles are disallowed.

Aside from the aforementioned restrictions, all connections within the workspace are assigned stochastically in order to remove any organizational biases that a more strictly engineered approach may produce. The process of making connections between the workspace neurons using these restrictions is described by the pseudo-code in Fig. 3.

Typically, we choose a distance parameter of $\alpha = 192$, which, for a workspace of 1280 excitatory neurons, gives an average circuit size of approximately 5 neurons and 256–260 total circuits in the workspace, with very little variation. This allows for natural comparison with the one-to-one connections between five distinct workspace areas of 256 neurons each in Shanahan (2008). Despite this natural comparison, however, it should be noted that there are several significant deviations from the work described in Shanahan (2008):

1. The treatment of the GNW as the one set of neurons with stochastic connections, governed by certain rules, which are chosen to emphasize the aforementioned dynamics. This is in place of the highly engineered 'workspace areas' with one-to-one connections and direct mappings in Shanahan (2008).
2. The exclusion of special 'access areas' to act as mediums between neuron clusters and the workspace itself.
3. The connectivity of each workspace neuron is reduced—each excitatory workspace neuron stimulates only one other excitatory neuron, while each excitatory workspace neuron in Shanahan (2008) stimulates two other excitatory workspace neurons.

Both approaches, and indeed an entire spectrum of possible approaches, can provide different functionality and raise different questions; however, these are beyond the scope of the current work. These particular choices were made in an attempt to create a less constrained model: a single general space of neurons for the GNW rather than an arbitrary number of workspace areas, the removal of excess areas with non-obvious justification or function, and the use of the most reduced connectivity scenario. While some constraints still remain in this model, the removal of constraints from previous models is an important step. Fewer constraints can clarify the mechanistic role of parameter choices in the model, as well as of the remaining constraints. Ultimately, continued work in this direction could lead to conditions (necessary and/or sufficient) for a model to retain the important dynamics observed here and in previous work.

The organization of the excitatory workspace neurons is based on the principles of promoting reverberation and broadcast, and keeping a state of activation distinct from other possible states of activation. The connections of the workspace inhibitory neurons, on the other hand, are organized to reflect the ability of a new state entering the workspace to inhibit and subsequently knock out a currently reverberating state. The desired effect can be summarized as the one that samples the total activation of the workspace, independent of the circuit construction and inhibits diffusely proportional to the activation in the workspace. With this in mind, the inhibitory neurons are set as a separate ring, concentric to the first (Fig. 4 (a)). The inhibitory neurons

[1] An actual distance (in millimeters, for example) could be assigned to the ring and neurons placed on the ring according to a density function. However, we find this to be unnecessary in this model. Such an approach can be compared with Izhikevich, Gally, and Edelman (2004) where neurons are placed on the surface of a sphere and biologically significant distances are used. An accurate model of the brain would be much larger in scale than the order of thousands of neurons used here, and would benefit from such an organization. However, the aim of this paper is to explore some of the underlying concepts, parameters and behaviours of this type of network in depth, and this is made feasible by using a smaller model.

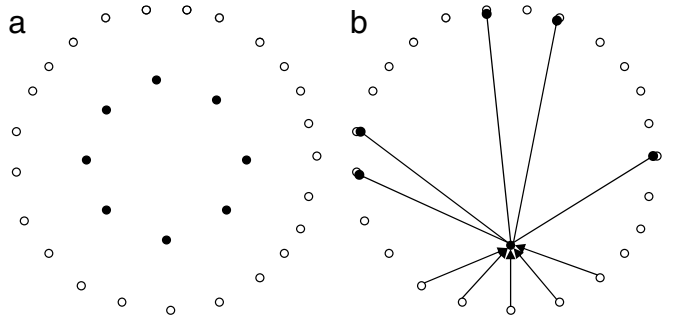

Fig. 4. (a) Spatial organization of workspace inhibitory neurons. (b) The inputs and outputs of a typical inhibitory neuron. Inputs are locally organized, while outputs are randomly assigned. For a workspace of 1280 excitatory neurons, there are 320 inhibitory neurons, each with 200 inputs and outputs.

take diffuse input from many local excitatory neurons, and then randomly and diffusely inhibit the rest of the workspace when a certain level of activity has been reached (Fig. 4(b)).

The theoretical ideal setup is the one where there is *no* inhibition for an accepted level of workspace activity, and non-zero inhibitory activation when workspace activity exceeds this—representing, for example, an additional state entering the workspace. This would inhibit activation in the workspace until it is reduced back to the acceptable level. It is assumed that a strongly driven state of activation would remain, while only traces of activation would remain from previous 'competitions'. This would be reminiscent of the strictly bistable states discussed throughout the literature on reverberating neural activity and working memory—a state of activation is in the workspace and persistent, then it is quashed, replaced by a different state. In this case, each workspace state could be considered bistable (active or not), while the workspace as a whole would has multistability (one stable regime for each possible active state).

However, it is found that this particular model of inhibitory neurons always collectively produce a minimum (non-zero) level of inhibition with just a minimal base current input, and this increases proportional to any added input. A more complicated neural setup could simulate the strict inhibitory behaviour of the type described above. However, in the interest of keeping the model as general as possible, it is necessary to maintain a minimal number of constraints. In addition, the recent literature (e.g. Durstewitz & Seamans, 2006) has pointed out the growing evidence that strict 'on–off' dynamics are rare during working memory tasks, and that more complicated temporal dynamics are likely at work.

We instead accept that a reverberating state in the workspace is not strictly persistent and will slowly decay with time. Additional activation in the workspace (from a new state of activation, for example) will invoke a larger inhibitory response and greatly increase the rate of decay, resembling an 'off' switch but not completely acting as one.

2.2. State of activation

We choose a simple form for spatiotemporal patterns of activation in order to test the basics of this model. With a workspace divided into 256–260 cycles of connections, these cycles can be divided by number into four 'quadrants' of 64 cycles each (1–64, 65–128, 129–192, 193–256), with cycles above 256 being ignored for these purposes. Each pattern is represented by simultaneous activity of a specific group of cycles within one of the quadrants firing (approximately 60% of the quadrant, or 37 cycles). These groups of cycles are chosen for ease of visual representation only, and realistically there is no reason that the entire quadrant could not be active for each state.

States of activation which make use of overlapping circuits (e.g. cycles 20–40 in one state and 30–50 in another) could also be considered. In addition to the spectrum of possible network structures with different constraints previously mentioned, there is a spectrum of possible spatiotemporal patterns to consider. This represents a significant departure in complexity from this model, however, and poses questions for potential future work with this kind of model.

2.3. Izhikevich neurons

Individual neurons were simulated using Izhikevich's "simple model" of spiking behaviour (Izhikevich, 2003, 2007). This model is able to generate a large range of empirically accurate spiking behaviours, like the Hodgkin–Huxley equations, while being much easier to compute with. It is thus well suited to a large-scale, biologically plausible simulation. Moreover, the behaviour of the model is governed by four parameters (a, b, c, and d in Eqs. (1)–(3) below), which can be varied to emulate the signalling properties of a wide variety of known neuron types. The model is defined by the following three equations:

$$\dot{v} = 0.04v^2 + 5v + 140 - u + I \tag{1}$$

$$\dot{u} = a(bv - u) \tag{2}$$

$$\text{if } v \geq 30 \text{ then } \begin{cases} v \leftarrow c \\ u \leftarrow u + d \end{cases} \tag{3}$$

where v is the neuron's membrane potential, I is its input current, and u is a variable that regulates the recovery time of the neuron after spiking. Eq. (3) describes the way the neuron is reset after spiking, which is assumed to occur when its membrane potential reaches 30 mV.

The values of the four parameters a, b, c, and d were lifted from Izhikevich (2003). For excitatory neurons, these were $a = 0.02$, $b = 0.2, c = -65 + 16r^2$, and $d = 8 - 6r^2$, where r is a uniformly distributed random variable in the interval [0, 1]. For inhibitory neurons, the values used were $a = 0.02 + 0.08r$, $b = 0.25 - 0.05r$, $c = -65$, and $d = 2$, with r as above. The random variable r introduces a degree of variation into each population giving a range of different biologically plausible neural firing patterns. These parameter choices place all neurons on an appropriate scale of behaviours with a random factor introduced to ensure that model dynamics are not a serendipitous coincidence of a homogeneous neural distribution. Excitatory neurons can be between regular spiking or chattering types, with a bias towards regular spiking, while inhibitory neurons are distributed evenly between fast spiking and low-threshold spiking dynamics. This is not intended to model distributions in a particular area of the cortex, instead it reflects neural types found in the cortex with a heterogeneous distribution to avoid unintended dynamical artifacts (but see Izhikevich, 2003; Izhikevich et al., 2004, for more detail). For example, in the excitatory case, if $r = 0$ we get the regular spiking behaviour shown on the left of Fig. 5, while if $r = 1$ we get the chattering behaviour shown on the right of the figure.

Consider a time t and a neuron i, and let Φ be the set of all neurons j that fired at time $t - d$ where d is the conduction delay from neuron j to i. Then, the input current I for neuron i at time t is given by

$$I(t) = I_b + \sum_{j \in \Phi} S_{i,j} F, \tag{4}$$

where I_b is the base current, $S_{i,j}$ is the synaptic weight of the connection from neuron j to i, and F is a scaling factor whose value depends on the type of population to which i and j, respectively, belong (e.g.: workspace area, lateral inhibitory pool, etc.) and does not change with time. These scaling factors are an

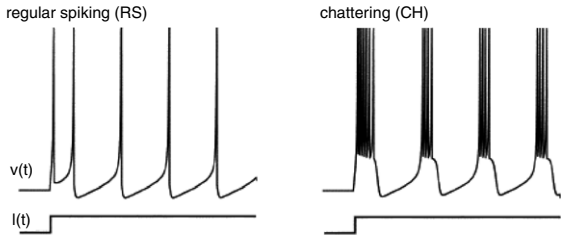

Fig. 5. Varieties of excitatory neurons using Izhikevich's simple model (from Izhikevich (2003)—electronic version of the figure and reproduction permissions are freely available at www.izhikevich.com).

important parameter of the model and are the focus of the reported investigation of the effect of parameter choice on model dynamics.

Conduction delays in the workspace are chosen to be either 5 or 6 ms at random. These are the minimum conduction delays for which reverberation is consistently present—reverberation is limited by neuron recovery times when lower conduction delays are used. Higher conduction delays were tested (up to 20 ms) and the overall dynamics were found to be robust, so these are chosen in order to minimize run times. Two delays with a millisecond difference are chosen so there are fewer manufactured timing effects. For example, one can imagine very specific resonances or frequencies caused entirely by an arbitrary choice of a single delay. Possible variations in future models could include a much larger spread of delays, perhaps dependent on the distance of the connection. (See Izhikevich & Edelman, 2008; Izhikevich et al., 2004, for this kind of treatment), or conduction delays directly taken from biological data (inspired by precisely tuned myelinated axons, for example).

The scaling factor of a connection, together with the synaptic weight, represents the overall strength of that connection—how much input current the postsynaptic neuron receives from a spiking presynaptic neuron. Ignoring for the moment unit scaling, this can be considered to represent a few things. In the current model, we effectively absorb the synaptic weight into the scaling factor by setting the synaptic weights in workspace connections to one. This means that the scaling factor can be seen to represent the strength of a connection, or the amount of influence the presynaptic neuron holds over the postsynaptic neuron, perhaps affected by some kind of long-term potentiation or depression. Another interpretation, in light of the relatively small neuron population sizes, might consider each connection to represent many related connections, where a higher scaling factor could signify more presynaptic neurons represented by that single model connection. Here we consider it generally as a measure of strength of connection between neuron populations, regardless of whether this comes about from more influential synapses or simply more neurons and more connections.

3. Studying the workspace

The balance between excitation and inhibition in the workspace is a key feature of this model; this balance governs the strength of reverberation of a state, the sensitivity of the workspace to new patterns of activation, and the speed of decay of a state due to inhibition. The level of excitation and inhibition in the workspace is largely moderated by the choice of scaling factors to and from excitatory and inhibitory neurons. Consequently, we present a systematic investigation of the effect of different choices of scaling factors on the dynamics of states of activation in the workspace. This is done with the intention of making informed decisions from what would otherwise be arbitrary parameter choices in future work.

With only the workspace excitatory and inhibitory groups of neurons, there are three kinds of connections—workspace excitatory to excitatory, workspace excitatory to inhibitory, and workspace inhibitory to excitatory, noting that there are no inhibitory to inhibitory connections. For each of these a different scaling factor (F_{ww}, F_{wi}, and F_{iw} respectively) is set for all pulses travelling via those connections.

The progression and reverberation of a single state in the workspace is studied with a single pulse of 35 mA delivered to neurons belonging to the appropriate 37 cycles of connections in the workspace after 20 ms of rest.

3.1. Workspace excitatory to excitatory

Even without inhibition (F_{wi} and $F_{iw} = 0$), activation in the workspace does not reverberate in any circuits for workspace to workspace scaling factors (F_{ww}) less than 25, and does not consistently reverberate in all stimulated circuits for F_{ww} less than 40 (Fig. 6). For F_{ww} greater than this, activation in the workspace has an initial peak followed by a steady state, both of which increase in magnitude as the scaling factor is increased (Fig. 7(a)).

The relationship between F_{ww} and activity in the workspace (measured as the average steady-state firing rate) is displayed in Fig. 7(b). As one would expect, as F_{ww} increases, so does activity in the workspace; however, the direct relationship between the two is not obvious. Initially, activity in the workspace shows a shallow linear progression as F_{ww} increases, up to $F_{ww} \approx 80$. At this point, a phase shift occurs and a steeper linear increase is observed. This continues until $F_{ww} \approx 130$, where there is a significant change in the dynamic between F_{ww} and activity in the workspace. While it is unknown why these phase shifts occur, there are several likely possibilities. The size and nature of the self-excitatory cycles of connections, and neuron firing strongly overlapping with neuron recovery times seem the most likely contributors, though further investigation is needed to pin down exactly what effect these factors have.

In order to keep consistent and robust reverberation in the workspace which is still sensitive to inhibition, typically $60 \leq F_{ww} \leq 100$ is used in experiments with specialist areas attached, though it is noted that a much more pronounced change in excitation, especially in the initial peak, occurs in the interval between 80 and 100. To best study the effect of increased activation in the workspace (and thus F_{ww}) on other aspects of the model, $60 \leq F_{ww} \leq 100$ is also used for subsequent experiments with only the workspace neurons.

3.2. Workspace excitatory to inhibitory

We next introduce non-zero F_{wi} in order to study the dynamics of the workspace inhibitory neurons with unchanging activation in the workspace excitatory neurons ($F_{iw} = 0$). With these parameters, the inhibitory neurons fire due to input from the excitatory population, but without resulting in actual inhibition of the excitatory population. Unlike the excitatory neurons, there is a low ambient level of firing in the inhibitory population with just a low base current (2 mA). This firing is very small compared to increased firing due to workspace excitation (2.5 firings per ms), but is enough to cause a slow decay of workspace activation with non-zero F_{iw}.

As connections from the excitatory workspace population to the inhibitory are diffuse (200 per inhibitory neuron), we expect a much stronger proportional increase in inhibitory firing from a change in F_{wi} when compared with the increase in excitatory firings from a change in F_{ww}. Thus, the interval over which F_{wi} is studied is lower than F_{ww} with smaller increments (1–10).

With non-zero F_{wi}, the dynamics of firing of the inhibitory neurons become tied to the dynamics of the excitatory neurons.

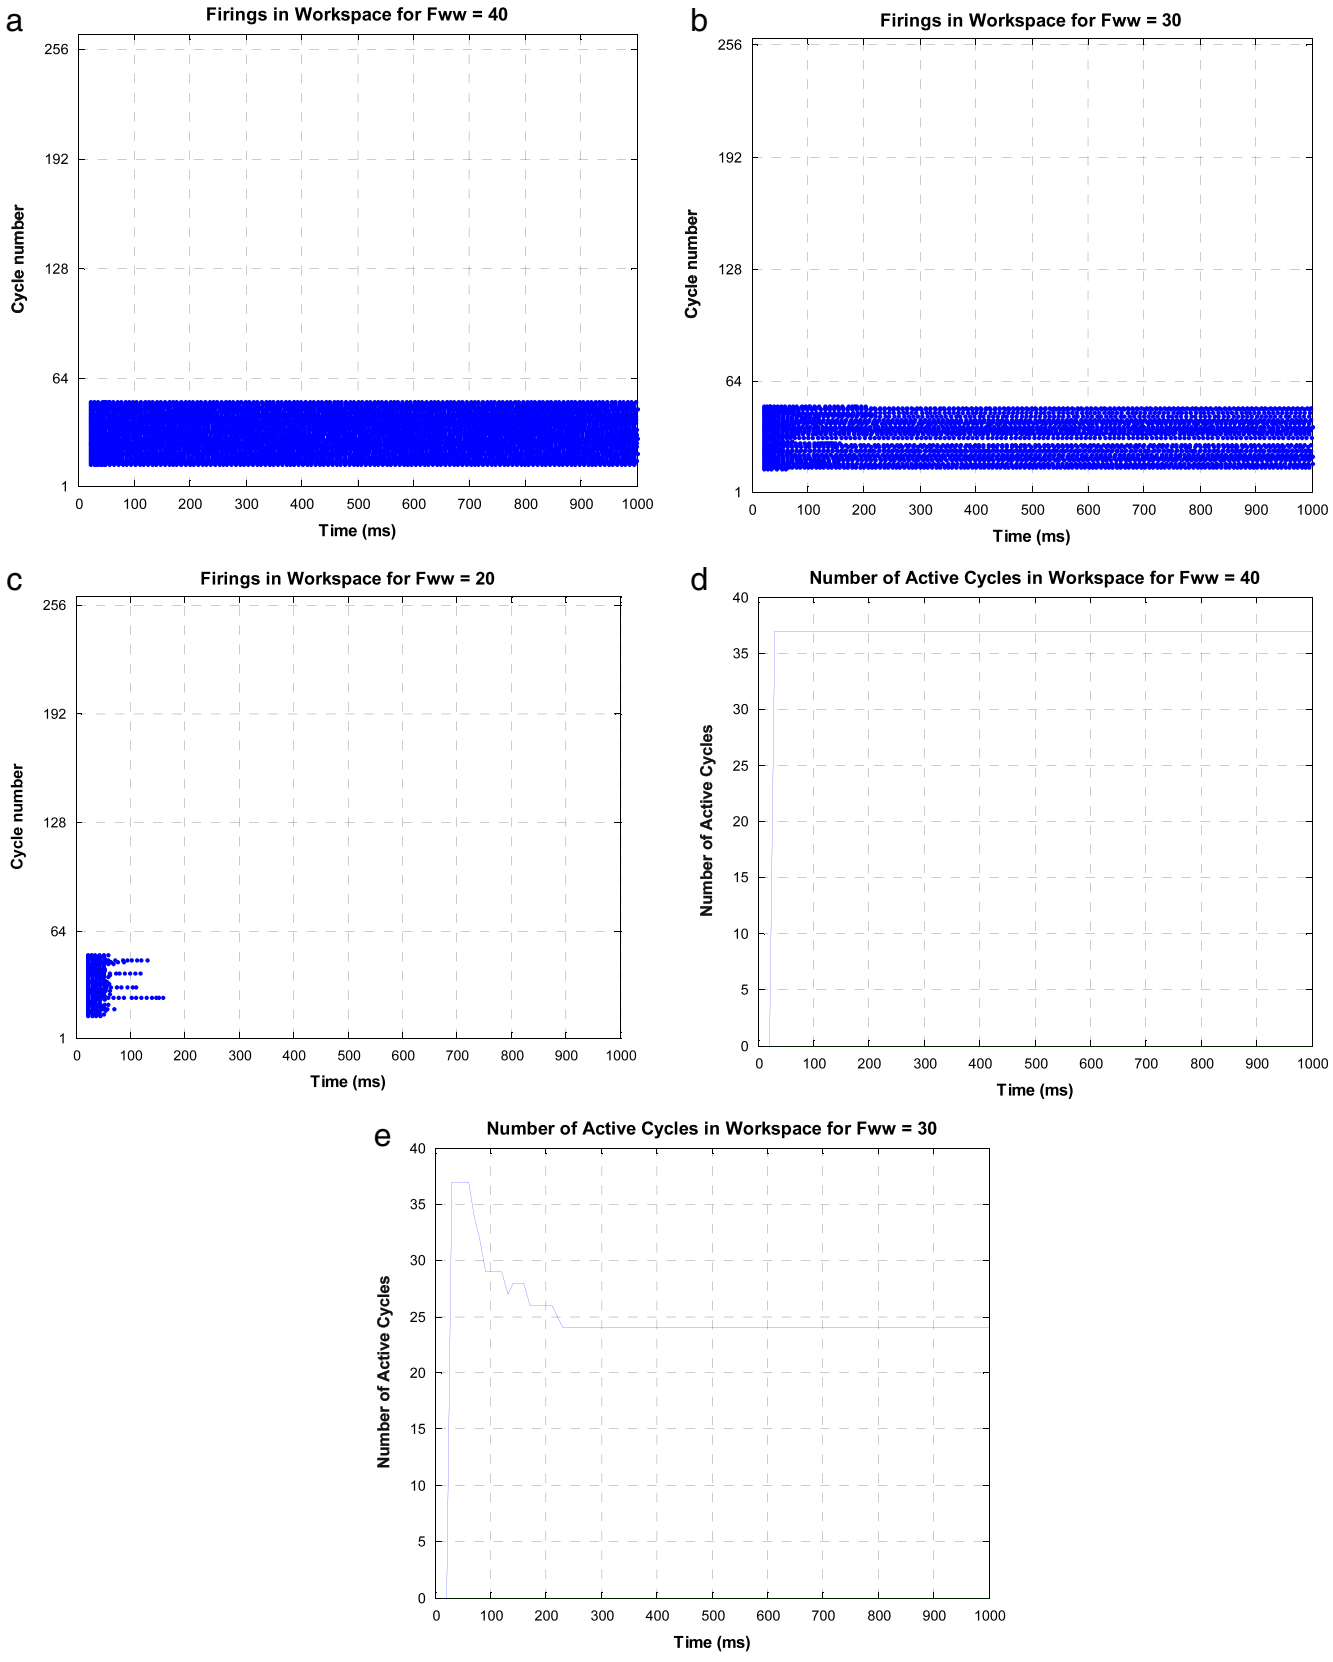

Fig. 6. (a)–(c) Representative raster plots for $F_{ww} = 40, 30, 20$. Each point represents that at least one of the neurons belonging to that cycle has fired. (d)–(e) A plot of the number of active cycles (measured by at least one neuron of that cycle firing in a 10 ms sampling of time) for $F_{ww} = 40$ and 30. 37 cycles are initially stimulated.

A small peak or a high peak followed by a steady state in the excitatory neurons leads to the same in the inhibitory neurons (Fig. 8(a)). Much like the analysis of the excitatory dynamics with F_{ww}, the activity of the inhibitory population of workspace neurons is measured via the average firing rate of the inhibitory neurons once it has settled from the initial peak.

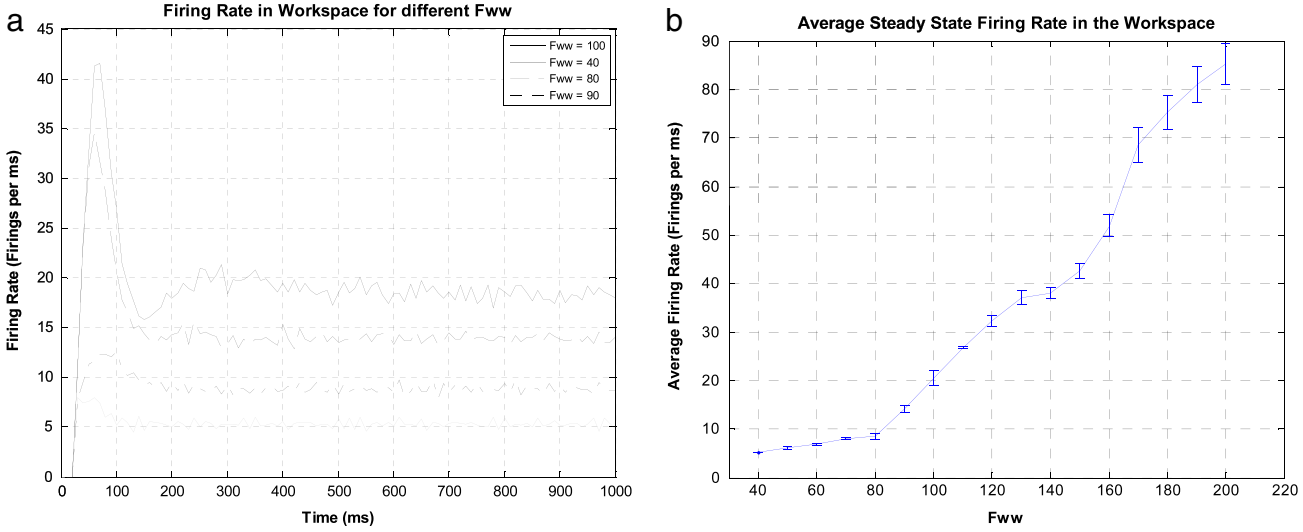

Fig. 7. (a) Progression of the firing rate in the workspace excitatory population with time (averaged in 10 ms windows of time) in a few typical model runs with F_{wi} and $F_{iw} = 0$, and varying F_{ww}. There is an initial peak in activity immediately after presentation of a stimulus, which drops into a steady state. (b) A plot of the average steady-state firing rate as changed by F_{ww}. There seem to be several phase shifts where the increased scaling factor causes a shift in the dynamics of the model. All data points are averages over several trials with different stochastic workspace connections.

Unlike the phase shifts in the dependence on F_{ww} of the activity of excitatory workspace neurons, the inhibitory neurons show a much more direct increasing dependence on both the activity in the workspace and the scaling factor F_{wi} (Fig. 8(b)–(c)). The clear change in dependence of inhibitory activity on F_{ww} between $F_{ww} = 80$ and $F_{ww} = 90$ closely mirrors the observed phase shift in excitatory activity with F_{ww} in that range. This indicates, as one would expect, the inhibitory firing is directly dependent on the amount of activation in the workspace, not the scaling factor F_{ww} itself.

While it may initially seem as if the inhibitory neurons fire far more often than the excitatory neurons, this can be deceptive. With only one active state in the workspace, only 37 cycles of excitatory neurons are active. With 4–5 neurons per cycle, this gives a maximum of 185 active excitatory neurons. On the other hand, with highly diffuse connections and activity spread across the workspace, it is possible that all 320 inhibitory neurons are firing. Considering the differences in connectivity between the two groups as well as differing conduction delays, it is not obvious how to scale the firing rate data so that the two different groups are comparable.

3.3. Workspace inhibitory to excitatory

Due to the isolated nature of the cycles in the workspace (there is no excitatory interconnection between them), an active cycle which stops firing due to inhibition will not start firing again without a new stimulus. With non-zero F_{iw}, the previously observed steady state of activation after a period of time instead becomes a steady decay in the firing rate for both excitatory and inhibitory workspace neurons. Activation in the excitatory neurons increases activity in the inhibitory neurons, which then feed back into the excitatory neurons randomly and diffusely, thus inhibiting some of the reverberating cycles. The decreased activity in the excitatory neurons drives less activity in the inhibitory population, which in turn inhibits less, causing the speed of decay of the state of activation to drop.

As with F_{wi}, F_{iw} are chosen in a range lower than F_{ww} (1–10). Some representative trials with different parameter values are shown in Fig. 9. The most significant decay has typically finished by $t = 200$ ms, and there is rarely observable ongoing decay after $t = 300$ ms except in cases of high inhibition. After this, there seems to be a balance of excitation and inhibition, sufficient to sustain roughly constant reverberation. Activation in the workspace for times 500 ms $\leq t \leq$ 600 ms, measured via the average firing rate of the excitatory neurons, is shown in Fig. 10 for the full range of inhibitory parameters tested.

The scaling factor F_{wi} is nearly interchangeable with F_{iw} in contribution to the total inhibition of the workspace, but it seems changes in F_{iw} have a slightly stronger effect (Figs. 8 and 9). In general, we can choose equivalent inhibitory scaling factors ($F_{wi} = F_{iw}$) for simplicity, with incremental changes possible in small (change in F_{wi}) or large (change in F_{iw}) increments. While there is strong decay in the state of activation for high combined inhibitory factors ($F_{wi} + F_{iw} \geq 10$), there are very few parameter choices which lead to total loss of activation (Fig. 9). This suggests there will often be at least small remnants of previously reverberating states if a series of states is considered.

How, if at all, does the excitatory scaling factor, F_{ww}, contribute to the balance between excitation and inhibition? Comparing $F_{ww} = 80$ to $F_{ww} = 100$, where there is a known large increase in activation due to the scaling factor, we see a typically increased amount of activation, nearly twice as much in all cases (Fig. 10(a)–(b)). Interestingly, however, if the same comparison is made while observing only the number of active cycles, not the total activation, there is less difference between the two. Trials with $F_{ww} = 100$ do not maintain all active cycles for much higher inhibition than those with $F_{ww} = 80$, and in trials with increased inhibition, the increased excitatory scaling factor only serves to keep a few more circuits active (Fig. 10(c)–(d)). The largest difference can be observed with mid-range inhibition ($5 \leq F_{wi} + F_{iw} \leq 10$), where trials with $F_{ww} = 100$ maintain up to 15% more active cycles than trials with $F_{ww} = 80$.

This exploration of the parameter space of this model gives a general picture of the reverberation of a single state in the workspace. There is a balance between excitation and inhibition—while strong reverberation is desired for a state of activation, favouring strong excitation, a mechanism is required to allow a new state to take over from an old one, highlighting the need for inhibition. In order to maintain a state of activation in times of the order of hundreds of milliseconds, a combined inhibitory factor less than 10 seems ideal—beyond this there is a quick drop off in the reverberation of a state.

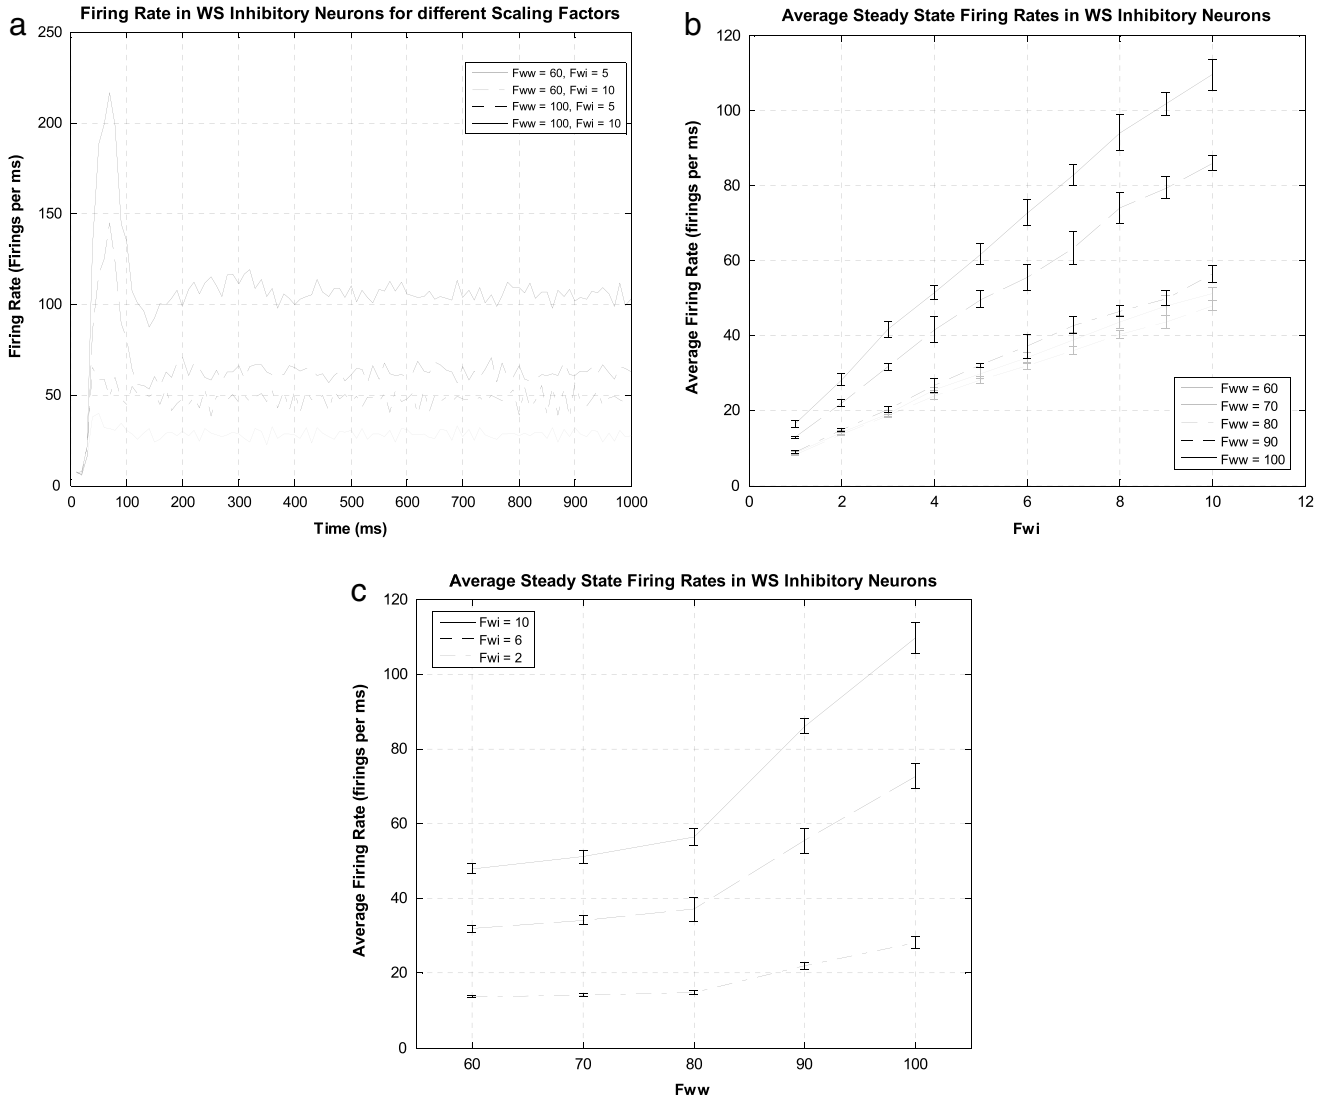

Fig. 8. (a) Progression of the firing rate in the workspace inhibitory population with time (as averaged in 10 ms windows of time) in a few typical model runs with $F_{iw} = 0$, and varying F_{wi} and F_{ww}. (b) Plots of the average firing rate at the steady state as changed by F_{wi}, for different F_{ww}. (c) Plots of the average firing rate at the steady state as changed by F_{ww}, for different F_{wi}.

3.4. Dynamics of multiple stimuli

All of these experiments have focused on the effect of varying scaling factors on the dynamics of a single stimulus. Although reverberation of a state of activation is a key predicted feature of the GNW model, it is also expected that such a structure is sensitive to new states, allowing them to enter the workspace and begin their own reverberation and broadcast. A new state of activation competes with previously lingering states, each trying to inhibit the other while sustaining its own reverberation. In order to study the effect of parameter choice in this particular model on this kind of competition, the interaction between different states of activation entering the workspace at different times is important.

Many of the effects of varying scaling factors on the dynamics of multiple states of activation are similar to the effects on a single state. The previous observations with regard to F_{ww} and F_{wi} ($F_{iw} = 0$) can be extended to multiple different states of activation, because the cycles themselves do not interact without inhibitory influence, simply by considering states with twice the amount of activation. Even with a non-zero choice of both inhibitory parameters, activation from two stimuli presented at or very near the same time (two pulses, each delivered to 37 different cycles of neurons with no overlap) inhibits each other in exactly the same way activation from a single stimulus with twice as much activation (a single pulse delivered to 74 different cycles of neurons) decays.

As previously mentioned, different sizes and forms of states of activation (for example, states with 74 active cycles instead of 37) in the workspace represent a significant avenue of future investigation. Here we limit the scope of investigation to temporally distinct stimuli, presented to the workspace greater than 100 ms apart, with non-zero choices of all scaling factors, such that the two states can interact due to the random, diffuse inhibition, which is not restricted to particular cycles. Some examples of these trials are shown in Figs. 11 and 12.

With non-zero F_{iw}, there is still a similar balance between excitation and inhibition as in the trials with a single presented stimulus. The inhibition still comprises roughly equal contributions from F_{iw} and F_{wi} with slightly greater contribution from a change in F_{iw} (Fig. 11(a)–(d)). The total amount of activation in the workspace which can be sustained over longer periods of time (more than 300 ms) for a given amount of inhibition seems to be conserved—the total activation caused by both stimuli as they reach a steady state is roughly equal to the amount of activation observed at the

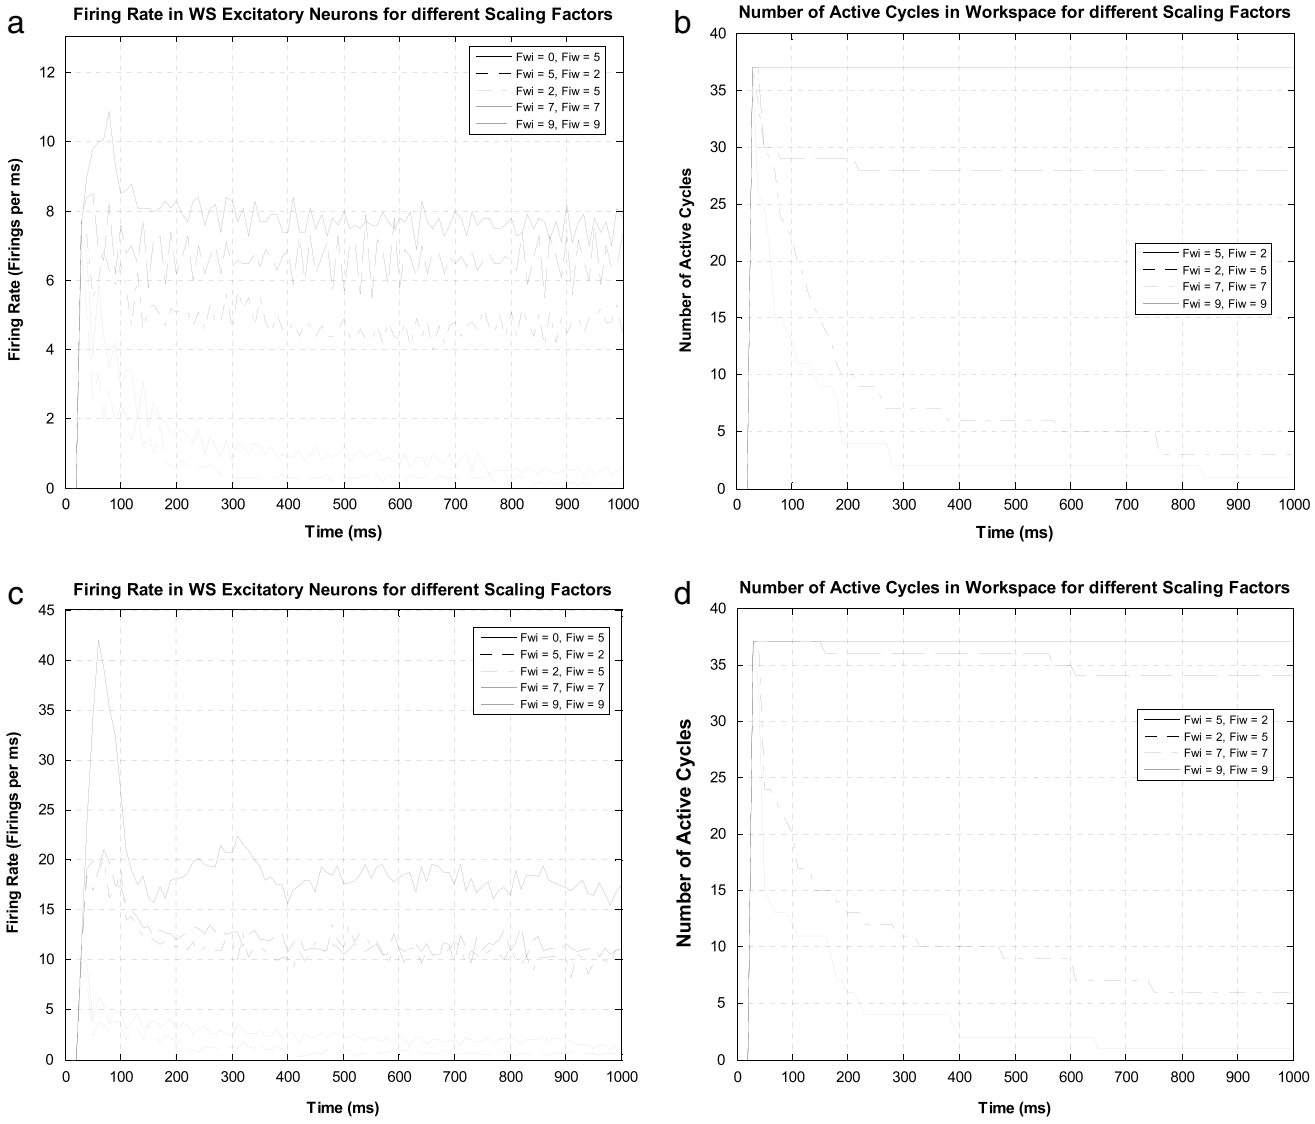

Fig. 9. (a)–(b) Some representative trial firing rate plots for $F_{ww} = 80$ and 100, with varying F_{wi} and F_{iw}. There is increased inhibition in the $F_{wi} = 2$, $F_{iw} = 5$ case relative to the $F_{wi} = 5$, $F_{iw} = 2$ case. (c)–(d) Corresponding plots of the number of active cycles in the workspace as time progresses. In the $F_{wi} = 0$ case (not shown) all stimulated cycles remained active for the full time.

steady state for trials with the same scaling factors and only a single initial stimulus (Fig. 11(c)–(d)).

It is also interesting to note that the activation in the workspace a few hundred milliseconds after presentation of the second stimulus is often split evenly between the two different stimuli (Fig. 11(c)–(d)). This demonstrates that a novel state of activation entering the workspace is inhibited just as much from an older reverberating state as the older reverberating state is inhibited by the incoming novel state of activation. This seems to hold true even though one state of activation can strongly dominate the other for short periods of time (Fig. 12(c)–(f)). In such cases, it is not obvious from parameter choice alone which state will dominate (Fig. 12(a)–(f)).

These latter cases (Fig. 12) resemble the predicted competition of a GNW, but only appear with high inhibitory parameter choices—in the range where significant reverberation of a single state is impossible. This suggests a combined inhibitory parameter choice around 10, the boundary where the opposing two effects of strong reverberation and competition between states meet, for future experimentation with this kind of workspace model. However, even with this restricted choice of scaling factors, the observed behaviour of the model is less than ideal. There is limited reverberation of a state over the time span of a few hundred milliseconds (more than half the cycles of the state often cease activity), and the competition between different states in the workspace consists more of finding an equilibrium between the two than the predicted 'winner-takes-all'.

These issues can be partially addressed through a more thorough look at the question 'What is a state of activation?', but a simpler effect can be tested which may account for the differences between a predicted workspace and the current model: driven stimulation. Though it is certainly possible a stimulus is presented to the workspace once, it seems more likely that a stimulus is presented over a span of time, as regularly timed pulses, for example. In this way, a distinction can be made between a strongly driven stimulus (presented many times or over a longer period of time) and a weakly driven stimulus. A driven stimulus may promote strong reverberation over the span of time in which it is presented. Once it ceases to be driven, reverberation continues but tails off with time so that a new driven stimulus presented to the workspace can compete and almost completely take over the state of the workspace.

Rather than presenting a single pulse at $t = 20$ ms and a different pulse at $t = 300$ ms, a series of pulses is used, 10 ms apart from

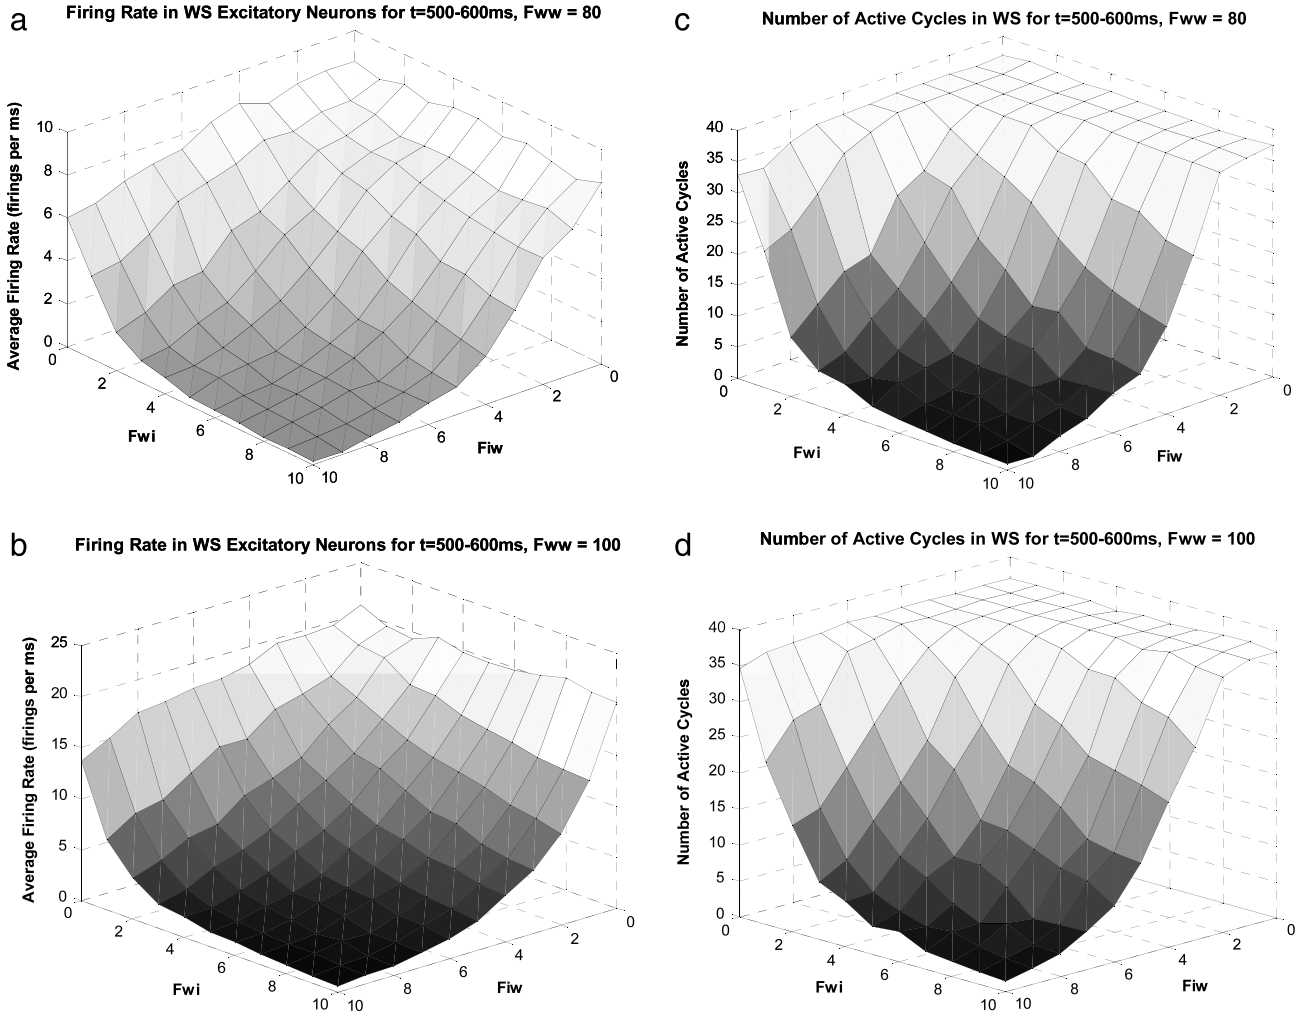

Fig. 10. (a)–(b) Surface plots showing the relative change in workspace activation (measured as the average firing rate of excitatory neurons for time $500 \leq t \leq 600$ ms) over the full range of inhibitory parameters tested for $F_{ww} = 80$ and $F_{ww} = 100$. (c)–(d) The number of cycles still actively reverberating in the measured timeframe for $F_{ww} = 80$ and $F_{ww} = 100$. Interestingly, though trials with $F_{ww} = 100$ show double the activation of trials with $F_{ww} = 80$, there are only consistently more active cycles in the $F_{ww} = 100$ case at the mid-range of inhibition ($5 \leq F_{wi} + F_{iw} \leq 10$).

each other for a duration of 160 ms.[2] As predicted, driven stimulation promotes stronger reverberation and more clear competition. Fig. 13 shows a typical trial with the same scaling factors as in Fig. 12, but this time with driven stimuli. Unlike the single pulse stimulus trials, the second driven stimulus consistently 'wins' the competition, and both states of activation consistently reverberate with more than 27 active circuits (roughly 75% of the full state) while driven.

4. A more complete model

The discussion thus far has revolved around the workspace neurons: broadcast and reverberation, how the balance between excitation and inhibition facilitates these phenomena, how modifying the scaling factors affect the excitation and inhibition, and the important role of driven stimulation. While stimuli were simulated by a direct injection of current to the appropriate set of neurons, the model is designed for these to come instead from functionally distinct neuronal clusters—specialist areas—which are attached to the workspace. Specialist areas in this model are groups of neurons with many diffuse internal connections, representing the close-range connections of the model, in contrast to the long-range connections of the workspace neurons.

Each specialist area goes through a period of training before an experiment, during which that group of neurons learns to associate a particular input state of activation with a different output state of activation. Several specialist areas, each trained in this way with different inputs and outputs, and connected to each other via the workspace discussed above, represent a model which can go through alternating periods of broadcast and competition, as theorized might occur in a GNW by Dehaene et al. (2003).

Here we present the construction of such specialist areas in the context of this model and produce similar results to the comparable model of Shanahan (2008). The goal of this is to demonstrate the similarity of these models despite the modifications, as well as to highlight any differences in behaviour in the context of these same modifications.

4.1. Specialist structure and connectivity

Each specialist area can be divided into three groups: input, output, and generic internal neurons (Fig. 14). The input and

[2] The exact choices involved with a 'driven stimulus', such as the time between pulses and the overall time of the series of pulses, were varied to test for robustness. As one would expect, longer duration of the series of pulses and shorter times between pulses have a stronger effect. The particular values used here were chosen as sufficient to convey the overall point—a driven stimulus promotes stronger reverberation and clear-cut competition.

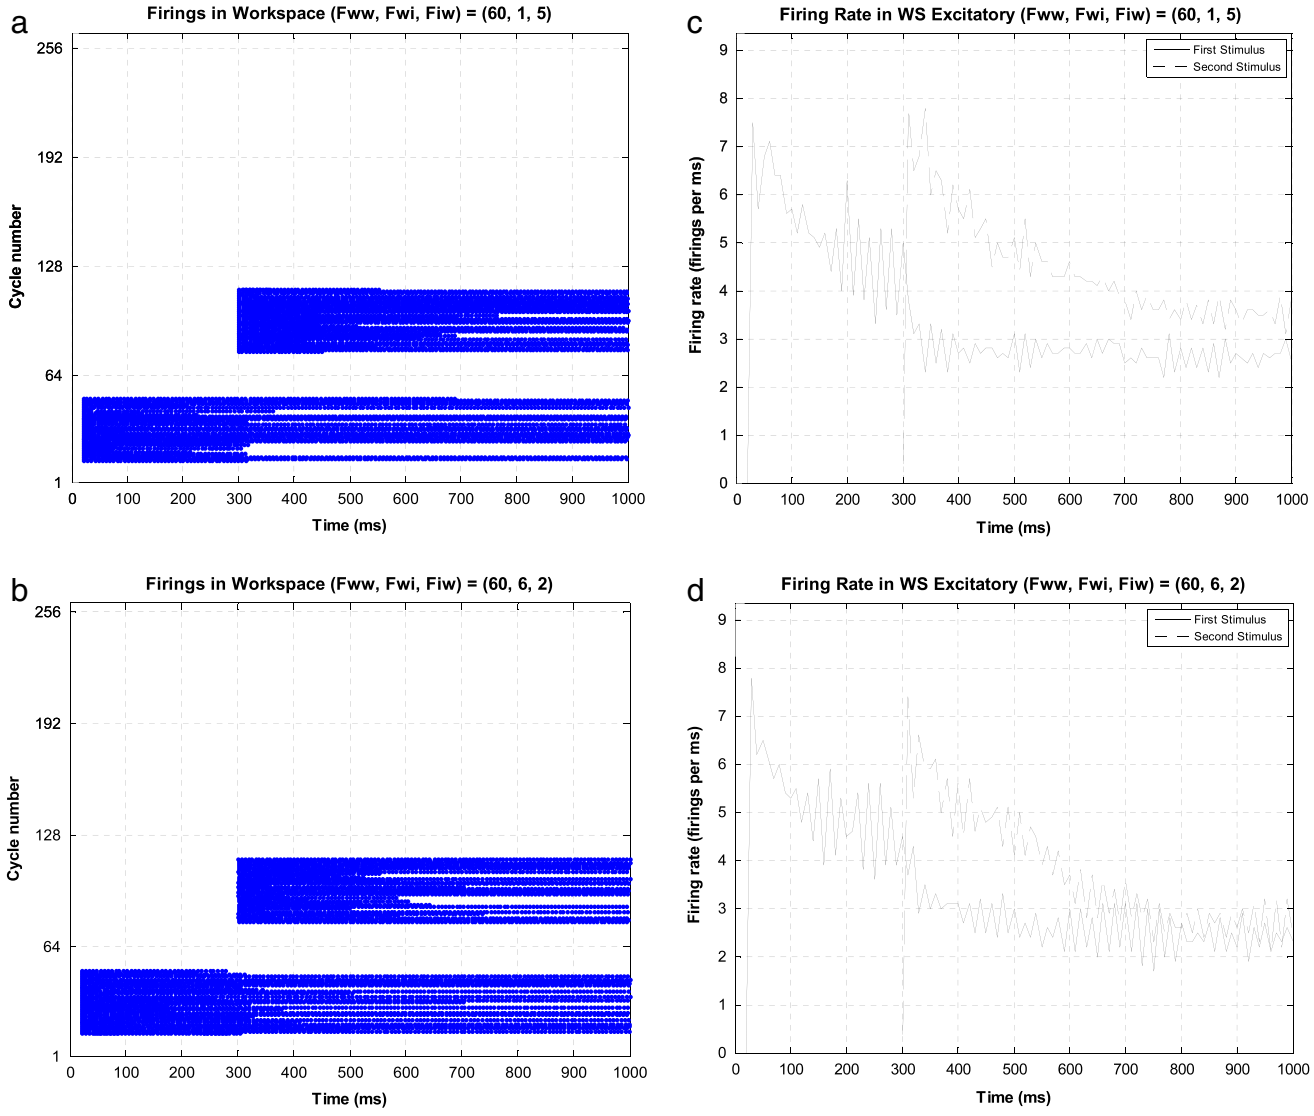

Fig. 11. (a)–(b) Raster plots for two trials with multiple stimuli that characteristically progress similarly over time. Inhibitory scaling factors, F_{wi} and F_{iw}, are roughly reversed, with $(F_{wi}, F_{iw}) = (1, 5)$ in one and (6, 2) in the other. Increased F_{iw} typically inhibits slightly more than increased F_{wi}. (c)–(d) Corresponding firing rate plots for the two trials shown. After some time it is difficult to distinguish which state of activation is 'winning'.

output areas are groups of 256 excitatory neurons, each of which represents a mapping of possible workspace states. Each neuron in these areas is numbered for organizational purposes (1–256), and neurons in one area map directly to neurons with the same number in the other areas. As explained in Section 2, the cycles of the workspace are similarly numbered; they also share this mapping.

Specialist areas are each designated a length of 384 workspace neurons (twice the short-distance parameter α) which are 'spatially close' (see Fig. 1 for a reminder of the overall arrangement of specialist areas and workspace). There are one-to-one connections from nearby neurons of the workspace to the input neurons of the specialist areas, and from the output neurons of the specialists to the nearby neurons of the workspace. These one-to-one connections are organized by the cycle numbers of the workspace and the appropriately mapped neuron numbers of the specialist input and output neurons. For example, activity in neuron number two of the specialist input is stimulated by activity of a nearby neuron in the workspace belonging to cycle number two, ensuring that the specialist area 'sees' states of activation from the workspace. Firings of neuron number two in a specialist output area will stimulate firings in a nearby neuron in cycle number two of the workspace, thus attempting to pass a state of activation on to the workspace.

The group of generic internal neurons comprises 1126 generic excitatory neurons and 410 inhibitory neurons.[3] Connections between the generic neurons, the input neurons and the output neurons, as well as internally within these subdivisions, are diffuse and their connection weights are initially drawn from a uniform random distribution in the interval [0, 1]. These connections are then strengthened or weakened according to a spike timing-dependent plasticity (STDP) rule during a training period prior to the full trial.

4.2. Specialist training

All specialist areas go through a pre-trial training period. The purpose of this is to strengthen and weaken the internal

[3] Note that, excluding the laterally inhibiting groups of neurons, the specialist areas contain the standard proportion of excitatory to inhibitory neurons in a real mammalian cortex, 4:1 (Binzegger et al., 2004; Gabbott & Somogyi, 1986). This ratio is maintained throughout the specialist areas and workspace.

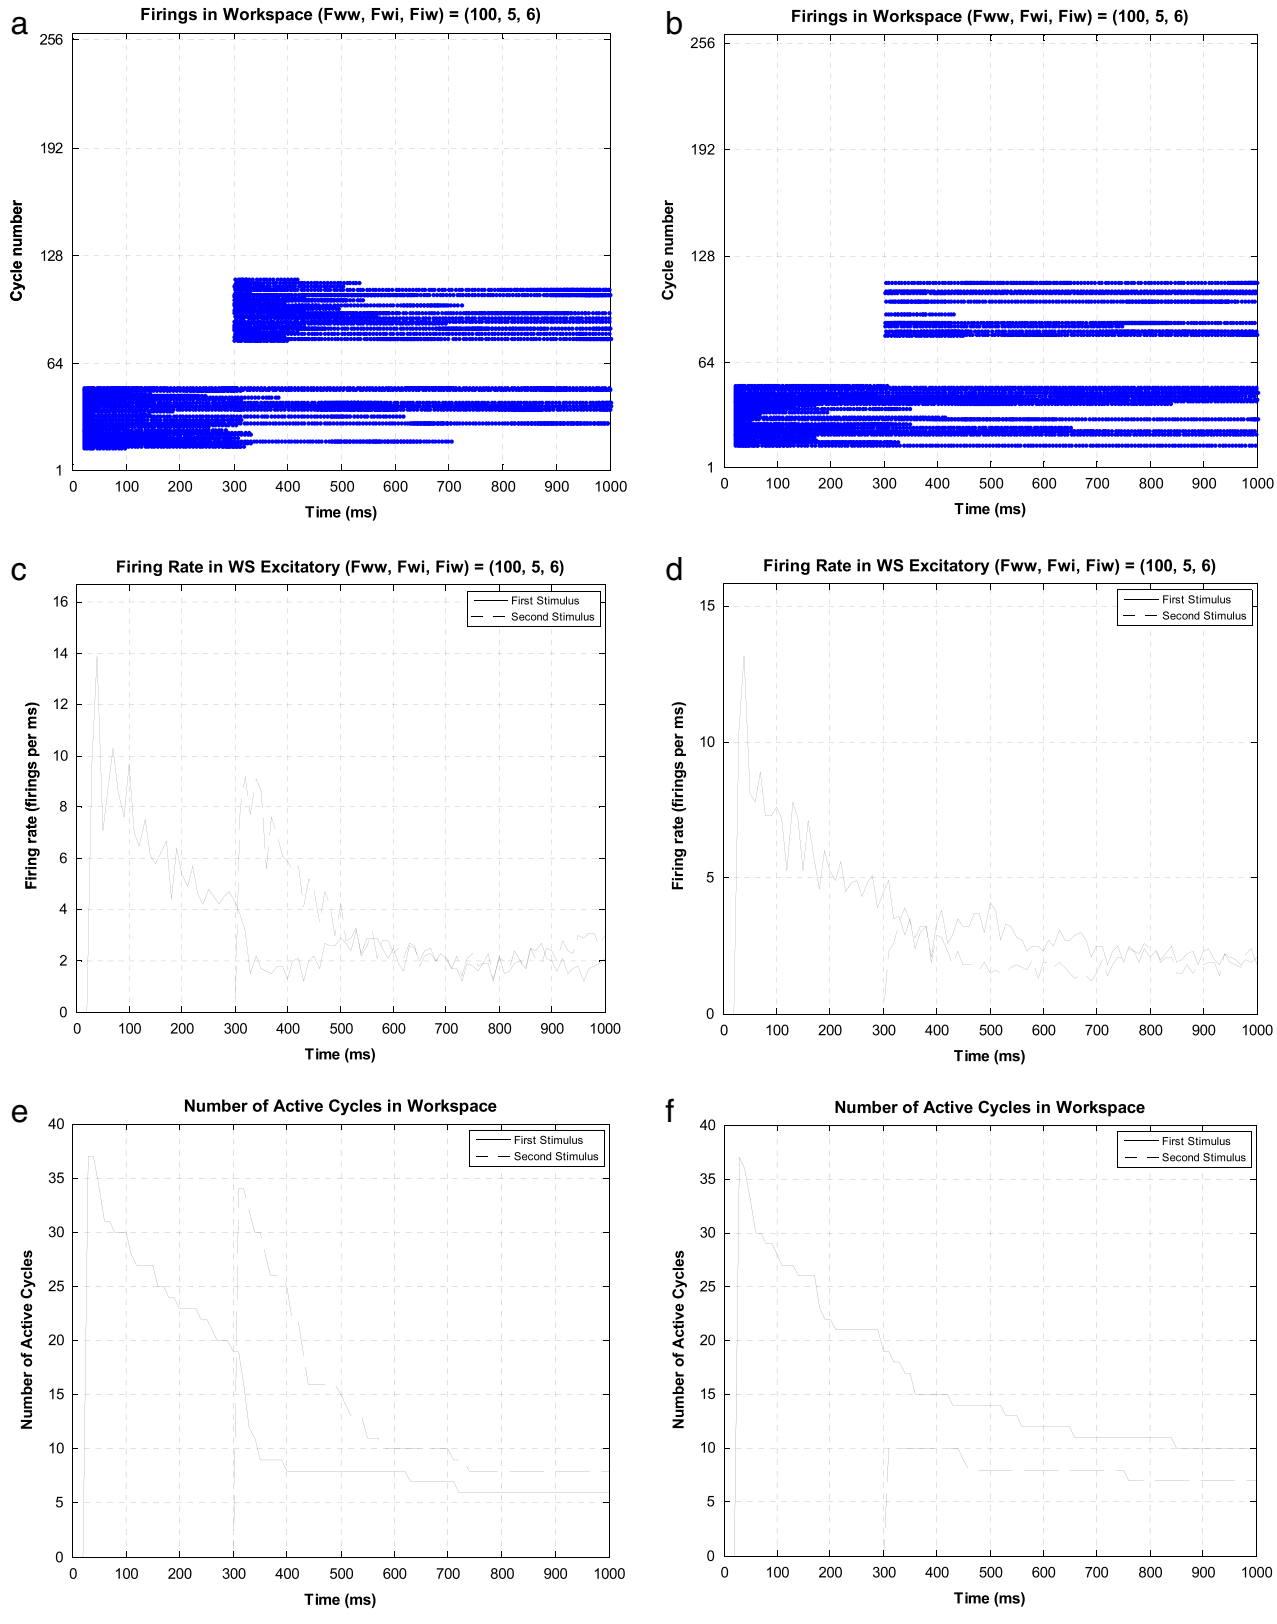

Fig. 12. Two trials in which identical scaling factors were chosen and very different results are observed. (a)–(b) Raster plots for the two trials. In one there is clearly strong activation from the second stimulus while in the other the second stimulus is almost completely blocked out. (c)–(d) Firing rate plots of the two trials. Despite only weak activation of the second stimulus in one trial, there is still a nearly even balance of activation from both stimuli. (e)–(f) Plots of the active cycles of each state for both trials. Neither trial shows strong reverberation of a state over hundreds of millisecond timescales, nor is there a clear 'winner' of the competition between the two, despite nearly complete initial blocking out of the second stimulus in one trial.

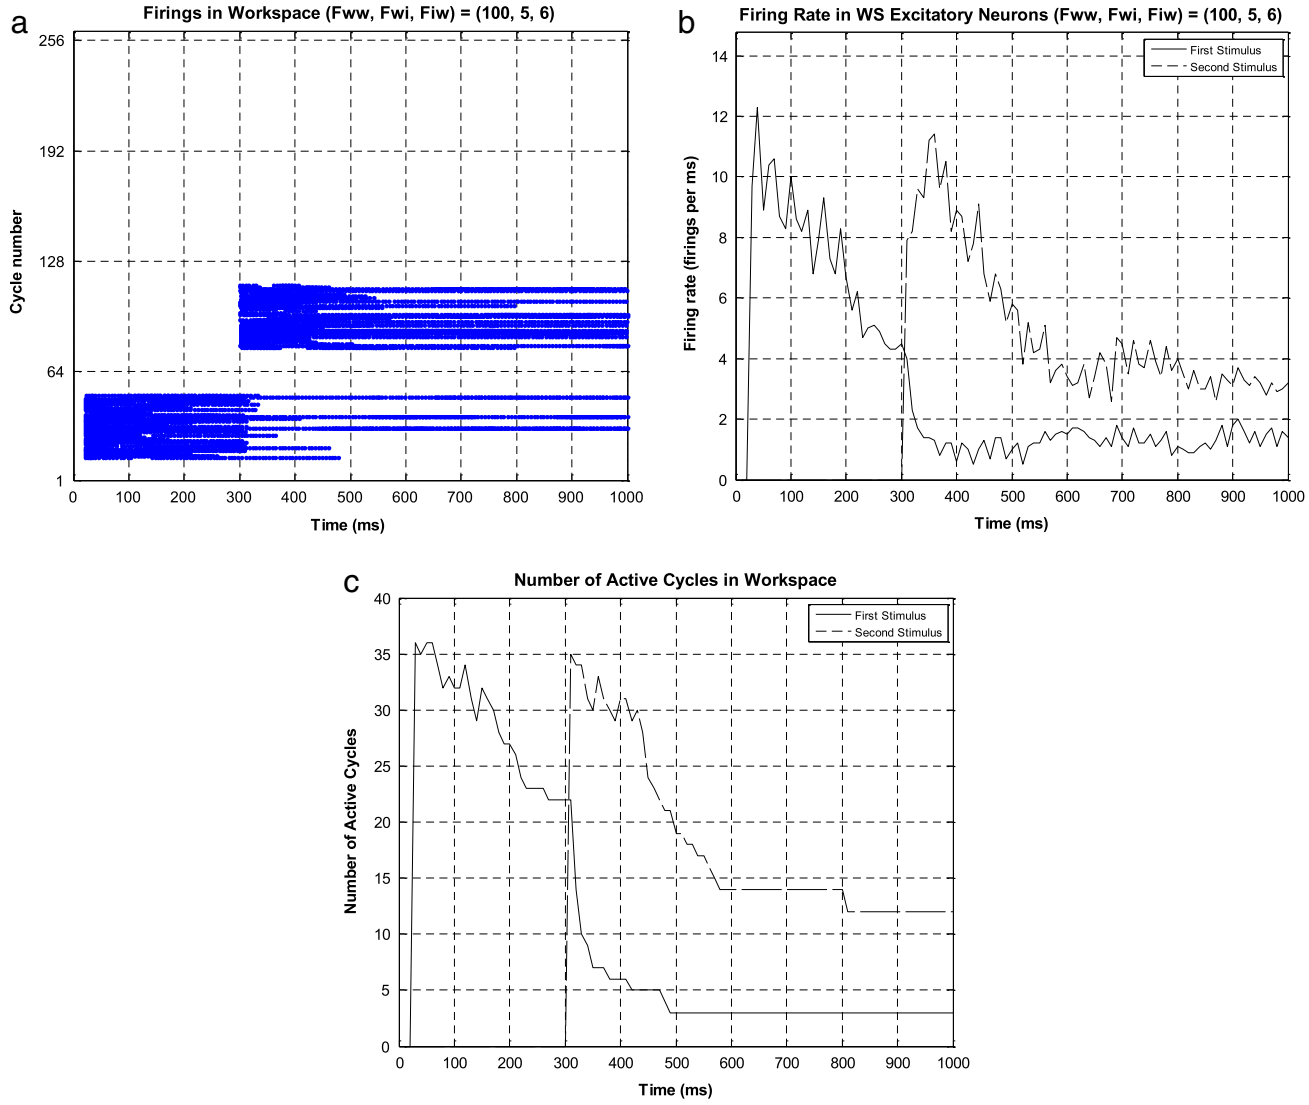

Fig. 13. (a)–(c) Raster, firing rate and active cycle plots for a representative trial with driven stimulation. While being driven (a duration of 160 ms from the start of stimulation) each state shows significantly improved reverberation. In addition, when the second state of activation, driven by repeated stimulation, enters the workspace it strongly inhibits the un-driven firststate of activation, creating competition with a much clearer 'winner'.

connections of the specialist area such that it has some overall functionality. In this model, 'functionality' takes the form of transforming a state of activation—given a particular state of activation of the input neurons, a different state of activation fires in the output neurons. These states of activation are identical in organization to the states of activation of the workspace in the previous section (a subset of one of four quadrants simultaneously firing—see Section 2.2).

During the training period, the specialist areas are disconnected from all other structures and are each alternately presented with a stimulus in the input area, followed by a different stimulus in the output area. This is repeated while a form of STDP is applied (Song & Abbott, 2001). After several iterations, the internal neural connections which lead the specialist area to fire the response state in the output area after the appropriate state appears in the input area are strengthened, while other connections are weakened.

The neurons for an input state are excited by four 10 mA pulses, 5 ms apart from each other, followed by the same treatment for a response in the output, after a delay of 40 ms. Each training lasts for 200 ms of model time, and this process is performed three times for each specialist area before every trial.

Even without changing the input and output states or generic set of neurons, different forms of training are possible. Rather than four pulse input, delay, and four pulse output, many other choices exist and do indeed provide slightly different dynamics. One such choice attempted in this study was alternating from single pulse input to single pulse output, with a delay between each repetition of this. The overall behaviour of the model was found to be robust, but the choices involved in the STDP training of the model are numerous (see Caporale & Dan, 2008, for a recent review of many of the currently studied STDP rules observed in real neurons). How the so-called specialist clusters of neurons reach a state of having this functionality is easily the topic of a much broader study—here we assume that it has reached the state of having such functionality and use the aforementioned STDP rule just as a method of emulating this process in a less engineered fashion than setting the connection weights as a parameter.

Usually the different specialists have different inputs to which they respond and can together be trained to give a series of activation states if the appropriate response states reverberate in the workspace, thus becoming inputs for other specialist areas. Each specialist is given a unique response state, so it is possible to tell at any one time where the state currently in the workspace originated.

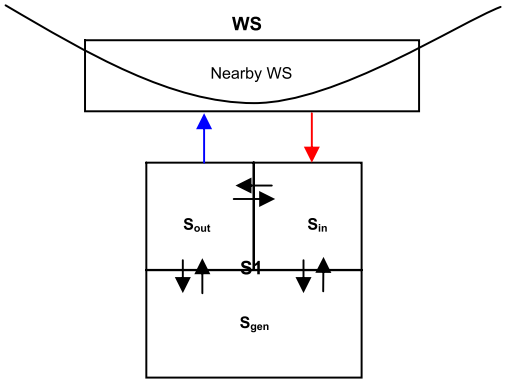

Fig. 14. Connections to, from and within a specialist area. A specialist area with a competing neighbour would have an additional group of laterally inhibiting neurons receiving its input from the S_{out} group of neurons (not shown). Connections between the workspace and S_{in} or S_{out} are one to one. Within the main body of the specialist area, connections are initially diffuse and random and are modified before the main trial using a form of STDP and test pulses.

4.3. Competing specialists

In addition to the competition within the workspace of simultaneous states of activation, direct competition between two neighbouring specialists for access to the workspace can be modelled. Two competing specialists are trained to react to the same input, but with different responses, and in all trials presented here the two specialists share the same stretch of local workspace neurons (though this is not necessary). An additional group of inhibitory neurons (one fifth the size of a specialist area) is connected randomly and diffusely from the output group of neurons of one specialist to the entirety of the other specialist. This is appropriately mirrored from the other specialist with a second group of laterally inhibiting neurons. Thus, when there is activity in the output group of neurons of either specialist, indicating an attempt to influence the workspace with a new state, the other specialist receives a diffuse inhibitory signal, attempting to shut down any competing attempt to access the workspace.

4.4. Sample trials

Consider a single run with this model. The first specialist area is trained to associate an input state *A* with an output state *B*, the second specialist area associates an input state *B* with an output state *C*, and the third specialist area associates an input state *B* with an output state *D*. The experiment begins with the injection of a pulse representing state *A* directly into the workspace. With sufficient reverberation, the first specialist area picks this up and responds by attempting to influence the workspace with state B. If this reverberates and is broadcast to the second and third specialist areas, they will compete to place states *C* or *D*, respectively, into the workspace.

Fig. 15 shows several representative trials with this behaviour using scaling factors $F_{ww} = 100$, $F_{wi} = 5$, $F_{iw} = 6$, with some sample trials taken from Shanahan (2008) for comparison. In agreement with the idea that these two different models both contain an underlying GNW structure, the same qualities of competition between states, both within the GNW and directly between specialist areas, reverberation, and broadcast appear in both models. There are some minor differences between the two (e.g. the appearance of the eventual 'loser' early in competition in the presented model as opposed to little trace at all in Shanahan (2008)) but these are probably due to structural choices without a weighty interpretation.

In the case shown, the response is mostly bistable between states *C* and *D* (there are some traces of a 'loser' state which avoid inhibition due to the statistical nature of the inhibition in the model). However, this is sensitive to parameter choices—without sufficient excitation, no state is stable, and without sufficient inhibition, multiple states do not compete and the workspace acts like a single 'on' switch. Note that additional specialist areas and competitions could be included in the model, in which case there would be multistability for the workspace as a whole.

5. Discussion

With both the GNW-only model and the larger model with specialist areas, we aim to present a possible instantiation of the GNW concept originally proposed by Dehaene et al.. We have shown that three major components of the GNW theory (broadcast, reverberation, and competition between different states of activation) can come about in a model different from those already proposed. With this in mind, it seems that there is a potentially large unexplored space of possible instantiations of a GNW. While this model attempts to remove some of the arbitrary choices made in the previous work, it is clear that other choices can be made, perhaps including a biologically accurate instantiation of a GNW.

In the spirit of further removing constraints from the model, certain choices made here can be analysed with an eye to future work. In particular, the use of stochastic wiring with isolated cycles in the workspace, though it was chosen as an alternative to an arbitrarily engineered network, could be changed. Perhaps the right balance of inhibition and excitation could be used to overcome blurring of states of activation without the use of completely topographic connections (Wang, 2001). Alternatively, topographic connections of the kind used both here and in Shanahan (2008) may be justified as emergent through the correct application of a synaptic learning rule like STDP in a model like that used in Izhikevich et al. (2004).

Here we choose to separate the neurons with long-range connections into a distinct group, the workspace. However, this is not necessary. A more biologically plausible model might show the distinct GNW qualities described here, but amongst clusters of neurons where no preference is given to the neurons with connections between clusters (long-range connections). An interesting, and potentially vast, avenue of study lies in determining the minimum rules and restrictions on the construction of connections within a set of neurons such that there is activity resembling a GNW within that set.

It is also not yet clear what role the particular neuron model used plays in the overall network dynamics. The model used here (Izhikevich, 2003) was chosen for its balance between diverse, realistic neural dynamics and computational tractability. It would be an interesting exercise to confirm that these results remain largely unchanged for different neuron models, such as a simple integrate-and-fire model. Furthermore, any differences in dynamics could be analysed in order to judge whether there are interesting effects at work, or simply artifacts of the model. This could potentially lead to a better understanding of the limits or pitfalls of using one model rather than another, and its usefulness would extend beyond the scope of GNW models.

The effectiveness of this model as a GNW depends on the choice of certain parameters, and we have shown how the behaviour changes with regard to one particular parameter (the scaling factor). Though the effects of modifying this parameter can at times be obvious (e.g. increasing an excitatory scaling factor leads to more activation), the range of parameters for which GNW-like behaviour is observed is not. Here we have shown, in the context of this model, how modifying the scaling factors of connections within the GNW can lead to regimes of overly strong reverberation (such that inhibition does not stop activation from increasing as

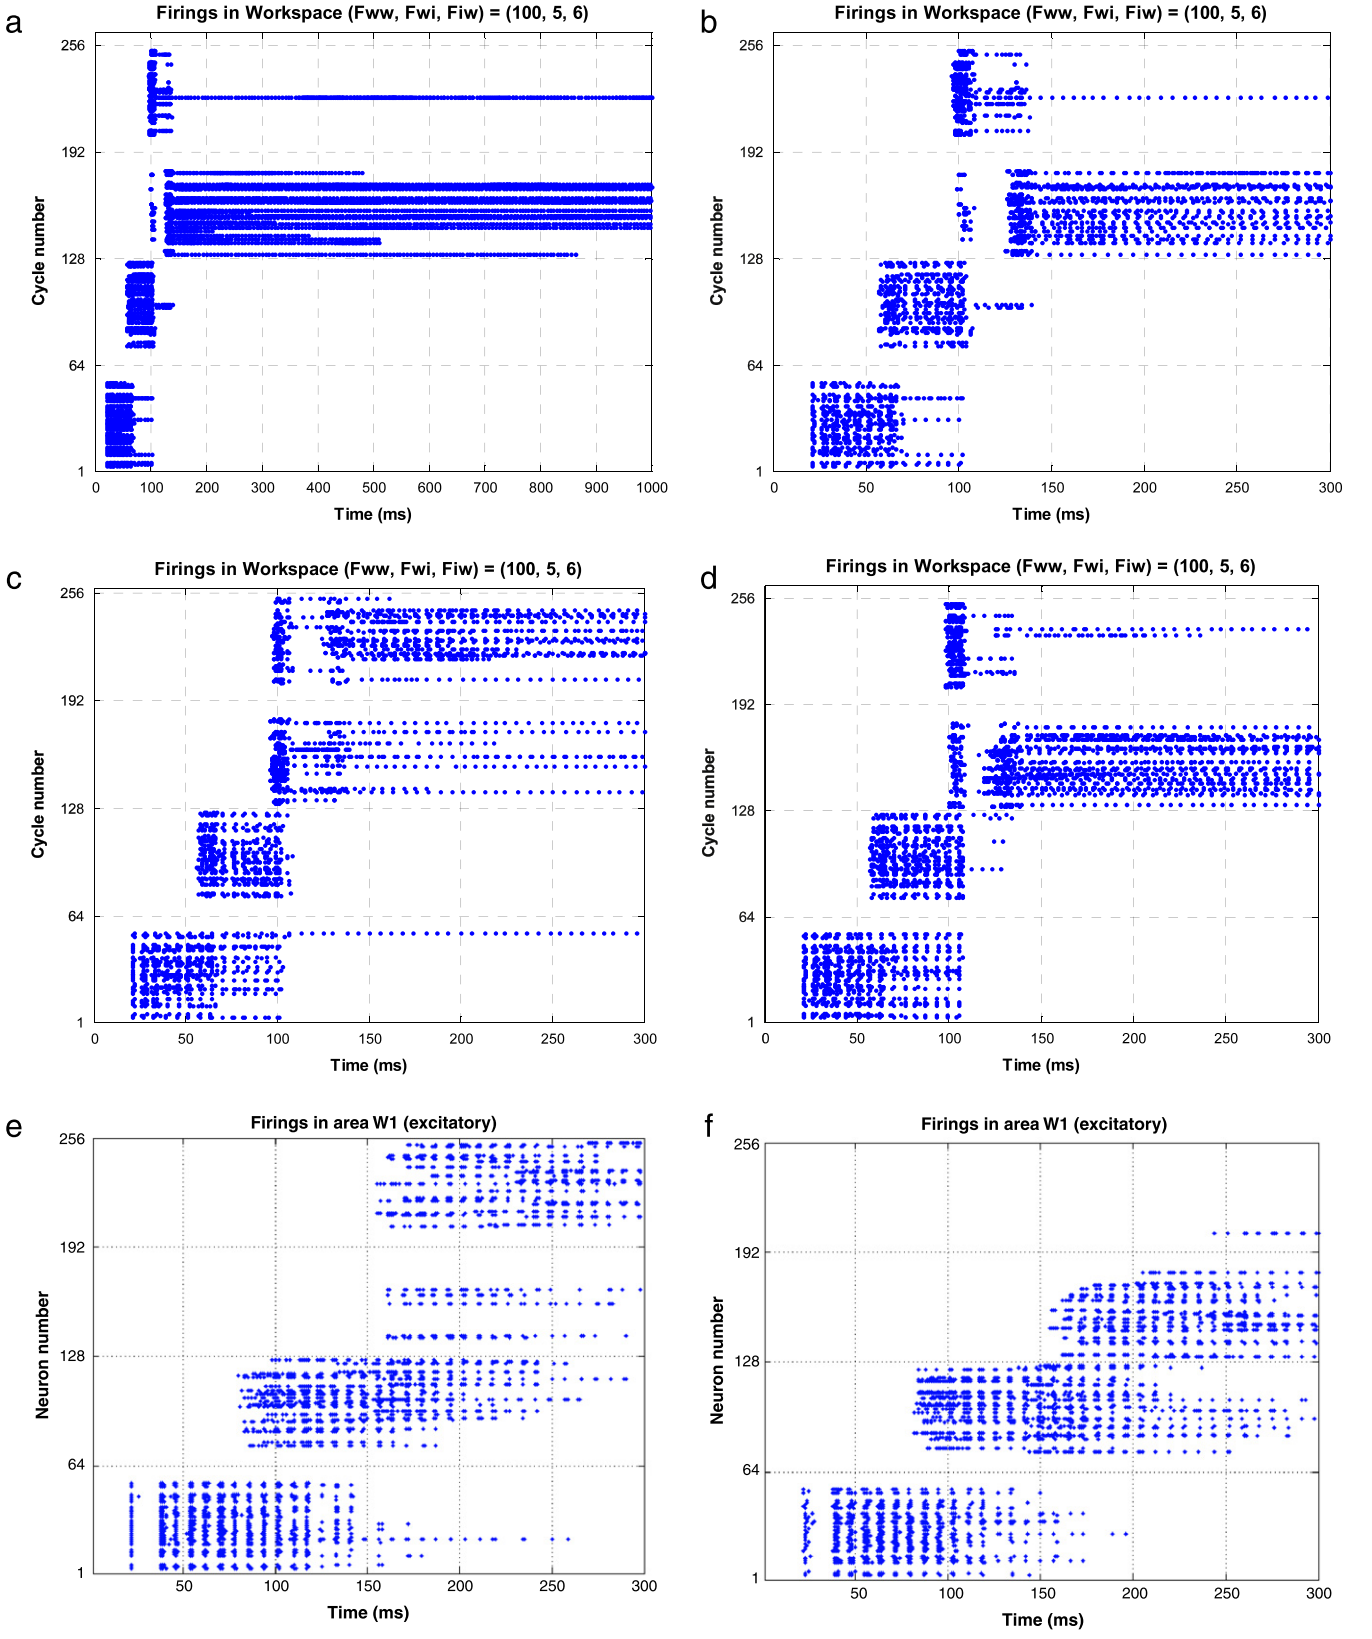

Fig. 15. (a)–(b) Raster plots for a single trial shown over the whole 1000 ms of the trial and just for the first 300 ms as in Shanahan (2008). (c)–(d) Raster plots for two different trials with the same scaling factors, but with different competition outcomes. (e)–(f) Raster plots for two different trials with different competition outcomes using the Shanahan (2008) model for comparison (from Shanahan (2008)).

more states of activation enter the workspace) or quick decay of states, such that they never have a chance to be broadcast around the entire workspace.

Other parameters, such as conduction delays between neurons and the connectivity of neurons, play a similarly important role in modulating this behaviour. The choice to have tightly distributed

delays could potentially be lifted with few consequences to the overall dynamics, though our own preliminary work in this area has shown that delays which are too low can lead to network dynamics, such as reverberation, being disrupted by neural dynamics, such as individual neural recovery times. Work towards exploring more general principles governing the overall changes in behaviour due to changes in these parameters within the space of possible GNW instantiations is warranted.

Acknowledgements

Thanks to Stanislas Dehaene, for his insight and hard work with GNW theory, and Eugene Izhikevich, for his generosity in making his MATLAB code publicly available. Thanks also to Yang Xu and Vidu Vinayagamoorthy for their complimentary work on STDP.

References

Baars, B. J. (1988). *A cognitive theory of consciousness*. Cambridge: Cambridge University Press.

Baars, B. J. (2002). The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Sciences*, *6*, 47–52.

Bentley, N. M., & Salinas, E. (2004). Bistability in oscillatory cortical modules. *Neurocomputing*, *58–60*, 769–774.

Binzegger, T., Douglas, R. J., & Martin, K. A. C. (2004). A quantitative map of the circuit of cat primary visual cortex. *The Journal of Neuroscience*, *24*(39), 8441–8453.

Caporale, N., & Dan, Y. (2008). Spike timing-dependent plasticity: a hebbian learning rule. *Annual Review of Neuroscience*, *31*, 25–46.

Dehaene, S., Kerszberg, M., & Changeux, J. P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences USA*, *95*, 14529–14534.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, *79*, 1–37.

Dehaene, S., Sergent, C., & Changeux, J. (2003). A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences USA*, *100*, 8520–8525.

Durstewitz, D., & Seamans, J. K. (2006). Beyond bistability: biophysics and temporal dynamics of working memory. *Neuroscience*, *139*, 119–133.

Gabbott, P. L. A., & Somogyi, P. (1986). Quantitative distribution of GABA-immunoreactive neurons in the visual cortex, area 17, of the cat. *Experimental Brain Research*, *61*(2), 323–331.

Izhikevich, E. M. (2003). Simple model of spiking neurons. *IEEE Transactions on Neural Networks*, *14*, 1569–1572.

Izhikevich, E. M. (2007). *Dynamical systems in neuroscience: the geometry of excitability and bursting*. MIT Press.

Izhikevich, E. M., & Edelman, G. M. (2008). Large-scale model of mammalian thalamocortical systems. *Proceedings of the National Academy of Sciences USA*, *105*, 3593–3598.

Izhikevich, E. M., Gally, J. A., & Edelman, G. M. (2004). Spike-timing dynamics of neuronal groups. *Cerebral Cortex*, *14*, 933–944.

Kalitzin, S., van Dijk, B. W., & Spekreijse, H. (2000). Self-organized dynamics in plastic neural networks: bistability and coherence. *Biological Cybernetics*, *83*, 139–150.

Lo, C.-C., & Wang, X.-J. (2006). Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nature Neuroscience*, *9*(7), 956–963.

Marzouki, Y., Grainger, J., et al. (2007). Exogenous spatial cueing modulates subliminal masked priming. *Acta Psycholica (Amsterdam)*, *126*(1), 34–45.

Shanahan, M. (2008). A spiking neuron model of cortical broadcast and competition, and supplementary note. *Consciousness and Cognition*, *17*(1), 288–306.

Song, S., & Abbott, L. F. (2001). Cortical development and remapping through spike timing-dependent plasticity. *Neuron*, *32*, 339–350.

Van den Bussche, E., Segers, G., & Reynvoet, B. (2008). Conscious and unconscious proportion effects in masked priming. *Consciousness and Cognition*, *17*(4), 1345–1358.

Wang, X.-J. (2001). Synaptic reverberation underlying mnemonic persistent activity. *Trends in Neurosciences*, *24*(8), 455–463.

Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, *36*, 955–968.